Photographer: Jayme Okerblom

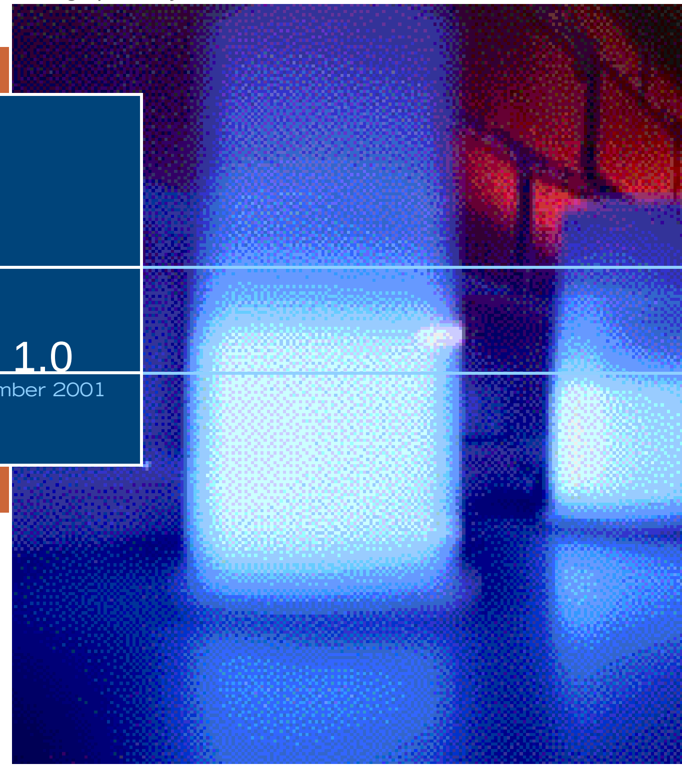

Luminosity

Spotlight on Emerging Technologies

Vol. 1.0
September 2001

Table of Contents

luminant

In This Issue of Luminosity

In this inaugural issue of Luminosity, we have two enlightening columns.

Luminant's EAI practice leader, Ankur Laroia, introduces us to the world of EAI in the first of his three part series on EAI, in the column *"Introduction to EAI"*. Ankur is a recognized expert in the implementation of EAI strategies and solutions.

In the column *"Optical Internet Infrastructure"*, Anupam Singal discusses the advances being made in Optical Internet technologies. He has over ten years of experience in designing and implementing broadband optical network solutions.

Enjoy
Your comments are always welcome.

Chetan Sharma
Luminosity Editor

Editor's Opinion: 3Ps of Pervasive Computing

Privacy, Personalization and Protection

As we tread toward the next step in computing, the three Ps of pervasive computing are becoming increasingly important. It's imperative to provide the following for all applications and services:

- **Privacy of user data**
- **Personalized user experience**
- **Protection of information assets**

But, we haven't seen the three components together just yet. In this column, we will discuss what these three elements mean to the future of commerce and Internet applications and services.

Privacy

Let's get to the most important issue first–privacy. Lately, a lot of companies, organizations, and politicians have been jousting for their stance on privacy. A battle is brewing between Passport (Microsoft) and Magic Carpet (AOL). There is talk of privacy legislation on various fronts–E911 and online–and organizations like EPIC (Electronic Privacy Information Center) have launched scathing attacks on Microsoft. So, besides appearing ethically and morally astute, what's at stake here? **Transactions/user.** The conventional wisdom is the more you know about as many customers as possible, the more dollars per user will flow toward you. The grab for consumer data is going to be an interesting battle in the coming years. AOL has five times more subscribers than MSN while 90 million PCs are projected to have Windows XP by 2002 (source: Business 2.0). All of this lays the groundwork for interesting times ahead.

Privacy is all about "trust." If someone masquerades as someone I should trust with my "information," then turns around and sells that information to the highest bidder, without my consent, it's WRONG – plain and simple. They can muddle around with verbiage all they want, but the fact remains it's wrong. Unfortunately, it is a common practice. The Internet is here, and collecting user information is absolutely essential to provide a meaningful user experience. Otherwise you are bound to get junk more often than not, and people don't have time to sort through it. (More on personalization later.) You might find it hard to believe, but an amazing amount of information can be collected, stored, and mined to build a pretty good user profile. If I get value for what information I provide, hey, I'm all for it – for example, purchasing airline tickets and books online with minimal clicks, or viewing customized news and sports scores, and so forth. Who has the time to enter addresses and credit card info again and again? God bless Amazon and Expedia for making our lives easier.

However, problems arise if these companies go beyond using the information to improve my user experience, and then do underhand dealings with spammers. That's a betrayal of my trust and needs to stop.

Let's take a look at some of the technology solutions being introduced to address privacy.

P3P

Good progress has been made on the Platform for Privacy Preferences Project (P3P) at W3C, which allows automatic, computerized reading of a Web site's privacy policy by browsers. There are two key components:

- The client side that allows P3P clients to automatically fetch and read P3P privacy policies on Web sites
- The server side (Web site) that allows Web sites to translate their human-readable privacy practices into a standard, machine-readable format that can be retrieved and read by browsers

However, the problems with P3P include the following:

- It's not a complete solution: We need more customized capability. Instead of generically stating, "I don't want mining of my data," you can say, "I don't want mining of my data from XYZ and there needs to be a way to track my private information from changing hands on demand."
- I need a way for my information to go into hibernation and have the ability to delete it from any web site, if that's what I choose to do. (Vendors could also consider destroying post-transaction personal data, depending on non-repudiation requirements.) P3P is largely unknown to consumers and businesses alike. There is no automatic way to equip legacy browsers with P3P capability.
- It's a chicken-and-egg dilemma: Companies won't make the translations until customers have the tools and demand P3P capability on the server side, but consumers won't bother to download and configure tools just to interpret a privacy policy that they don't read anyway.
- Also, P3P categorizes the types of information handed over by the user in the following ways:
  - ◊ The purpose for which it is collected,
  - ◊ The recipients of the information and
  - ◊ The duration of the information's retention.

These categories can be misused. So, if the purpose category is <current/> or <stated-purpose/>, the user has to dig for more details of the vendor's meaning.

- Very complex user scenarios exist with the class of devices that most need privacy – wireless phones. This is especially important in location-based data sharing scenarios.
- The most serious problem is of course the inability to enforce privacy. We need to devise mechanisms so the policy agreement between consumer and vendor is legally binding.

How many users do you think will actually change the browser default?

When users visit a site that uses P3P, they can click on the privacy icon in their browser to "privacy check" the site (Figure 1.) This brings up a window that explains any areas where a site's policy conflicts with a user's preferences. Users can also use this window to jump directly to a site's privacy policy, as well as to see whether the site has a privacy seal.

However, P3P is a good first step. A Web-agent-based approach is best for a privacy handshake, but we need to keep working with organizations like EPIC to make progress. Due to the pressures from such groups, Microsoft has now limited the information required to use Passport. Additionally, we need strict enforcement of privacy laws, so there is no doubt in anyone's mind about the repercussions. Unfortunately, legislative bodies move too slowly for the information age we live in. We need to design legislation keeping the next decade in mind. Then, we might have a faint chance of getting it right.

Figure 1

P3P1 Policy is Acceptable

P3P2 Policy is Not Acceptable

IDcide's Privacy Companion

XNS

Last year, XNSORG (Extensible Name Service Public Trust Organization) took the P3P concept to the next level by introducing a new open protocol and open-source platform, XNS. Based on XML and using Web agents, XNS is designed as a global solution for automatically exchanging XML data between two devices with privacy, security, and synchronization controls. XNS uses the concept of XNS business agent to first negotiate a legal contract between user and the business before doing the P3P step. Take a look at some interesting work done by OneName in this area.

In addition there are anonymity and pseudo-anonymity tools, encryption tools, filters, identity management tools and other devices available. However, they can't be used as generic solutions to address privacy concerns, because the average consumer won't go through the trouble of learning the nuances. For any solution to gain wide spread adoption, it needs to be part of the browser.

Regulations

There are current US regulations that protect consumers' financial (the Gramm-Leach-Bliley Act) and medical (Health Insurance Portability and Accountability Act) information from being sold to third parties for purposes of telemarketing or other direct marketing. The Children's Online Privacy Protection Act, enacted in law in 1998, requires that Web sites visited by children under age 13 post a privacy policy detailing any personally identifiable information collected from those children.

In Europe, most of the EU member states have implemented the EU 1995 Data Protection Directive that seeks to protect consumer data. There are also some laws in place for data that crosses borders. A "safe harbor" arrangement exists between the US and EU. It declares that personal data about EU citizens may be transferred to the US only if adequate protection is provided, such as obtaining consent for any sensitive information used for purposes other than originally stated in the privacy policy.

Because of the importance of the issue, we will see more regulations in the future.

Wireless

Earlier this year, I moderated a panel on "Harnessing the power of the wireless Web." It included senior executives of carriers, content providers and platform developers. We began discussing the subject of wireless advertising and how the revenue streams are just ready to flow from location-based advertising that consumers are going to love and can't live without. I suggested that "true" location-based services are not going to be here anytime soon. Second, I stated that unless we develop UI standards for devices that allow complete control over what comes to them (when and how users want it) wireless advertising is just not going to fly. For a carrier to monetize location services, they must develop end-user subscriber applications, giving users the control over what application gets which personal information and which personal information is off limits.

The Wireless Advertising Association (WAA) has made progress in defining standards and measurement definitions, but more needs to be done to cover a range of devices and technologies. Privacy filters are critical to the success of any consumer wireless application or service. It's an absolutely critical element that needs to happen before any location-based services see the light of day.

Do's and Don'ts of Privacy

It's also in your interest to get audited by third parties like PwC, E&Y, Truste and others. Everybody in the value chain should work proactively on the privacy issue.

Protect consumers from your partners as well. It's not permissible to ship user information to your partners without legally binding contracts that adhere to your privacy standards so that your partners won't misuse the data. User data should be guarded, just as you would guard any other sensitive information such as user ID and password lists. Also, you should stick to your privacy policy and not change it frivolously to suit your business needs.

The privacy issue is not just a desktop issue. It is an AORTA (always on real-time access) issue—anything connected to a network can potentially transmit personal information about usage habits. Recent issues with TiVo and other broadband platforms such as GPSs in rental cars, biometrics technology in stores/malls/companies to identify shoplifters, raise this issue. Aggregators such as Axiom, Experian, and Engage collect a broad range of customer data across various channels to build a lifestyle score – which could be used to either provide useful services or discriminate.

It is not in the best interest of product and service companies to stealthily record information without consent. It leads to PR nightmares, bad press, and a black mark on their privacy records. Ask RealNetworks, DoubleClick, Microsoft, Intel, and others. A corporation is NOT smarter than the society. Somebody somewhere is bound to figure things out, so why waste your time and effort in questionable activities?

Another risk that should be avoided is mixing collective (aggregate) data with personal information. As Mary Modal of Forrester once commented, "Companies should think of personal information and collective data as the church and state of Internet business. Keep the two separate."

We also need to keep in mind that when it comes to privacy, we can't generalize. We have to pay attention to both ends of privacy concerns. People who don't feel comfortable trusting the Internet need to be accommodated with consumer-friendly technologies and policies.

Privacy can be used as a competitive advantage mechanism. Companies will build up great brand value propositions around having the absolute reliable "privacy," and consumers will flock to these brands, as consumers won't have to question the reselling of their data. A new "brand value" will soon be total and unassailable end-user privacy. Computer makers can also help by bundling personal firewall products with their PCs.

Personalization

Now, let's talk about personalization a bit. It is all about instant gratification; meeting the customers one-step earlier than they expect you to. It is one of the more overused words but it essentially means that the computer tries to figure out who you are, what you want, and when you want it. Essentially, we are progressing to a nirvanic state of computing, where the computer on the other end can read our minds. Based on information about my devices, networks, computers, usage habits, navigation history, payment history, and so forth, it already knows what I want. It should not care how or when I get there.

The following need to be in place for this to happen:

- Substantial data needs to be collected and stored
- Effective data mining techniques need to be deployed
- On as-needed and as-permitted basis, information needs to be shared among various applications

User experience is a two-way street. To improve it, you have to have user information. If this doesn't happen, the promise of personalization falls through. Personalization can be used to wow and surprise consumers, which go a long way in building long-term trusting relationships.

Protection

The amount of security needed is directly proportional to the value of the information that needs to be protected. For example, under normal circumstances the value of a user ID and password list is high and should be protected at all costs, while items such as news releases or executive bios are not as critical as damage risks in the event of security breach are small (Figure 2.) $100 to the janitor (to wipe some confidential papers and passwords stuck on computers) is still cheaper than breaking complex algorithms and codes, yet people worry more about how many bits are being used for encryption rather than examining their policies and procedures or employing effective monitoring. Your security is only as strong as your weakest link. If you believe that authentication and encryption are enough, your e-security will e-vaporate at some point.

Figure 2

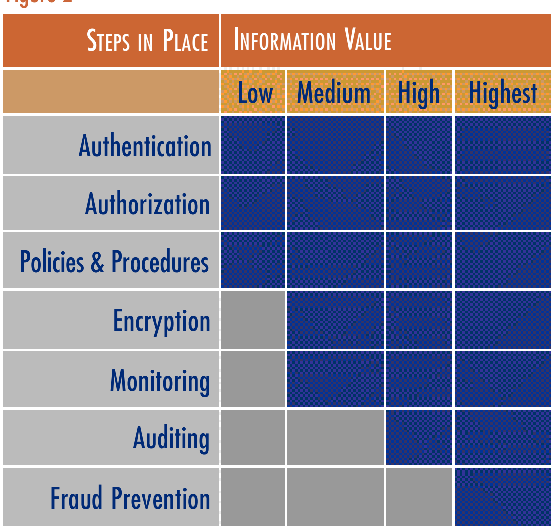

There are plenty of items to worry about, including hacking, IP spoofing, cyber squatting, viruses, worms, social engineering, salami techniques, piggybacking, packet sniffing, masquerading, logic bombs, software piracy, spamming, Trojan horses and more. Such activities can have serious impact on an organization, from loss of revenue to reputation. Hence it is wise to invest in security that provides adequate coverage corresponding to the value of the information you wish to protect.

So, there you have it. If you keep the three Ps of pervasive computing in mind when you design a service or an application, good things will happen to you.

Destinations of Interest

www.xns.org
www.epic.org
www.onename.com
www.w3c.org/p3p
www.privacyalliance.org
www.counterpane.com

Chetan Sharma

Chetan Sharma is Director of R&D and leads the Emerging Solutions Practice at Luminant Worldwide. He is author of the best-selling book "Wireless Internet Enterprise Applications" and co-author of the upcoming book "VoiceXML: Strategies and Techniques for Effective Voice Application Development" (both published by John Wiley & Sons.) He is frequently invited to speak at industry events worldwide and often quoted in media publications.

In Depth: Introduction to EAI (part 1 of 3)

By Ankur Laroia

Overview

Enterprise Application Integration, is a term that describes a process in practice for years: EAI is the combination of technology, methodology and business processes. It involves rethinking technologies and methodologies to make application integration a viable, cost-effective solution. EAI at its core needs to be driven by business issues, which require systems to seamlessly share information both, internally and externally across the enterprise. For EAI projects to deliver success, the following three key drivers must be present:

- Willingness to adopt new technologies
- Implementation of new methodologies, in support of new technologies
- A desire to fix years of architectural neglect

If an organization has all three drivers, it is then poised to take first steps into achieving "EAI Nirvana." Keep in mind that the road to nirvana is full of show-stopping obstacles. One must learn what EAI is really about—what processes, technologies and methodologies actually deliver success. This article serves to shed light on these key drivers and explore concepts related to EAI. It is the first in a series of three articles to set the stage and provide an overview of EAI, its beginnings and its future.

The Evolution of EAI

EAI has its roots in the era of distributed computing. With distributed computing came the challenge of exchanging information among disjointed systems because the mainframe was no longer home to all of the business applications. Specifically, the challenge was to seamlessly integrate these closed applications, which operated as islands of functionality. The approach taken to resolve this integration issue was simple: utilize synchronous FTP or another mode of file based data exchange, propagated in batch mode. In the beginning, this point-to-point approach worked beautifully, although it required a lot of custom coding on the application side to handle different types of files, applications and data formats. This point-to-point, file based methodology worked when the frequency of data transfers was low and the packet sizes were relatively small.

As distributed computing grew in popularity; applications supporting this computing platform became increasingly available. Now businesses were using out-of-the box, best of breed applications to support their core business processes. As these applications were implemented, the old point-to-point, file based mode of exchanging information was no longer robust enough to support the ever-growing distributed enterprise (Figure 3.) The reason was the file size and the complexity of data formats grew exponentially, not to mention the number of transports between each system. Synchronous FTP was no longer robust enough to send files from one system to another because the source and target systems would burn up cycle time contending with issues related to data transfers, rather than serving their intended purpose.

Figure 3

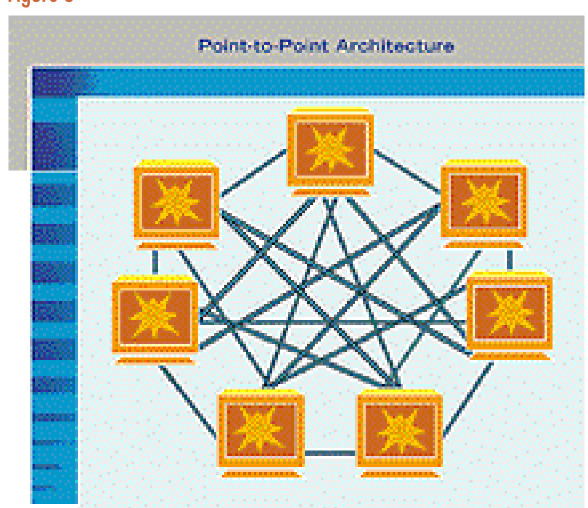

To compound the problem, as distributed systems grew, each application needed connections to other applications that supported a particular business process, such as order fulfillment. Soon, the enterprise became entangled in a point-to-point-based nightmare. Troubleshooting data transfers was tedious, time consuming, and frustrating. It was impossible for the information systems group to figure out which transfer had failed. ERP systems further compounded the problem because they introduced themselves as another variable into the integration equation. Now not only was there a distributed enterprise to contend with, but the addition of a mammoth system with its own share of idiosyncrasies led to the development of what is referred to as “middleware.”

Middleware: Silver Bullet?

Middleware was hailed as the silver bullet that would clean up the point-to-point, file-based architectural nightmare. Middleware provided a mechanism to effectively de-couple applications, data and business processes. Traditional middleware used asynchronous message queuing, coupled with a hub, to achieve seamless Enterprise Application Integration. Traditional middleware exploited the asynchronous ability of a message queuing-based middleware layer to propagate data from one system to another. The asynchronous ability of message queuing freed up the applications to serve business process instead of burning cycle time contending with data transfer issues.

The middleware layer undertook the responsibility of transporting the data from one system to another. The applications now sent data to one another in an asynchronous fashion by putting the export data on a queue and “forgetting about it”. Using a series of queues, the middleware layer seamlessly propagated data from one application to another. This allowed systems to be de-coupled, but getting the data from one system to another only solve half of the problem. What about data transformation and routing? A hub or message broker is used to examine data from multiple systems, transform it accordingly and finally route it to the intended target. Using a message queuing-based middlware layer, coupled with a content-based routing and transformation hub, solved back-end integration woes. The enterprise was robust once again.

IBM’s MQSeries coupled with MQSeries Integrator is a classic example of this technology. However, this “hub and spoke” architecture is facing new challenges. The Internet is fueling the need for reacting to business processes and events in real-time (Figure 4.) The Internet, in effect, is forcing applications to process ever increasing volumes of data and seamlessly support Web-based transactions over multiple systems in real-time.

Figure 4

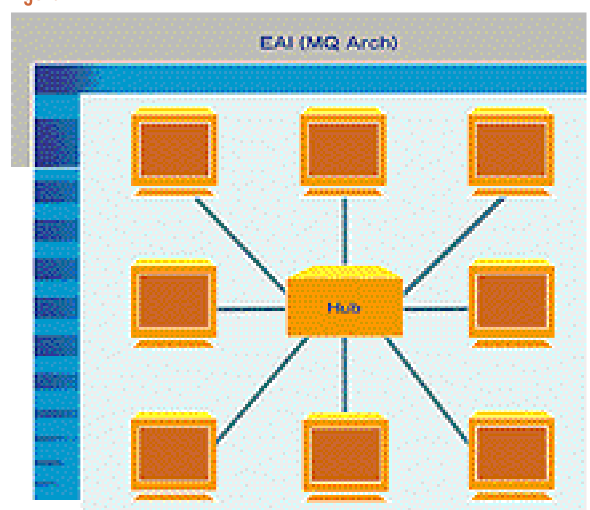

Tomorrow’s Business Drivers Fueling EAI Today

The EAI infrastructures of today are facing greater challenges. The need for implementing e-commerce initiatives and B2B ventures are two of the biggest business drivers, forcing application integration to new heights. Both of these drivers share a common theme, which deals with having to integrate the existing enterprise and extending data to the Web. With the advent of the Internet, e-commerce is a reality. For those organizations with the existing infrastructure, the challenge of integrating internally and extending out to the Web is a formidable one. Most CIOs are left pondering the merits and methodologies, which involve integrating existing back-end systems with "cyber-storefronts." The greatest challenge lies in "Web-enabling" an existing enterprise and extending it to employees, partners, suppliers and customers.

There are several issues to deal with, ranging from security, data integrity, data propagation, to assuring the scalability of distributed transactions. These issues are heuristics which most systems integrators face when undertaking a Web based EAI engagement. Most systems integrators now have set methodologies and approaches, that delineate a series of processes and steps to follow on the road to EAI. The key driver is the need for information to be processed in real-time. The entire enterprise must be able to react in real-time, to changing market and industry dynamics. This means that the enterprise must be wired for real-time, with decoupled data propagation and transformation based on business processes and events.

To support these business drivers, the enterprise must be integrated. I before E: "integration before e-commerce" is a popular tagline among systems integrators who have delivered success on e-commerce and B2B engagements. The integration of applications and systems in effect liberates the data that resides in them. To enjoy the benefits of having an enterprise that supports e-commerce or B2B initiatives, one must have the line-of-business applications integrated and exchanging data in real-time. Thus, if there is some level of integration present, one can leverage the data contained within the enterprise for e-commerce initiatives or one can share it with suppliers, vendors, and partners in a B2B setting. Either way, the enterprise still has to be wired for the exchange of information in real-time. EAI provides the foundation that makes this happen.

Wiring The Enterprise For Real-time

Integrating the enterprise in real-time allows for the decoupling of business processes, events and the line-of-business applications that support them. The middleware layer uses events in the business processes as triggers or alerts to perform actions to serve business processes. This methodology allows for the abstraction of business logic into the middleware layer, where it can function efficiently in a real-time manner, rather than embedding it into specific applications. This achieves two directives: It decouples an enterprise's processes from its systems and it allows one to plug in or replace line-of-business applications without having to contend with serious integration issues.

When the enterprise is wired for real-time, it can react to changes in real-time. The key take-away derived from wiring the enterprise for real-time is creating a zero latency environment. The propagation of data in a real-time environment is valuable for businesses involved in the energy, telecommunications, utilities, commodities, financial and manufacturing sectors as it can make a tremendous difference in the way transactions are done. Real-time data propagation enables transactions to leverage straight through processing.

Straight Through Processing (STP), provides a real-time picture of transactions as they occur. For example, a trader can track a counterparty's risk in real-time if the systems (trading and risk management) are wired for real-time data propagation. Another example is tracking a company's financial position on a daily or hourly basis. Once the applications in an enterprise are integrated, you can then leverage straight through processing techniques to enjoy lower transaction costs. STP offers a great way to leverage the existing EAI infrastructure to enjoy a high level of ROI.

In the next issue, we will explore different methodologies, technologies and architectures as they relate to EAI. Stay tuned!

Ankur Laroia

Ankur Laroia is responsible for the leadership and vision in defining and implementing Enterprise Application Integration (EAI)-based solutions and business strategies for Global 1000 companies. Under his leadership, Luminant's industry-recognized EAI experts help businesses successfully integrate disparate line-of-business applications. Laroia's experience includes Fortune 500 clients and spans a wide range of industries including pharmaceuticals, telecommunications, energy, messaging, utilities and high tech.

Laroia's accomplishments include contributing to the first patented EAI-centric methodology. A recognized expert in his industry, he is one of 2000 IBM MQSeries Certified professionals worldwide. He is also one of only 200 IBM Certified eBusiness architects in the world. Laroia has served as technical editor for two books published by Macmillian Publishing, "Realizing e-Business with Application Service Providers" by Louis Columbus and "Integrating Your Enterprise Applications" by Andre Yee. His work has also been published in EAI Journal.

Guest Opinion: Optical Internet Infrastructure

By Anupam Singal

In the past decades the digital telecommunications network has seen an evolution in the areas of transport and switching systems that spurred phenomenal growth. Newer technologies quickly came to life and matured. The biggest of these was the Internet, a medium that changed our lives.

The growth of the Internet was due to an organic process of change fueled by emerging technologies in the optical world. Ironically the past year has seen the optical bubble burst with stocks in this sector experiencing the worst returns, but the lasting effects of optical technologies for the Internet remain. Moreover we have yet to witness the revolution that optical technologies will bring to the Internet. The next generation of Internet infrastructure should see more growth and better services for users. There are challenges and technological barriers to overcome, but a gradual process of development will see new and interesting technologies emerge in the next decade.

The network for this infrastructure is experiencing change – network elements and systems capable of providing valuable services at lower costs. These services require more bandwidth and will run in real time. The goal is to establish one, integrated infrastructure to support all voice, data and video service that will become all pervasive. Users want global, real-time applications and want more desktop-to-desktop video conferencing, stock trading and medical advice. The Internet has increased the participation of users in advanced and advancing countries and brings people closer together through better communications.

Of course these changes will not take place overnight. One can examine the growth of the Internet going hand in hand with the development of optical networks and related technologies. The speed with which optical networking technologies have evolved is relative to the Internet's growth. Advances in the Internet were made possible by technological advances in the optical communications area. One can see a direct correlation between this. The content, applications, and the backbone are the building blocks for this revolution.

The Telecom world is divided into two broad categories:

- Transport systems
- Switching (routing) systems

The objectives of transport systems are to transport multiplexed (multiple) signals transparently across wide geographical areas without any corruption of the signal.

The function of switching and routing is to allow users to connect to each other with the ability to keep complex tasks at manageable levels while delivering good performance. Connecting about 10 billion worldwide users to services, or each other, is no small task. It takes the entire telecom industry to support this infrastructure. The network architecture designed to support this is comparable to building an airplane or a road for an unlimited number of passengers.

An important leap for the modern-day Internet was the creation of optical fiber in the 1970s and 80s. First used to transport electrical signals, it evolved to its commercial use as a single mode fiber capable of transporting high capacity signals over longer distances. Until optical fiber was widely accepted, most transport was carried by cable and limited to short distances and low capacity. It took a while to accept the idea that transporting electrical pulses through glass was not crazy, but this laid the foundation for the modern-day Internet. It has gained such large acceptance that each day about 4,000 miles of fiber is laid in the US alone.

Standard bodies such as:

- ANSI (American National Standards Institute www.ansi.com),
- ITU-T (International Telecommunications Union Telecommunications Standardization Sector www.itu.int),
- Bellcore (www.telcordia.com), and
- IETF (Internet Engineering Task Force www.ietf.org)

work vigorously to standardize the equipment interfaces across the US, Europe and Japan to accelerate growth of network equipment design and development.

SONET (Synchronous Optical Network) and SDH (Synchronous Digital Hierarchy) are two families of closely related standards used worldwide. They form the core of modern optical transport systems. Since 1985, T1X1 (an ANSI standards body) and ITU-T jointly worked with Bellcore Research to create SONET and SDH standards in 1988-89. Development of optical systems in the past decade revolved around this standard and formed the entire infrastructure to transport Internet packets.

Earlier systems from 1991-95 used SONET/SDH network elements with capacities of 51Mbps to 622 Mbps at STS-1 to STS-12 rates. Such equipment was connected to form linear, ring or mesh networks. SONET was the first form of network elements with standards stating that it would exclusively use fiber optics as the medium of transmission. This led to wider deployment of fiber with SONET, the first choice of network operators and carriers for transmission equipment. Since 1991 SONET has experienced the fastest growth in the North American transport business. North American market spending on SONET equipment was $8 billion in 2000. Similar market spending was seen in Europe where spending was around $6 billion in 2000. Asia Pacific market spending was $6.2 billion in 2001 and expected to grow to about $14.8 billion by 2004 (source: RHK, Telecommunications Industry Analysts.) This global market continues to grow at a healthy annual rate of 12-14 percent. Network construction by new and established carriers is driving demand for SONET/SDH network elements. Some of the preceding years saw growth rates of 60 percent in this market.

SONET/SDH systems carried about 800-1800 voice channels in 1991-95 as they evolved and increased in size and capacity. The data transport rates increased from 622Mbps, called as STS-12, to 10Gbps, called as STS-192, in the year 2000. Port and system capacities thus increased as initial systems carried voice as circuit switched and data as Frame relay. SONET lines now carry most Internet data through the network at about 1000 billion bits per second in a single fiber.

The hunger for higher bandwidth still remains, as does fiber capacity, which is around 40 Tera hertz for a single fiber. Compared to what is being exploited, fiber capacity is hardly nearing complete utilization. Variants of SONET systems evolved to carry more channels, each upon a different wavelength, and attempted to pack 100-200 different wavelengths of SONET on a single fiber, with 200 times capacity utilization. Such systems are called DWDM systems. Attempts are underway to pack more channels upon a single fiber at higher rates for larger distances. Even with 10-15 million subscribers, data, or voice, is now capable of traveling on a single fiber.

The optical transport area was the key focus of venture capital funding in the last decade. In 2000, venture capital funding totaled $130 billion in this sector and the market has not dampened this trend In Q1-Q2 2001 this optical transport sector has already seen about $30 billion of investment (source: RHK.)

The switching area has also seen unparalleled growth. Switches like ESS #5 (AT&T) and FETEX –150 (Fujitsu) or OCB-283 (Alcatel) were used until 1990 to primarily carry voice. Very little data was carried across the network, mainly as X.25 and other Frame relay technologies. The bulk of data remained in the LAN environment or within an enterprise. Local area networks were on Ethernet. In essence, there was a separation of the voice and data networks.

ISDN attempts to carry both voice and data gained some prevalence, but mainly data and voice switches and routers remained as separate network technologies rather than public voice switches. Four companies, Northern Telecom (now NORTEL), SPRINT, Sun Microsystems and Digital Equipment Corporation (DEC) founded a forum in October 1991 called ATM forum (www.atmforum.org.) The ATM (Asynchronous Transfer Mode) was a service independent technology of the entire network. This was essentially a packet technology that enabled convergence of voice and data through the same network.

During its first five years ATM forum produced more than 60 specifications. ATM as a technology attempted to create a converged network for voice and data services using the same infrastructure for both. This created an easy interface of the LAN to the wide area network (WAN) and became the early form of the Internet. IETF split into eight areas in 1989 and advanced the growth of Internet protocol in the last decade. By 1999, the ATM and the IP world merged to create a unified structure for voice, video and data. In the last decade the growth of ATM switches, IP routers, gateways, and bridges increased. The size of routing tables grew from 50,000 in 1999 to more than a million in 2001. New routers and networks were put in the network so fast that the IP protocol was running out of address space because the 32 bit IP addressing scheme was insufficient. The rigid, class-based IP addressing scheme was modified to Classless Inter Domain Routing Scheme (CIDR routing) to cope with network growth. The Ipv6, which is not yet fully deployed, will allow the Internet to grow even further.

This is the exponential growth area for switching and routing infrastructure. Ethernet packets are passed from a LAN to WAN environment, then carried by an ISP after converting them to IP or ATM through the core of SONET transport network. There is a big maze of switches and routers along the entire path over which a data packet travels. Bigger and bigger routers are being placed in the network. The ISP and data carriers require bigger and higher capacity routers every year. Systems that used to carry only voice at the advent of ATM and fast packet switching can now carry voice itself as packets. In North America the total traffic generated for data has far exceeded that carried as voice on circuit switches. In 2000, the data traffic was four times the voice traffic in the US market. Single routers that used to be called Core Routers handled total data of 10Gbps.

Today they are moving towards the edge in about six months and newer, higher capacity routers are required for the core. Bigger and higher capacity routers such as M40 (Juniper 40Gbps) M160 (Juniper 160Gbps), TSR (Avici 200Gbps), TN250 (Tenor 250Gbps) are the routers developed in 2000 that are serving the Internet today. These routers are running at full capacity. Network operators are requesting higher capacity, such as Terabit routers, for the core by 2002-03. The revenues for data carrying operators are also increasing. Until the year 2000, 90 percent of the revenue came from voice services. Although data revenues were insignificant when compared to voice, they are on the rise. Revenues in data services were estimated around $3.0 billion in 2000, a growth of about 40 percent for transporting data. Other than this, the enormous revenue generation comes from services that carry data.

Next generation system growth will stem from the following:

- Requirements for higher capacity systems
- Larger switching and routing fabrics
- Higher input and output port speeds
- Faster transport
- Faster network processing

All this is required for the infrastructure to keep pace with demand.

What does the picture in the coming years look like? How will the landscape of this optical Internet change in the next decade? It is difficult to predict what new services will be created, but a few things are almost certain in the infrastructure sector. System architects and developers of the next generation will help provide better, faster and cheaper services. This would enable the use of the Internet in ways previously thought to be impossible.

Big and fat pipes of 40Gbps shall start carrying data by 2003-04 in the core of transport infrastructure. This will be coupled with advanced DWDM systems working in bigger wavelength ranges, such as in C, L and S bands with support of about 200 – 300 wavelengths. Optical switches using MEMS or other techniques shall improve optical switching of the previously mentioned wavelengths. Multiservice transport boxes that can carry TDM voice, ATM, Ethernet or IP boxes will be deployed at the edge of the network to enable service independent transport to the core. Ethernet and IP will be carried in big fat pipes of 40G or more in the core of the network.

The Ethernet – once limited to LAN – will be transported to existing SONET infrastructure by standardization on SONET standards by the IEEE 802.3ae (IEEE working group) this year. Long-haul and ultra-long-haul transport will become more reliable. Its effects shall be seen in undersea fiber transport for global networks. Soliton-based systems will be used for Optical Signal Transport and would enable larger distances to be covered without signal regeneration.

Data and voice convergence will happen more frequently. New service architectures such as carrying voice on IP would be maturing by 2003. The core of the network shall enable convergence by use of technologies including as MPLS (Multi protocol Label Switching). Newer applications, such as those used to create virtual LANS or VLANS in the public Internet, should gain momentum. Newer protocols such as DiffServ will provide better QoS (Quality of Service) guarantees. IP networks systems will not only be capable of switching or routing but also be able to inspect flow of packets, and monitor and analyze them for billing and accounting based on packet flow.

Systems will be able to classify flows better, based on service requirements and priorities as well as provide some guarantees of QoS for the IP world. Deep packet inspection shall be performed. Firewall security will increase. All this packet analysis should be carried on huge data pipes of 10G to 40Gs by 2003. Higher capacity Terabit routers will rule the core and the existing core routers would move to the edge. Ipv6 will start to gain ground and deployment.

All of this should enable newer services, such as improved medical image transmission with QoS to serve hospitals and their patients. Desktop-to-desktop video conferencing using Internet VLANS will benefit enterprise customers, while video and broadcast on the Internet will not see packet drops and delays. Voice services on the Internet could be become as clear as today's voice on the circuit switched public network.

This picture may be a forecast of the next generation of optical Internet infrastructure; a future that may bring benefits to customers and operators alike.

Anupam Singal

Anupam Singal currently works as a Senior Design Engineer at Multilink Corporation, New Jersey. He has been working in the area of Broadband Optical Networks for about 10 years involved with the design and engineering of SONET/SDH and related products in the Optical Networking area. Prior to this he was working with FORE systems (now Marconi Communications) in the High capacity Terabit IP Router group. He has been actively involved with design, manufacturing and testing for Telecom equipment for IP in the Optical Infrastructure area . Prior to this he worked for C-DOT (Centre for Development of Telematics in India) where he was responsible for design of a MultiService SONET box to handle IP and ATM.

Multilink Technology Corporation *is a leading designer, developer and marketer of advanced mixed signal integrated circuits, modules, VLSI products and higher-level assemblies designed to enable the next generation of high-speed optical networking systems. Products span the markets from Metro to Ultra Long Haul optical transport equipment with a focus on the fastest commercially available speeds of 10Gb/s and higher. Multilink is headquartered in Somerset, New Jersey with additional offices located throughout North America and Europe.*

About Luminosity

Luminosity is a monthly electronic periodical focused on insights, trends, and advancements in emerging technologies.

About Luminant

Luminant Worldwide Corporation (Nasdaq:LUMT), a leading professional services firm focused on technology-enabled business solutions, helps Global 1000 companies capture increased revenue, improve productivity and enhance customer loyalty from the Internet and other enabling technologies. The Company's approach combines deep industry, functional business and technology skills with proven processes and strong vendor alliances to deliver solutions of significant economic value. Luminant offices are located in Atlanta, Chicago, Dallas, Houston, New York City, Seattle, and Washington, D.C. For more information, visit www.luminant.com or call 866.809.5868.

Call for Articles

If you would like to contribute an article to Luminosity, please send email to luminosity@luminant.com. Topics should be focused on emerging technologies and should not be of a promotional or marketing nature. Authors maintain ownership of all submissions.

Feedback

Your comments are always welcome. Please send your suggestions for future issues to luminosity@luminant.com.

To Subscribe

To receive Luminosity via email on a regular basis, please send email to luminosity@luminant.com.

To Unsubscribe

To cancel your subscription to Luminosity, please send an email to luminosity@luminant.com