

Position Paper for the
OOPSLA 2003 Workshop on
“Pervasive Computing
Going beyond Internet for small screens”

Prepared By: Tony Willis

E-mail: tony.willis@t-systems.co.za

Company: T-Systems (South Africa)

Pervasive Computing

Introduction

The use of computers over the past 50 years can be broadly categorized into 3 eras. The first era, 1950 to 1970, saw computers being used to solve mathematical problems and perform logical functions. The second era, 1970 to the mid-1990s, applied computers to automating business processes. The third era, which seems to be taking shape, is seeing computers becoming an essential component of social interaction.

Within these 3 eras, there has been a continual fluid shift in focus between Personal, Enterprise and extra-Enterprise computing. These shifts in focus have been predominantly directed by two factors - the capabilities of the computer itself and the connectivity options available for that computer.

Mainframes and dumb terminals resulted in a focus on specific business operations that were available from the mainframe applications. These had limited impact on personal productivity and personal computing was not a mainstream reality.

The advent of the Personal Computer ushered in an era where applications could be focused on the individual and personal computing became a real option. Once a PC could be connected directly to a mainframe, the individual could suddenly achieve business and personal productivity from a single workstation, but was still restricted in terms of collaborative capability by the lack of connectivity to colleagues within the organization.

With the wide-scale adoption of Local Area Network technology in organizations and even in private homes, where many IT personnel have more than one PC networked together, departmental and inter-departmental collaboration became a reality. The additional application capabilities of e-mail, knowledge management, document management and workflow have greatly enhanced the organization's productivity options (or at least the potential thereof). Testament to this is the turnaround time available today for business processes that used to take days and even weeks, that now can be achieved within hours. One of the engagements that I have been working on with a South African bank involves a 15-minute Loan application decision, regardless of where in the country the applicant is at the time. This kind of capability is not new, as Enterprise computing and the automation of business processes has now been widely available and widely used for the last 5-10 years.

The advent of the Internet had a number of major impacts. Initially it provided access to global information that previously was not available due to limits placed on users by both the location of the information and also the location of the user. Even today, I am still astounded at how much information is (freely) available on the Internet. The second great benefit that the Internet introduced was almost limitless connectivity for e-mail distribution and then e-commerce. (I can buy Tennis rackets and Tennis strings from an e-commerce site in the USA, pay for shipping to South Africa as well as the import duties and still save almost 40% as compared to buying the same goods in South Africa. What makes this ability even better, is that I have a far wider selection of product than is available in South Africa.)

However, the IT industry soon cottoned on to one more benefit that the Internet produced, the connectivity options that were introduced could be used to transmit almost any data – including inter-company transactions (that rivaled EDI’s capabilities) and at a fraction of the cost. Thus the point was reached where global connectivity options became available for both personal and extra-Enterprise computing in addition to intra-Enterprise capabilities. Thus, it was now possible to collaborate and share information between organizations, resulting in Supply Chains that could now be completely automated and optimized – as long as connectivity to a network was achieved. What if this connectivity was not available?

The advent of PDAs (Personal Digital Assistants) and multi-functional Cellular phones in the last 5 years has taken away these connectivity restrictions. With the new devices that have become available and the connectivity options provided by Mobile Network Operators that support connection to these devices, the reach of computing capability and access to both Personal, Enterprise and extra-Enterprise functionality has become pervasive. So we enter the era of Pervasive Computing.

Definition of Pervasive Computing

Pervasive Computing has thus evolved primarily from advances in the areas of networks, connectivity and device capabilities in the last 5 years.

At last year’s workshop, the question around “What is Pervasive Computing” was raised and in fact consumed the majority of the day’s discussion. After last year’s workshop, I took a conscious decision to focus less on the definition of Pervasive Computing and more on the capabilities that it provides. In addition, my focus up until last year’s workshop had been on the human side of Pervasive Computing.


After last year’s workshop, I developed a matrix that included machines as part of a potential operating or business model for Pervasive Computing. This matrix is depicted in the diagram below:

Pervasive Computing Business Models

B2B Formerly the domain of EDI	B2C Marketing, Spam	B2E HR, Company News & Events	B2M Software upgrades, OTA
C2B Orders, Buy/Sell, Updated Personal Info	C2C Auctions, SMS, MMS, payments, data sharing	C2E Orders Documents Mail	C2M Home Appliances/ Security/ Environment
E2B HR, Leave Application	E2C Customer Service, Technical Data	E2E Trade Unions, Collaboration	E2M Office Tools, Presentations, Security
M2B Problem ID & Reporting	M2C Home Alarms, Problem Notification	M2E Office Maintenance, Outsourcing, Early warning	M2M Sensors, Intelligent Systems

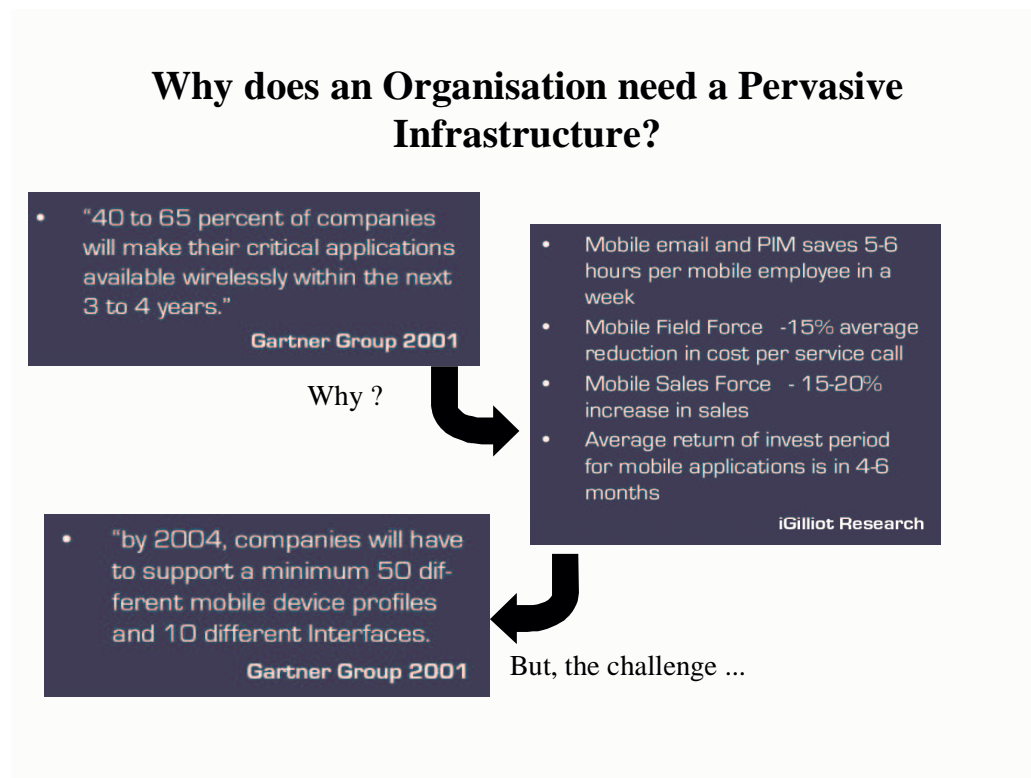
• The Players :

- Business
- Consumer
- Employee
- Machine


Convergence is our business. Page 1

Change in focus on Pervasive Computing

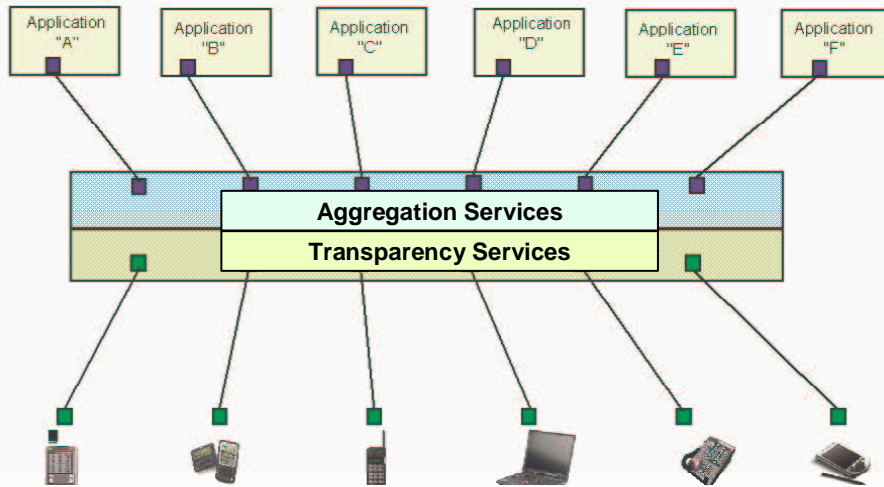
Up until recently, one of the challenges of Pervasive Computing has been to find the correct infrastructure solution for Pervasive Computing in the Enterprise. This is due to the wide diversity in the range of available devices as well as the different connectivity and processing models. This problem is well enunciated in the following diagram:



The generally accepted solution to this problem has been fairly consistently architected by most vendors. This architecture essentially provides a single interface into every back-end system and a translation layer (generally from XML) into the markup language or UI format required by the target device.

Most of the major vendors seem to have solved the basic “Transcoding” problem that was associated with the early days of Pervasive Computing for Personal users. IBM, Microsoft, Sybase, 724.com and even T-Systems International have Pervasive Infrastructure products that allow a user or employee to access back-end systems from a wide range of devices and a combination of technologies. These vendors all adhere to the basic Transcoding Architecture depicted in the diagram below.

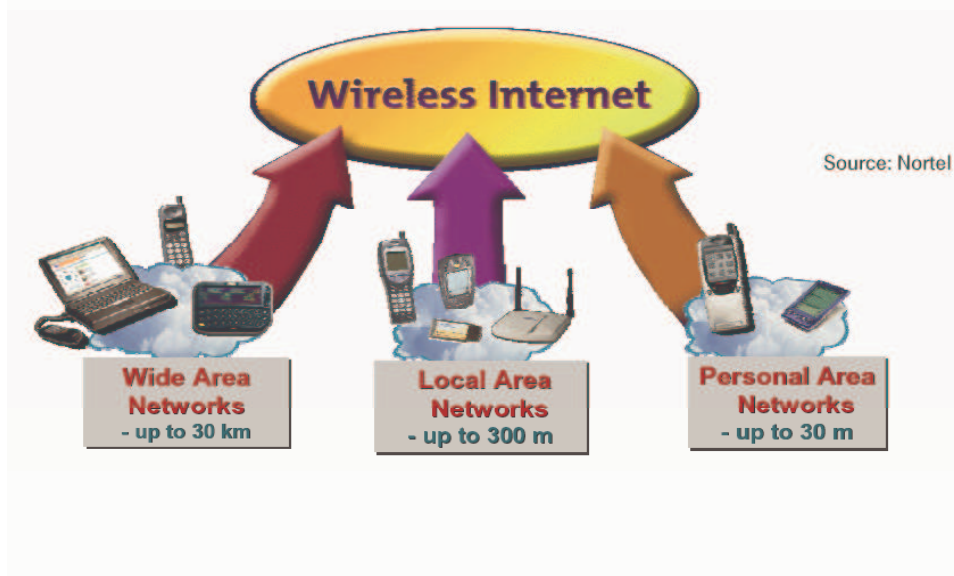
A Channel (Transparency) Layer abstracts device connectivity, while the Legacy (Aggregation) Layer provides a single integration point to applications



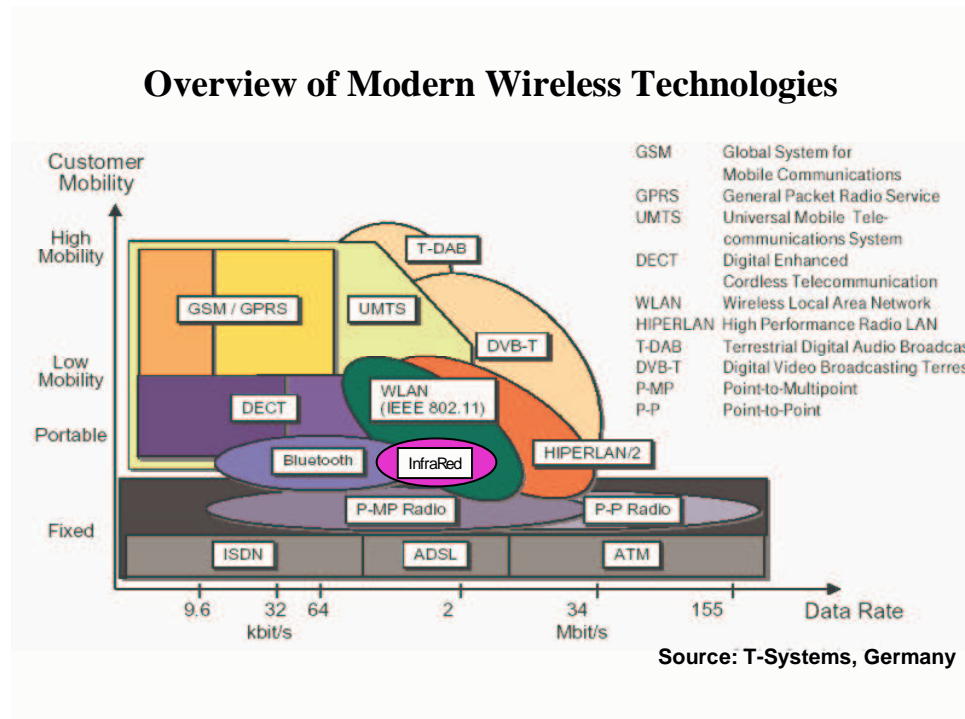
So if we have solved this basic problem of how to achieve Pervasive access to back-end systems and their information and services, what are the challenges left for achieving this Pervasive world?

The first problem involves the fact that there are essentially three different connectivity models as described in the diagram below and the coverage (in terms of distance) and speed implications associated with each are different, as are the security implications of each.

Pervasive Computing involves 3 main connectivity models



The technologies supporting these connectivity models are varied, not only in their underlying characteristics, but also in their deployment around the world. If one looks at Wireless Wide Area Networks, satellite and GSM technology are the most widely used and widely accepted globally. But GSM is not the de facto standard in the USA. In addition, a number of countries have progressed to 3G networks (Japan, Holland, Korea), whilst less developed countries are still operating with 2G and 2.5G (GPRS and EDGE) with no short term or even medium term plans to move to 3G. The range of technologies and their associated data speeds and mobility provided is depicted in the diagram below. This diagram is by no means comprehensive but is intended to depict the vast disparity in connectivity technologies in the Pervasive space.



1) Seamless Roaming between Connectivity Options

Given that we can now provide Pervasive access to information and services from a number of devices to the required back-end applications, the next move is towards the “Always On” world. This requires the ability to seamlessly hand over connectivity and the associated context and session management between the different connectivity options, but more importantly requires a management layer on the infrastructure side to track and control this handover. The problem is exacerbated by the fact that there are different security models employed by the different connectivity options and these need to be abstracted and hidden from the user to provide the ideal user experience.

One of the first products that seem to have achieved this goal is the Websphere Everyplace Connection Manager product from IBM. Further information on this is at (<http://www-1.ibm.com/industries/wireless/doc/content/solution/278603104.html>). We can expect to see a number of competing products appear in the near future. The end result of this capability is well depicted by a diagram from Gartner Group:

The Real-Time Enterprise in an 'Always-On' Society

Individuals

- "Always-on" customers and employees
- Many connection technologies and service providers
- Blurring of work and leisure
- Communication with a very broad range of people and devices



Enterprises

- Link people with information
- Capture information early, never freeze information
- Dynamic real-time collaboration
- Illuminate information shadows
- Flexibility at all times and in all places

© 2009 Gartner, Inc. and/or its Affiliates. All Rights Reserved.

Gartner

The greatest implication of this capability is actually nothing to do with technology, but rather is a social dynamic with severe consequences for quality of life. Even before we reach this "Always On" capability, many people find that they are "too connected", due to the fact that they have a mobile phone and can always be reached. Many people I know have taken to switching their mobile phones off on weekends, simply so that they can disconnect from the world at large. This social dynamic should not be lost sight of – many people want the connectivity, but not necessarily all the time and this has significant socio-dynamic implications for developing applications and solutions that take cognizance of the end-users context at a particular point in time.

2) The Holistic Pervasive Infrastructure

In the past 12 months I have encountered a number of organizations that have started Mobile, Wireless or Pervasive initiatives with a limited scope in mind. The Architecture of these solutions has not allowed for a growth path and essentially they have painted themselves into a corner (either by opting for a single device solution or for a single connectivity channel). Within 3-6 months of going live with the solution, these organizations suddenly need to support new Channels and new Devices AND their initial Architecture did not include any function/channel mappings, different notification models, device management etc.

I believe it is essential for organizations to architect holistically when it comes to Mobile and Pervasive solutions, even if they deploy selectively to begin with. The growth path or expansion path for these organizations is vastly simplified by doing so.

Factors that should be considered in arriving at a more holistic Architecture include:

- Notification Services – different mechanisms to notify users of occurrences, confirmations, reservations
- Location Based and User Context Services – providing information only relevant to a user or his context /location
- Synchronisation Services – to allow users to synchronise mobile devices with data stored on a central server or laptop
- Device Management Services – providing the ability to manage device configuration, gather hardware, software, and configuration data about mobile devices, distribute applications and files to mobile devices and recover configurations after device loss or theft
- Function/Channel Mapping – allowing the management of which functionality is available over which channels, thus preventing transactions that require security from being shipped over insecure channels
- Device/Channel Mapping – allowing certain devices and device types to be used on certain Channels only
- User Connectivity Management – allowing users to “hop” between connectivity options without losing context or data and providing a “truly mobile” experience
- Message Resend – allows a message to be resent where the message has not been received by the intended recipient
- User Authentication Mechanisms – different Channels allow for different User Authentication Mechanisms
- Security – due to the different Connectivity options and their varying mechanisms for securing data that is passed over these networks, a variety of Security mechanisms is required as well as the ability to apply them to different User scenarios

3) Return on Investment and a Business Case

Although this issue is by no means specific to Mobile or Pervasive Computing, it invariably crops up after a Mobile solution has been rolled out and is successfully being used. The organization concerned cannot understand why they are not making any money from this solution or achieving any business benefit.

It is vital for organizations to decide, before they enter into a Mobile or Pervasive project, whether they are intending to make money (profits) from a solution, whether they are going to save money from a solution or whether the solution becomes a necessity to stay competitive with other organizations in the same industry. Once this decision has been made it is essential to calculate the actual (direct and hidden) costs associated with the specific Mobile or Pervasive solution.

There are a number of big success stories around Mobile and Pervasive computing, where organizations have achieved a Return on Investment within 3-4 months. But

there are other examples of organizations that will never recoup any Return on the Investment made, nor will they achieve any significant business benefit. For the sustained success of Pervasive Computing in an Enterprise, it is essential that we try to eliminate the “disaster stories” by pro-actively and realistically assessing the costs and benefits to be achieved from Pervasive projects.

Summary

The combination of holistically Architecting Pervasive Applications for the future and applying business (financial) logic to Pervasive projects will significantly improve the success rate of Pervasive initiatives in the industry. In addition, taking note of the move towards the Always On world, but understanding the social impact of this capability is essential to prevent a backlash against the Pervasive movement.

Pervasive Experience

My name is Tony Willis and I work for T-systems in South Africa as an IT Architect.

I have been involved in Pervasive Computing projects since 1999. In that time, a number of architectural projects have been undertaken, including:

- Blueprint for a Pervasive Gateway infrastructure for a leading South African organization
- Architecture for a Mobile Payment Gateway that currently provides Pervasive access from 7 different channels
- Development of the South African Government's e-Gateway
- Reviewing the Pervasive strategy and Architecture for one of the largest two banks in South Africa
- Evaluating T-systems International's Mobile Corporate Portal

I have also been involved in a number of initiatives to develop and deploy Pervasive Computing solutions. These include:

- Innovation laboratories with a number of customers to demonstrate Pervasive functionality
- A Mobile project for a leading beverage manufacturer
- Providing T-systems executives with Pervasive access to their e-mail and diaries.

Lastly, I have participated in the Pervasive Computing workshops at OOPSLA in 2000 and 2002.