

AOR TA

3Ps of Pervasive Computing *Privacy, Personalization, Protection*

"Soon, a person's importance in an information society will be measured by how **unreachable** they are"
- Anonymous

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

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

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 Pri-

vacancy 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 towards 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 usable 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.

Inside This Issue:

3Ps of Pervasive Computing—Privacy, Personalization, Protection	1
Phone Hacking: The Next Generation	5
Introduction to EAI	7

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) project 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) component, 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.*

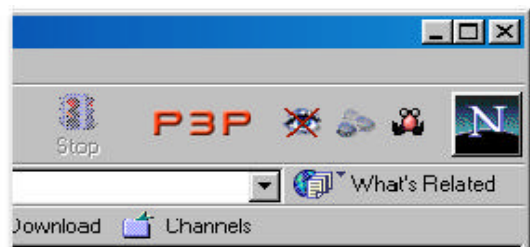
However, the problems with P3P include the following:

- *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 at any website, 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*
 - *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 vendor's meaning.



P3P Policy is acceptable



P3P Policy is not acceptable

Idcide's Privacy companion

- *Very complex user scenarios 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. 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.

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 (see AORTA Issue 5, E911 dials for help!). 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. In the last issue, we briefly touched on the "User Awareness Components" of which 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.

Destinations of Interest

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

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, Experion, 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 Real Networks, 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 steps need to be in place for this to happen:

- a) *Substantial data needs to be collected and stored,*
- b) *Effective data mining techniques need to be deployed,*
- c) *On as needed and as permitted basis, information needs to be shared between 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 goes 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 (see AORTA 3, “Adventures in Information Security” and, “Phone Hacking: The Next Generation” by Bruce Schneier). 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. \$100 to the janitor (to wipe some confidential papers and passwords stuck on computers)

Steps in Place	Information Value			
	Low	Medium	High	Highest
Authentication	√	√	√	√
Authorization	√	√	√	√
Policies and Procedures	√	√	√	√
Encryption		√	√	√
Monitoring		√	√	√
Auditing			√	√
Fraud Prevention				√

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.

There are plenty of items to worry about, including hacking, ip spoofing, cybersquatting, viruses, worms, social engineering, salami techniques, piggybacking, packet sniffing, masquerading, logic bombs, s/w 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.

There are two exciting columns in this issue of AORTA.

The first “*Phone Hacking: The Next Generation*” comes from Bruce Schneier of Counterpane Internet Security, a well-known authority on computer security & cryptography. His practical advice has benefited companies worldwide.

Luminant’s EAI practice leader, Ankur Laroia introduces us to the world of EAI in the first of his three part series, “*Introduction to EAI*”.

Enjoy.

My deepest gratitude to Bruce (and Counterpane), and Ankur for sharing their knowledge and for their promptness. I would also like to thank Joe Herzog of InfoSpace for helping with several ideas mentioned in this column.

If you would like to share what you know, please contact me.

Your comments are always welcome.

Best wishes,

Chetan Sharma

Phone Hacking: The Next Generation – Bruce Schneier

*Reprinted with permission from Counterpane Internet Security, Inc.
This article appeared in the montly "Crypto-Gram" newsletter in July.*

The phone network and the Internet are converging. That's good news for smart telephones, new telephony services, and customer convenience, and bad news for security. If you think that phone hacking is bad now, take a gander at what's coming.

During the last fifteen years or so, there has been a trend toward intelligent telephone networking. We've seen ISDN. We've seen SS7. We've seen IN (Intelligent Networking). These protocols are responsible for all the cool telephony features we've come to know and love: call forwarding, call following, local number portability, caller ID, etc. These features work fine, but are limited because they are all controlled by the phone company. If you want to initiate caller ID, you need to get the phone company involved. If you want your business calls forwarded to your home after 5:00 PM, you need to turn that on and off every day.

On the corporate side, we've seen Computer Telephony Integration (CTI), which didn't work very well because it was so big and clunky. It might be fine if you're a huge call center, but it just wasn't cost-effective for your average business. Development cycles were long, and service creation horrendously expensive; usage was rare.

But along came the Internet, and everything changed. The notion of intelligent endpoints (computers) and a dumb network (routers) turns the telephony model upside down. There are several consortiums and standards bodies working on bringing the Internet model to the telephone network, and allowing Internet-based control of telephone switching. The idea is to turn the telephone network into a giant networking resource that people outside the telephone network can control and manage. The benefit to the enterprise is more features and control: cost savings, better sales and marketing, improved customer service, etc.

The Parlay Group is a major player in this space. A consortium of software, hardware, and telephony companies, they are creating a specification and API to enable phone-system control from outside the secure telco network. This API will allow software to do such things as reroute calls, get notified of call attempts, retrieve the location of mobile users, and more. Even access to telco billing systems is planned. The idea is that computer applications can have integrated telephone components.

Even more fundamentally, all the switching protocols will interoperate at multiple points. Switches, gatekeepers, proxies, and call control agents will all be components of the new telephony control system. Control can be distributed or centralized, depending on the application.

Meanwhile, the IETF is defining the Session Initiation Protocol (SIP) for Voice over IP (VoIP) and more. This protocol will allow a user to define complicated ways to redirect calls: between 9 AM and 5 PM ring my office number, between 5 and 6 PM call my cell phone, after 6 PM call my home phone, and if my mother calls at any time, send her directly to voice mail. The protocol even includes a programming language, so a user can write a program to handle phone calls to match his own needs. While these features are nominally controlled by the user, the programs are stored in the telco network, and a DNS-like service is used to handle the profile and call forwarding. SIP is becoming a big thing; it's currently being used for VoIP telephony, will control calls in 3G wireless networks, and is being envisaged for all sorts of other uses like Instant Messaging.

The big idea here is to leverage the development techniques of the Web to services for telephony. New services are essential, because all the carriers have cut their collective throats on per-minute long-distance rates. Premium services are seen by many as the only source of meaningful revenue in the future. This means that telephony, which has heretofore been slow and methodical and reliable, will become as freewheeling as the Internet.

I am terrified at the security implications of these services. Sure, the Parlay spec says that communication between the Parlay client and Parlay server in the telco network is encrypted, and authentication will be enforced, but I don't believe for a minute that this will remain unhacked. SIP contains security provisions, but I don't trust them.

It's not the details of the protocols. It doesn't matter how many bits the key is, or what authentication protocol they employ: we've learned from experience that all systems like this are hackable. The worry is that these protocols open a huge hole into the telephone system. The problem is that these telephony control systems will sit on top of insecure operating systems. They will be hacked, and then things will get ugly.

Think about the possibilities for a minute. Denial-of-service attacks are a breeze: just reroute all calls to a person elsewhere. Or reroute all calls to a popular phone-sex service to another person. Or maybe just eavesdrop: set up a three-way conference bridge whenever someone receives a phone call. Remember the Trojan program that quietly made the modem dial Moldavia; this kind of system would make that hack a lot easier. And don't you think all of those hackers who chat on IRC would much rather take over a PBX and set up a conference call? You don't need me to think up the possibilities; there are lots and lots of them, none of them good.

One of the biggest backward steps is the re-merging of the control and voice channels. Switch and PBX hacking used to be very easy when signaling was done in-band. SS7 is an out-of-band signaling system, which separated the voice from the telephone control and made "beeping into the receiver" hacking impossible. These new IP telephony systems rebuild that old, vulnerable model.

It gets worse. The FCC is mandating that cell phone companies pinpoint phone locations to within 50-100 meters (for use with 911 calls). The carriers plan to use this information to create new data services based on location. The location information will also be available through services like Parlay for third parties to use. Imagine the security implications of that information getting into unauthorized hands. What if someone correlated a person's cell phone with his online identity? Could he pinpoint locations of desktop computers on the Internet? (This is actually a serious issue for 911 services. Unless one can somehow manage location information for endpoints, there's no hope of fielding a reasonable life-critical communications system based on the Internet.)

And think about reliability. The one thing about the telephone system is that it just works. That reliability is very hard to engineer using Internet protocols. As the phone system starts to look more and more like the Internet, it will become as reliable as the Internet. This means that it will forever be in beta. This means there will be software incompatibilities, upgrade problems, and random weird errors. This means that it will fail, catastrophically, once in a while.

Telephone hacking is not new. There have been decades of allegations and investigations into Las Vegas crime syndicates surreptitiously rerouting escort-service phone numbers, and the automatic telephone exchange was invented in the late 1800s by someone convinced that operators were rerouting his calls to rival businesses. Before the Internet, the phone network was the primary focus of hackers.

But it's a hard network to hack. Telephony is still a controlled closed universe. The protocols are often proprietary, access is limited, and information is scarce. You need to speak SS7, have the right physical connections, etc. There is nominally no interconnect to the TCP/IP Internet. Even with knowledge, it is the limited physical access that provides the most constraint. Voice and control are on separate channels. None of this provides absolute security, but it helps keep the number of hackers down.

The Internet, on the other hand, is much easier to hack. It's public. It's available. Anyone can connect a computer up to the Internet. Anyone can download boatloads of hacking tools. Anyone can become a script kiddie.

What we're seeing is another example of the tension between functionality and security. Opening the network is a good thing from the perspective of creating innovative new services, speeding up development cycles, adding value to data and voice. Yet when we do this, we open up the potential for the bad things as well. It's impossible to get the one without the other.

Soon the phone network will become just like the Internet. Putting control of telephony networks on the Internet means anyone can hack `chicago.switch.uswest.net`. These protocols will turn control over to both authorized and unauthorized Internet control. If you think phone phreaking was bad, just wait until anyone can do it.

Standards and companies active in this area:

<<http://www.parlay.org>>
 <<http://www.telecomsys.com>>
 <<http://www.invertix.com>>
 <<http://www.locationnet.com>>
 <<http://www.openls.org>>
 <<http://www.locationforum.org>>
 <<http://www.3gpp.org>>
 <<http://www.sipforum.org>>
 <<http://www.sipcenter.com>>
 <<http://www.etsi.org/tiphon>>

Steve Bass and John Ladwig both helped with this article.

Bruce Schneier is founder and CTO of Counterpane Internet Security Inc., the author of "Secrets and Lies" and "Applied Cryptography," and an inventor of the Blowfish, Twofish, and Yarrow algorithms. He served on the board of the International Association for Cryptologic Research, EPIC, and VTW. He is a frequent writer and lecturer on computer security and cryptography.

Counterpane Internet Security, Inc is the world leader in Managed Security Monitoring. Counterpane's expert security analysts protect networks for Fortune 2000 companies world-wide.

Introduction to EAI — Ankur Laroia

(Part 1 of 3 parts)

Overview

EAI is a term that describes a process in practice for years: EAI is the combination of technology, methodology and business processes. Enterprise Application Integration 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*
- *Implementing 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

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.

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 uses asynchronous message queuing coupled with a hub, to achieve seamless Enterprise Application Integration. Traditional middleware exploits 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 solves 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

point-to-point, file based mode of exchanging information was no longer robust enough to support the ever-growing distributed enterprise. The reason was the file size and the route it to the intended target. Using a message queuing based middleware 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**. 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**.

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, which 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 de-coupled 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

Thusly, 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 de-coupling 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 de-couples 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.


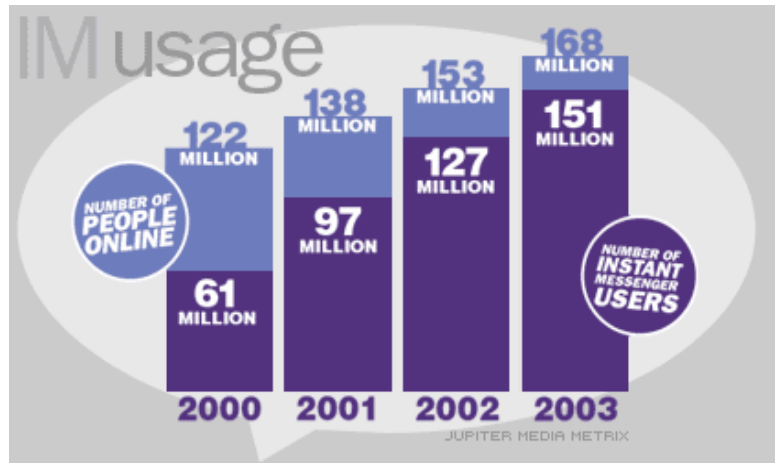
When the enterprise is wired for real-time, it can react to changes in real-time. The key take-away derived from wiring the enterprise for real-time is creating a zero latency environment. The propagation of data in a real-time environment is valuable for businesses involved in the energy, telecommunications, 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 & 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 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.

Stat Focus (Thanks to Tod Knight)



Next Issue

Preview of an upcoming book on building voice web applications

About Luminant

Luminant Worldwide Corporation, a leading professional services firm focused on technology-enabled business solutions, helps Global 1000 companies capture increased revenue, improved productivity and enhanced customer loyalty from the Internet and other emerging technologies. (www.luminant.com).

Call for Articles

If you would like to contribute articles, please send a note to aorta@luminant.com. Topic should be focused on emerging technologies. Articles should not be of a promotional or marketing nature. Authors maintain ownership of all submissions.