
BTECH 450DT 2003

Computers New Zealand:
End of Semester Report

Anisha RAMJI

2521221

CONTENTS

Abstract	2
Personal Profile	3
Company Profile	4
Project Objectives.....	6
Scope of Learning Topics.....	7
Project Deliverables.....	8
Research & Development.....	9
Knowledge Base.....	34
Conclusion.....	47

Appendices: -

Knowledge Base Initial Assessment Report.....	50
Gantt Chart.....	52
Discussion Papers;	
Microsoft Greenwich.....	53
Jabber Instant Messaging.....	56
Windows XP Deployment Site.....	59

Abstract

This report provides an account of the project I have completed, along with Vincent Yee and Chris Chou, for Computers New Zealand (CNZ) under the supervision of TN Chan. The first section of the report gives an outline of the actual project including the objectives, learning scope and deliverables. It also provides a brief background of CNZ and myself.

The report details the first part of the project in which we were acting as part of the Research and Development team for CNZ, collecting & collating information and writing discussion papers for reference within the company structure. The final versions of these discussion papers are included in the appendices section of the report.

The second part of the report contains further information relating to the actual implementation of the CNZ Knowledge base, as a continuation of work done by the previous years BTech project team. It provides details of the current state of the Knowledge base and related documentation, modifications and work completed during the implementation phase.

Also included in the report is a summary of the practical skills and knowledge I have learnt throughout the project, as well as previous knowledge gained at University that we were able to apply to the project.

Personal Profile

Name: Anisha RAMJI
UPI: aram033
ID: 2521221
Degree: BTech(IT)

Papers completed to-date:

- PHYSICS 120, 150, 243
- MATHS 108, 162, 208
- STATS 107
- INFOSYS 110, 222, 224, 329
- ACCNTG 101
- COMPSCI 101, 105, 210, 220, 230, 314, 334, 335, 375

Papers enrolled in for 2003:

- COMPSCI 708, 742
- SCIGEN 201
- BTECH 450 A & B
- INFOSYS 339
- PHIL 101

I chose this BTech project to give me the opportunity to apply my theoretical knowledge gained over the last 3 years of my BTech degree at The University of Auckland. Throughout the duration of the project I hoped to gain invaluable industry experience and learn and develop skills relating to project management and working in a team environment.

Company Profile



Computers New Zealand is a total information solutions provider with skills and capabilities to implement information and automation systems for meeting clients' business objectives and increasing their productivity and competitiveness.

Started in 1997, CNZ is a conglomerate of IT companies and practitioners in New Zealand. The company has a strong reputation and is known as being the only engineering-based computer company in New Zealand. The CNZ group is directed and managed by principles who each have more than 10 years of experience within the information technology industry. As a result, many IT or PC service providers are looking to CNZ as the source of knowledge and the vehicle for technology development. Clients already served include large corporations, local governments, down to small business and home-based business operations.

The founding member, Compucon New Zealand, is well known in the industry for its vigorous quality management approach for computer system production. The company has achieved ISO-9002 quality management standard since 1995. Compucon and CNZ believe that IT practitioners should take the same approach in providing IT solutions for clients in an increasingly complex IT market. More recently, the company has started developing high-tech solutions to complement the range of system hardware. In 2001 CNZ worked in partnership with Internet Service Provider NZNet to offer a Linux based firewall and mail server. Since 2002 the company has also been working with a software developer on an integrated business application database.

CNZ maintains a standard set of operational policies and procedures for activities including proposal, consultancy, order processing, product sourcing, servicing, delivery, documentation and warranty claim handling. All members are committed to zero mistakes within a strict QA Management programme. The Group maintains a code of ethics prescribing total dedication of service, honesty and openness with clients and colleagues, and a positive contribution to society at large.

Challenges undertaken by the group for clients include:

- o migrating to new platforms
- o aligning IT to business
- o reducing costs of ownership
- o keeping abreast of IT developments
- o develop effective IT investments
- o manage distributed data
- o manage vendors and suppliers

Products & Services offered include:

- Full Consultancy - review of client workflow structure, appraisal of existing information systems, planning of platform installation or migration, preparation of IT policies and procedures.
- Solution Provision - customized information technology solutions from conception to turnkey delivery to staff training.
- Full Liaison - a single channel for the supply of IT equipment and services including file servers, desktops, vendor installation, facility management, usage audits, on-site support and staff training.

Project Objectives

Project Details

Project Period:	Mid Feb 2003 - 26 Oct 2003
Workload:	10 hours per week
Work Allocation:	Each student will take up one technology scope and be involved jointly in other work areas

The company wishes to establish the infrastructure and content of an engineering and information technology knowledge base. The infrastructure refers to a framework that will identify, develop if appropriate and accumulate:

- (a) relevant technical expertise and know-how
- (b) state of the art of technology applications in New Zealand
- (c) procedures for handling system integration and implementation projects

The project consists of 2 main areas; Research & Development and Knowledge Base implementation.

In the Research & Development phase of the project we are required to apply and refine the researching and report writing we have been taught at University to write up discussion papers on various topics. These papers will be used at a later date and stored within the Knowledge Base (KB) for use by the company's consultants.

In the Knowledge Base implementation phase we were required to assess and document the current status of the KB and use this to determine what further work is required to complete the project started by the 2002 team. This will involve applying the practical and theoretical programming skills we have learnt over the last 4 years at University.

Scope of Learning Topics

<i>Technical:</i>	System architectures such as iA32 and iA64 System platforms such as Springdale and Grand Champion CPU such as P4 Northwood and AMD Hammer System Redundancy and Clustering Power User Computing platforms Data communication such as IEEE 802.11x and 1394 Mobile standards such as CDMA and 3G Internet access and security such as WEP LAN/WAN such as Active Directory and LDAP Windows 2000 and .NET Linux kernels and distributions Database development based on Open Source products Software programming and scripting Web Services components including XML System integration and interfacing design
<i>Commercial:</i>	The IT and engineering market Value adding concept and vertical solutions Worldwide technology developments Business models such as ASP (application service provider)
<i>Management:</i>	Time scheduling and control Specification control and Quality assurance Contracting and Cost control Human resource organization

Project Deliverables

Expected deliverables for the CNZ BTech project 2003 include:

- Business Management QA Manual
- Project Management QA Manual
- Web Site for Public Reference (restricted access)
- Knowledge Database of Technology Disciplines
- Database of Potential Customers
- Research Reports
- Industry Channel Presentation Slides

Research & Development

The first phase of the project involved acting as part of the Research & Development team for CNZ. The aim was to provide reliable and useful information on various areas in Information Technology, so that it could be used for future reference both within the company and with clients.

Each of us three group members were assigned a certain technology area, for Vincent it was networking & communications, Chris was assigned hardware platforms and myself, software. The topics on which we were required to research and report were determined by TN Chan and discussed at our weekly meetings. The objective was to understand what the technology was about, how it fits in with other complementary and competing technology and provide evidence of its promise. This was done from the perspective of an IT consultant to enable us to recommend the most appropriate technology to meet customers' requirements. We collected and collated information from various sources and produced reports, refining them as necessary until they reached a certain level of reliability and professionalism. This was a very beneficial task as we were able apply the research and report writing skills we had learnt at University, and further refine them during the process. We were also encouraged to provide an account of the sources from which we acquired information and to rate these sources in terms of our perception of their reliability.

The reports were to be used to populate knowledge base once the second stage of the project had been completed.

The following gives a brief outline of the topics I have been assigned & researched;

Microsoft Greenwich

Microsoft Greenwich is the latest in Real Time Communications (RTC) released by Microsoft, developed for enterprise use, to provide secure and manageable real-time communication, including instant messaging, video conferencing and voice communication. Greenwich was created to be a secure, manageable enterprise Instant messaging solution out of the box and also as an extensible, standards based, real-time communications platform.

Benefits of the Greenwich platform include:

- Presence based functionality
- Integration with MSN Messenger Connect for Enterprise service
- Strong end-to-end communications based on industry standards
- Easy administration based on industry standard technologies and familiar Windows based tools
- Platform that includes data collaboration, voice, video and allows for future communications applications
- Capabilities for logging and archiving instant messaging conversations for regulatory compliance

Jabber Instant Messaging

Jabber is an open source XML protocol for the real time exchange of messages and presence between any two points on the Internet. The main advantage of Jabber is its interoperability between other IM systems, such as AIM, MSN, ICQ and Yahoo. This is achieved by the Jabber protocol 'masquerading' as the foreign protocol, acting as a translator between the two systems and providing seamless integration at the user level. The XML platform is the basis for peer-to-peer collaboration, creating the interface for sharing enterprise computing resources in the P2P environment. This is done through the combination of XML and Simple API for XML (SAX) to create a way in which to pass data and call remote processes.

Jabber advantages over legacy Instant Messaging systems include:

- Open: the Jabber protocol is free, open, public, and easily understandable, and multiple open-source implementations exist for Jabber servers, clients, and development libraries.

- Extensible: using the power of XML namespaces, anyone can extend the Jabber protocol for custom functionality
- Decentralized: anyone can run their own Jabber server, enabling individuals and organizations to take control of their IM experience.
- Secure: any Jabber server may be isolated from the public Jabber network, many server implementations use SSL for client-server communications

Windows XP Deployment Site

The Desktop Deployment Centre site from Microsoft has been designed to aid IT professionals in the migration to the Windows XP operating system and its partner product, Office XP. It includes various resources and tools to assist in the evaluation and deployment stages, as well as documentation and best practise advice from current users and professionals.

Tools available on the website include:

- Microsoft Application Compatibility Analyzer 1.0
- Windows Application Verifier 2.5
- Compatibility Administrator

Linux Research

As a combined industry & technology research assignment I was given the task of researching the implementation of Linux into a small business environment. I was given the situation of a small business running a central server and 10 desktops, and required to research into implementing a business network using the open source Red Hat distribution and to also determine the skills and software required to support it. In addition to this I was also required to research the applications available to support various requirements such as file sharing, email, Internet access, accounting systems and general office productivity. The aim of this was to determine whether Red Hat Linux would be a commercially viable option for small businesses, and to provide information for users and IT professionals to make informed decisions based on this information.

Over the second half of the project I have been researching this topic quite in-depth, the following is a brief outline of some of the reports created for the project and a summary the information contained in them. Full versions of the reports can be found in the appendices.

Linux is an operating system based on the time-tested UNIX system. Linux is 'open source' software meaning it has no one particular owner, instead it is freely available and able to be modified by anyone. This concept of open source programming originated from when Universities needed to be able to share their work and allow others to adapt programs to meet their own needs. This was also the basis used by Linus Torvald, a student of the University of Helsinki at the time, when he started developing Linux in 1991 as a new version of UNIX. Torvald combined talents with a group of programmers to create the core operating system, which is now one of the most widely ported operating systems available for the PC.

The popularity of Linux, in both personal use and the corporate environment, has increased significantly as more users look to an alternative to the widely used Windows operating systems. Unlike proprietary operating systems, Linux can be installed and upgraded for free which makes it a good option for businesses that don't have a high budget but still want a reliable operating system. As Linux programs are also open source, they can be installed freely without major licensing concerns and offer the type of flexibility not found in other operating systems. As with open source software any bugs can be fixed quickly and easily, without having to wait for costly upgrades or service packs. Today Linux is being hailed as the fastest growing server-side operating system available and it provides some convincing arguments as to why.

Benefits of Linux:

- *Lower software product costs*

The software can be freely installed on an unlimited number of machines having the effect of an "unlimited-user, unlimited-installation" license.

- *Lower software support costs*

Hewlett-Packard has offered unlimited, 24/7 worldwide phone and e-mail support for Linux for a fee of \$130 per month per server

- *Lower hardware costs*

Another advantage of Linux is its ability to run on older machines with less memory and disk capacity, which translates into savings on hardware upgrades. Linux provides excellent performance on new hardware while still running adequately on older machines.

- *Better performance for lower cost*

Tests have revealed that Linux quite consistently beats Windows NT not only on single-processor machines, but also on multi-processor machines.

- *More stable*

The Unix design, which Linux shares, is a time-tested one. The open-source code model of Linux seems to ensure that bugs are detected and fixed early.

- *Excellent interoperability with other systems*

Linux is widely claimed to be able to coexist with other operating systems, including the Netware environment, Macintosh, and IBM mainframes.

Weaknesses of Linux:

- *User friendliness*

There are two graphical desktop environments for Linux. KDE ('K' Desktop Environment) and Gnome. Neither interface is as yet as polished and complete as the Windows or Mac desktops.

- *Installation difficulties*

Even though Linux does a reasonably good job of detecting all the hardware components on a PC, it does require some enlightened input from the installer. Linux installation is getting progressively easier, with reports of Red Hat version 6.0 and later being quite simple to install.

- *Scarce applications*

The most glaring weakness in the Linux story was the complete absence of strong database. On the desktop, Linux also has a long way to go. The applications are there, but the most important one, Microsoft Office, is not. However, there are many similar application suites available, and being developed.

- *Unviable business solution*

Independent software developers need to target Microsoft's Windows, because it is the most popular desktop platform, but they soon face competition from Microsoft's own products.

Microsoft Windows Based Product	Cost per machine	Linux product	Cost per machine
Microsoft Windows XP	New PC: \$100 Upgrade: \$100 Professional: \$300	Red Hat Linux Debian Linux	Red Hat dwnld: Free Debian Linux: Free Red Hat box: \$39.99
Microsoft Office	Small Business Edition: \$200 Office Professional: \$350 - \$600	Open Office Star Office	Open Office: Free Star Office: \$76
Microsoft Explorer	Free	Mozilla Netscape Opera Konqueror	Free
Microsoft Outlook	Free	Evolution	Free
Microsoft Project	\$600	MR Project	Free
Microsoft Visio	\$200	Dia	Free
Adobe Photoshop	\$600	GIMP	Free
Adobe Acrobat	Free	Adobe Acrobat	Free
cuteFTP	\$40	IglooFTP Pro	Downloaded: free Purchased: \$30
Digital Camera	Free	gtkam	Free

Linux Operating Systems

Red Hat Enterprise

The Red Hat Enterprise Linux family is a comprehensive suite of Linux operating systems, combining the innovation of open source technology with the stability of an enterprise-class platform. The suite consists of;

Red Hat Enterprise ES - Designed for smaller systems with up to two CPUs and 4GB of main memory. Provides the core operating system and networking infrastructure designed for entry-level and departmental server applications. Ideally suited for network, file, print, mail, Web, and custom or packaged business applications.

Red Hat Enterprise WS - Provides support for workstation/desktop systems with up to two CPUs. Red Hat Enterprise Linux WS is the desktop/client member of the Red Hat Enterprise. Ideal for client-server deployments, software development environments and targeted ISV client applications.

Red Hat Enterprise AS - Supports the largest commodity-architecture servers with up to 16 CPUs and 64GB of main memory. Red Hat Enterprise AS is the core operating system and infrastructure enterprise Linux solution.

Red Hat Linux 9

Based on the Linux kernel 2.4.20, Red Hat Linux 9 combines the latest technology from the Open Source community as well as the latest in Red Hat technology. Red Hat Linux 9 is available in 2 versions; Standard and Professional. Both versions incorporate the familiar Linux Bluecurve user interface as well as OpenOffice, the Linux software suite for office productivity. Red Hat Linux 9 also offers other tools and applications useful for any personal workstation ranging from basic e-mail and browsing to project management, these include:

- Ximian Evolution - e-mail client, browser and calendar
- Jpilot - pilot desktop software
- Mozilla - Internet browser, e-mail client, and address book

- gphoto2 - photo manipulation tool
- MrProject - project management software
- GIMP - graphic manipulation software

The user is able to choose the software packages most appropriate for their system during the install procedure for a more personalised system.

Red Hat Linux 9 Professional contains all the features of the Standard version with the added benefits of being able to choose a personal desktop, workstation, server or custom installation or upgrade an existing system.

Red Hat Linux Enterprise vs. Red Hat Linux 9

The most significant difference between the Red Hat Enterprise Linux package and the Red Hat Linux v9 distribution is in the support offered. Red Hat Enterprise Linux products are delivered through an annual subscription with one year of support services and access to Red Hat Enterprise Network. However, much documentation and support guides are available through the Internet from other sources, so that the Red Hat Linux 9 distribution would be a more than adequate solution in a small business network for a significantly less cost.

Linux Accounting Software

Given the task of investigating different accounting packages available for the Linux operating system and evaluating each based on varying criteria, I found the following;

AccountiX – AccountiX is a modular software accounting package based on VisionPoint 2000. It allows users to start with the basic system and add to it as necessary to create a more specialised system, tailored to the business' needs. Being open architecture means it can be further modified and customised to meet requirements and better accommodate the business. Good support and documentation are available for the software package, both online and in over 60 countries around the world. Available for single or multi user, in compiled or source code formats.

Linux General Ledger - Linux General Ledger is a general business bookkeeping system specifically designed for small to medium sized business and client write-up. Engineered for ease of use and functionality the software includes features such as Postgres File Transfer Protocol, Open source file structures and customisable source code. Linux General Ledger is a simple, yet effective program for business accounting that will run on Linux systems with the minimal of specifications.

GnuCash 1.8.0 - Designed to be easy to use, yet powerful and flexible, GnuCash allows the user to track bank accounts, stocks, income and expenses. Based on professional accounting principles to ensure balanced books & accurate reports, and backed by an active development community. GnuCash is compatible with many of the popular accounting data formats including Quicken and OFX and supports the import of these file types. The graphical user interface and support for internationalised dates and currencies makes the software very user friendly. Documentation is also available in 4 different languages. GnuCash, however, does have a few software dependencies, as it uses a number of libraries from the GNOME environment and may require a few extra packages to be installed to run correctly.

NOLA - NOLA is a General Public Licence accounting software created by Noguska. NOLA provides for the company's accounting, inventory, point of sale, contact management, billing, purchasing, and reporting all in one integrated package. It also allows for real time inventory quantity updates. The NOLA system is built around a secure, open platform and ships with the Apache Web Server. Also supplied is the MySQL database engine, a fast SQL server designed for large amounts of data. NOLA-Pro, a newer version with many improvements, has also been released, using PHP, PERL, Java, JavaScript and incorporating MYSQL and POSTGRESQL databases. The website contains much in the way of useful documentation, ranging from help files on installing the software, on both the Linux and Windows platform, to using the software for the first time.

SQL-Ledger - SQL-Ledger is a double entry accounting system. Accounting data is stored in a SQL Server, for the display any text or GUI browser can be used. SQL-Ledger can be used on any UNIX, Mac OS X and Windows computer. The application is written in Perl, developed on FreeBSD and Linux with Galeon, Konqueror, Netscape, Lynx, Links, W3M, Voyager, Explorer to render the display, Apache to communicate between the server and the browser, and PostgreSQL or Oracle to store accounting data. Templates are provided in html and tex format. The tex templates are processed with latex to produce postscript and PDF documents and can be sent to a printer, displayed in a PDF viewer or sent out via email. SQL-Ledger also provides a way to import customers, vendors and parts from Quickbooks. It is also available in 17 languages. SQL-Ledger is freeware under the GPL and the latest release, SQL-Ledger 2.0.8, is available for download from its website as well as excellent documentation and support.

Recommendations

I would recommend SQL-Ledger as the accounting software for Linux small business users. SQL-Ledger is developed under the General Public License so it has obvious cost benefits associated. The software package is well proven among Linux users and is backed by an active development community so is continually updated and refined. SQL-Ledger is also compatible with Quicken, allowing objects to be imported into SQL tables so that Quicken users can get started straight away. Excellent documentation and support is readily available both on the SQL-Ledger website as well as throughout the Internet.

As an additional topic, I was given the opportunity to investigate further into accounting packages that could be used in a Linux small business network, this time through the use of an emulator program called WINE.

Wine is an Open Source implementation of the Windows API on top of X and Unix. It does not require Microsoft Windows, as it is a completely alternative implementation consisting of 100% Microsoft-free code, but it can optionally use native system DLLs if they are available. It consists of a program loader, which loads and executes a Windows binary, and a library that implements Windows API calls using their UNIX or X11 equivalents. The library may also be

used for porting Win32 code into native UNIX executables, often without many changes in the source. Wine is free software, and its license is Linux General Public License style.

Wine's goal is to make it possible to run Windows applications on Unix. To this end it will provide replacements for just those DLLs and APIs that are needed by these Windows applications. This means that Wine will not provide replacements for DLLs that are not shipped with Windows or are always shipped with Windows application (e.g. the Visual Basic run time). As Wine is not yet complete there may still be programs that crash. However throughout the research I found there were a few additional criteria for successful emulation, including;

- Wine has to be able to 'see' an MS Windows binary (i.e. application) if it is to run it, in many cases you will need a directory and file infrastructure that is similar to an existing Windows installation. Some applications' installation programs want to distribute some of the package's files into the /windows and /windows/system directories in order to run, and unless these exist on your UNIX file system, those programs will not install correctly and probably will not run well, if at all.
- Wine is written to be file system independent, so MS Windows applications will install and run under virtually any file system supported by your brand of UNIX. Wine is also window manager independent, so the X window manager you choose to run has little or no bearing on your ability to run MS Windows programs under Wine.
- A Wine workstation will work with 16 MB of RAM and a 16 MB swap partition as long as you have a reasonable graphics card.

Red Hat Package Manager (RPM)

The RPM Package Manager (RPM) is a command line driven package management system which is capable of installing, uninstalling, verifying, querying, and updating computer software packages. RPM enables users to take the source code for new software and package it into source and binary form so that binaries can be easily installed and tracked, and allow the source to be easily rebuilt. Each software package consists of an archive of files along with information about the package, such as its version, a description, etc.. The open packaging system is available for anyone to use, and works on Linux and UNIX systems. The current release, RPM-4.1, ships with Red Hat Linux 8.0, but has some known errors and incompatibilities. Distributed under the GPL, Red Hat also encourages other distribution vendors to have a look and use it for their own distributions.

For the end user; - RPM provides many features to aid in maintaining systems. Installing, uninstalling, and upgrading RPM packages are all one line commands to aid simplicity. RPM also maintains a database of installed packages and their files, this can then be used to perform powerful queries and for the verification of your system.

For the developer; - RPM allows you to take source code for software and package it into source and binary packages for end users. This process is quite simple and is driven from a single file and optional patches that you create. This clear definition of sources and your patches and build instructions eases the maintenance of the package as new versions of the software are released. There is also a related API that permits developers to bypass 'shelling out' to a command line and to manage such transactions from within a native coding language.

During my research on Linux for the small business environment I was also given the task of finding more in-depth information on was the file systems that are compatible with open source operating systems, in particular Linux.

The Linux OS supports the 'UNIX file systems', which encompasses a few different file systems, including FAT, FAT32, HFPS and NTFS (read only). As Linux aims to improve its appeal, the Linux programmers who are continually working to update and improve Linux have put into the operating system compatibility support for most of the other operating systems out there. Newer file systems such as JFS & similar file systems are now supported. However, some of these may require special drivers or utilities to be added to a basic Linux install.

Previously Ext2 has been the file system for most Linux machines. Described as robust, reliable, and suitable for most deployments, however the need for new file systems arose due to requirements that included large hard-disk partitions, quick recovery from crashes, high-performance I/O, and the need to store thousands and thousands of files representing terabytes of data. To try to meet these requirements Linux now has 4 alternatives: - Ext3, ReiserFS, XFS, and JFS. All 4 of these file systems make use of a feature called 'journaling'. A journaling file system can simplify restarts, reduce fragmentation, and accelerate I/O. Journaling file systems use transactions, much like a database transaction, and treats changes to file system meta-data and/or user data as a single, atomic operation. Before the file system makes those changes, it creates a transaction that describes what it's about to do, the *journal* part of a journaling file

system is simply a list of transactions. Using the concept of a transaction guarantees that either *all* or *none* of the file system updates are done.

Features of Journaling file systems include:

- Faster allocation of free blocks
- Large numbers of files in a directory
- Large files
- Faster recovery & restart times

Ext3

Ext3 is an extension to Ext2, designed to provide higher availability and add support for journaling, while maintaining the robustness of Ext2. Ext3 uses the same disk layout and data structures as Ext2, and it is forward and backward-compatible. As it is an add-on to Ext2, Ext3 has the same limitations such as directory implementations not scaling well to large file systems (there is an upper limit of 32,768 subdirectories in a single directory, and a 'soft' upper limit of 10,000-15,000 files in a single directory). Another issue is the fixed internal structures being too small to capture large file sizes, extremely large partition sizes, and enormous numbers of files in a single directory.

ReiserFS

ReiserFS, like the other journaling file systems, is open source & available in most Linux distributions and supports meta-data journaling. One of the unique advantages of ReiserFS is support for *small* files. The philosophy behind this is that small files encourage coding simplicity, so rather than using a database or create your own file-caching scheme, use the file system to handle lots of small pieces of information.

In comparison to Ext2, ReiserFS is about 8 - 15 times faster at handling files which are less than 1K. ReiserFS can also store about 6% more data than Ext2 on the same physical file system. ReiserFS utilizes a B* tree to manage all file system meta-data, *and* stores and compresses *tails*, portions of files smaller than a block. It has excellent performance for large files, but it's especially adept at managing small files.

JFS

JFS is based on IBM's JFS file system originally designed for the OS/2 Warp system, and is well-suited to enterprise environments. JFS uses many advanced techniques to boost performance, provide for very large file systems, and journal changes to the file system.

JFS uses extent-based addressing structures, along with aggressive block allocation to produce compact, efficient, and scalable structures for mapping logical offsets within files to physical addresses on disk. JFS dynamically allocates space for disk inodes as required, allowing it to free space when it is no longer required. With dynamic inode allocation, users do not have to estimate the maximum number of files and directories that a file system will contain as with Ext2.

Two different directory organizations are provided: one is for small directories and the other for large directories. This eliminates the need for separate directory block I/O and the need to allocate separate storage. The contents of larger directories are organized in a B+ tree and provide faster directory lookup, insertion, and deletion compared to traditional unsorted directory organizations. JFS is a full 64-bit file system allowing support for large files and partitions.

XFS

Based on SGI's Irix XFS file system technology, XFS supports meta-data journaling, and extremely large disk farms. A single XFS file system can be 18,000 petabytes and a single file can be 9,000 petabytes. XFS is also capable of delivering excellent I/O performance. In addition to its better scale and speed, XFS uses many of the same techniques found in JFS.

Other reports I have created for the 'Linux for Small Businesses' topic include:

- Linux Implementation Guide
- Red Hat Linux 9: Installation Guide
- Red Hat Linux 9: Hardware Compatibility Listing
- Linux Kernel History & Hardware Compatibility
- Requirements for a Linux Network Administrator

Chris' Research

Chris was assigned to research further into the topic of Hardware. Over the project period he also produced many reports detailing different aspects of hardware and various technologies. A brief summary of these topics and his findings are given below;

Intel® Centrino™ mobile technology

- Intel® Pentium® M processor
 - New micro architecture including power optimised 400 MHz processor system bus
 - 1MB power managed L2 cache for higher CPU performance
- Intel® 855 Chipset Family
 - Support for up to 2 GB of DDR 266/200 memory
 - USB 2.0 support
 - Optional integrated graphics
- Intel® PRO/Wireless 2100 network connection
 - Supports 802.11b wireless networks and is Wi-Fi compliant
 - Industry standard and extended wireless security support

Intel® Springdale Chipset (865)

Features:

Springdale will feature support for 533/667 MHz buses, four DIMMs supporting up to 4GB of dual channel DDR 266/333 memory, AGP4X/8X, and CSA for GbE as a communications port. It also supports:

- UDMA 66/100 and Serial ATA 150
- Eight USB 2.0 ports, and eight PCI masters
- Support Pentium 4 and Celeron 478
- Hyper-Threading Technology

Intel® Canterwood Chipset (875P)

Features:

Canterwood chipset will support 3.06 GHz+ P4 CPU on 800MHz FSB, Dual DDR 400 memory, Intel Hyper-Threading Technology, Intel Performance Acceleration Technology.

HyperStreaming

1. **Single Stream with Low Latency Technology** enables a single stream to respond in prompt action and reduce latency in PC system.
2. **Multiple Stream with Pipelining and Concurrent Execution Technology** with its integrated concurrent parallel structure effectively widens the bandwidth of data transmission and simultaneously processes non-sequential data to improve the performance of a system, bringing about more flexible operations of all peripherals.
3. **Specific Stream with Prioritized Channel Technology** distinguishes music and images downloaded from the Internet and allocate bandwidth and piping for listening or viewing to ensure smooth playing.
4. **Smart Stream Flow Control Technology** with the smart arbiter can bring about more efficient streaming and even expand the space of processor access through effective throughput control, based on different protocols and characteristics, when dealing with the external chipset interface.

ACPI (Advanced Configuration and Power Interface)

The principal goals of ACPI are to:

1. Enable all computer systems to implement motherboard configuration and power management functions, using appropriate cost/function tradeoffs.
2. Enhance power management functionality and robustness.
3. Facilitate and accelerate industry-wide implementation of power management
4. Create a robust interface for configuring motherboard devices.

Wired for Management (WfM)

- *Power Management*

Power management allows a system to consume less power, but still be fully operational in a short period of time. This enables low power consumption and improved end-user experience when starting up (really 'waking up') a system.

- *Preboot Execution*

The Preboot Execution Environment (PXE) is a protocol and interfaces by which agents can be loaded remotely onto the client to perform management tasks in the absence of a running OS. The Preboot Execution Environment enables the automation of a number of management tasks.

- *Problem Resolution*

Today, every management environment must have some form of trouble ticketing. However, a particular strategy usually only works with explicit cross-developments by individual vendors. This means that management applications usually cannot obtain timely and accurate information about specific problems because the information is transient and not when the problem occurred.

Suspend to Ram:

Information on system configuration, open applications, and active files is stored in main memory (RAM); while most of the system's other components are turned off.

DMI (Desktop Management Interface)

DMI is designed to be:

- independent of a specific computer or operating system
- independent of a specific management protocol
- easy for vendors to adopt
- usable locally. No network required

DMI provides a common path for technical support, IT managers, and individual users to access information about all aspects of a computer - including processor type, installation date, attached printers and other peripherals, power sources, and maintenance history. It provides a

common format for describing products to aid vendors, systems integrators, and end users in enterprise desktop management.

DMI's four components are:

Management Information Format (MIF) – Each component has a *MIF file* to describe its manageable characteristics. When a component is initially installed into the system, the MIF is added to the (implementation-dependent) *MIF database*.

DMI Service Provider (AKA Service Layer) - an OS add-on that connects the management interface and the component interface and allows management and component software to access MIF files. The service layer also includes a common interface called the local agent, which is used to manage individual components.

Component interface (CI) - an application program interface (API) that sends status information to the appropriate MIF file via the service layer. Commands include Get, Set, and Event.

Management interface (MI) - the management software's interface to the service layer. Commands are Get, Set, and List.

CI, MI, and service layer drivers are available on the Internet. Intel's LANDesk Client Manager (LDCM) is based on DMI.

Active Directory

Active Directory, new in the Windows 2000 operating system, plays a major role in implementing your organization's network. To locate an Active Directory server, an Active Directory client queries DNS. That is, Active Directory uses DNS as a locator service, resolving Active Directory domain, site, and service names to an IP address. For example, to log on to an Active Directory domain, an Active Directory client queries its configured DNS server for the IP address of the LDAP service running on a domain controller for a specified domain.


iSCSI

iSCSI is an Internet Protocol (IP)-based storage networking standard for linking data storage facilities, developed by the Internet Engineering Task Force (IETF). By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. The iSCSI protocol is expected to help bring about rapid development of the storage area network (SAN) market, by increasing the capabilities and performance of storage data transmission. Because of the ubiquity of IP networks, iSCSI can be used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet and can enable location-independent data storage and retrieval.

Vincent's Research

Vincent was given the topic of Networking and Connectivity to research throughout the duration of the project. He successfully completed a number of technical reports as well as a sub-topic on Setting up an ISP. An overview of his research is given below;

The path from 2G to 3G mobile communication



From the figure above, we can see that each country or continent has their own mobile cellular system, which they are still upgrading or improving from the previous system over years and years.

Specifications for AGP 8X, PCI Express, PCI-X and iSCSI

AGP 8X: AGP 8X is the specification for the next-generation AGP interface. AGP 8X has a new specification which allows the doubling of speed to 533 MHz and supports a data rate of 2,133 MB/s. AGP 8X also introduces isochronous (time-dependent processes where data must be delivered within certain time constraints) operation and texturing abilities. Unfortunately, AGP 8X has reached the end of the road where no upgrade to performance is possible anymore.

PCI Express: PCI Express' main goal is to improve the speed of transmission by using a serial architecture rather than parallel architecture as used in PCI and PCI-X. PCI Express also tries to be backward-compatible with all other current PCI technologies.

PCI-X: PCI-X has evolved to version 2.0 and is an improvement over the PCI-X 1.0 and PCI, generally in the speed of transmission and also in error checking and accuracy in data transfer.

iSCSI: iSCSI is developed to overcome the distance and device support limitations. Besides this it is used to connect up PCs, Servers, Workstations, etc to other peripheral devices such as printers, scanners, storage applications, etc by using the TCP/IP protocol.

Long distance (20km) Wireless PABX

Cordless access in the PABX environment offers two key sets of benefits. First, it enables users to retain the wide range of conventional PABX features while physically roaming across a business site such as factory, office complex and campus or even across a multi-site private corporate network. Second, it offers the potential for deploying full PABX functionality but without much of the expense, delay, and inconvenience of wiring up new premises.

Introduction to the 802.11 Wireless Network Standard

The original 802.11 wireless standard was defined by IEEE way back in 1997. 802.11 is the basic wireless technology standard upon which all variations are based. The variations in use so far include: 802.11a, b, g, c, d, e, f and i. Each of the standards has an important role in the wireless network. However, I will just briefly detail 802.11 a, b and g.

802.11a: 802.11a represents a significant increase in transfer rates with a maximum theoretical speed of 54Mbps. 802.11a is operating in the 5.8 GHz Unlicensed National Information Infrastructure (UNII) spectrums with the implementation of Orthogonal Frequency Division Multiplexing (OFDM) which provides 12 non-overlapping channels.

802.11b: Heralded as a major leap forward in WiFi technology both in regards to ease of use, implementation flexibility and relative cost. 802.11b has a maximum speed of 22 Mbps and operating in the 2.4GHz ISM spectrum with the implementation of Direct Sequence Spread Spectrum (DSSS) which provides 11 channels with 3 non-overlapping.

802.11g: 802.11g is the current IEEE standard in question. As of the writing of this paper, it still has yet to be formally approved and ratified by members of the IEEE. 802.11g finds itself in a rather unique position as it has similarities and ties to both 11b and 11a. 802.11g has a maximum speed of 54Mbps even though it is running in the 2.4 GHz ISM spectrum. This is because 802.11g is implementing Complimentary Code Keying (CCK) and OFDM modulation scheme which provides only 3 channels of 1, 6 and 11.

802.16a Technology Development and Commercial Deployment

802.16 was developed by the IEEE's 802.16 Working Group on Broadband Wireless Access (BWA) to address the need for broadband connection solutions that are more economical than wired alternatives. The working group completed the BWA standard in October 2001 and published it in April. The most appealing aspect of the BWA technology is its ability to quickly expand a broadband network to customer sites that a conventional fibre network can't reach. It works by leveraging installed high-speed optical networks, combined with radio base stations mounted on buildings or towers. Moreover, the 802.11 wireless standards have driven the need for 802.16a in order to join up all the 802.11 WLAN.

VOIP cell phones

Voice over a WLAN network is a very attractive proposition to many vendors coming from the data side. In fact, some companies have suggested that mobile VoIP be the opportunity to gain wider VoIP acceptance in the fixed world. Intel is a good example. As part of its large and growing investments in WLAN, it is funding a VoIP startup, Telesym. In January 2003, Ericsson showcased a healthcare project, codenamed Guardian Angel, using IPv6 and roaming between 2G, 3G and WLANs. This was part of the EU-project IPv6 Wireless Internet Initiative (6WINIT). IPv6 has the huge advantages over the current IPv4, in that it provides a virtually infinite number of addresses allowing huge personalization opportunities.

Smartphone vs PDA

Smartphone is becoming very popular in 2003 compared to last year when it was first released. Likewise, the Smartphones based on Symbian OS produced by well-known cellphone companies such as Nokia, Sony Ericsson, Motorola, Samsung, etc. A lot of companies are creating software to improve the Smartphones such as PhotoBase, which is used to edit photos taken with the cellphone; MediaQ has launched a new graphics chip dubbed the MQ2100 Multimedia Platform Controller.

Setting up an ISP business in terms of technologies, equipments and skills required

This was a very large and tough research topic, and knowing where to start was a very important step. Therefore, understanding the Internet Architecture was the point chosen. Without knowing the architecture, you have no idea what and where an ISP is supposed to connect to.

Below are a few sections of the Internet Architecture:

- User PC – Multimedia PCs equipped to send and receive all variety of audio and video.
- User Communication Equipment – Connects the Users' PCs to the "Local Loop".
- Local Loop Carrier – Connects the User location to ISP's Point of Presence (POP).
- ISP's POP – Connections from the user are accepted and authenticated here.
- User Services – Used by the User for access (DNZ, Email, etc).
- ISP Backbone – Interconnects the ISP's POPs, and interconnects to ISP to other ISP's and online content.
- Online Content – These are the host sites that the user interacts with.
- Origins of Online Content – This is the original "real world" sources for the online information.


The understanding of the Internet Architecture leads to the internal operations of the ISP. Building an ISP requires a lot of equipment and technologies, dependant on the size of the ISP. However, an ISP must have a RADIUS server to identify its customers; a large amount of Routers to route many different users to many different destinations, each under the direction of a single autonomous system; Billing and Accounting Software to keep track of their customers' usages; DNS server to match Internet IP addresses with site names and so on.

After researching it was found the 5 popular ISPs in NZ; IHUG, XTRA, CLEAR, MAXNET and WORLDNET. Research was done on the types of connections they have provided, other services provided to their customers, installation and monthly charges, number of customers for each ISP and some extra information unique to the ISP.

Finally, research was done on the bandwidth purchasing and billing software, primarily with Orcon Company. A rough idea of calculation for bandwidth required for all the services provided by an ISP was learnt from Orcon. As for the billing software, most of the ISPs have their own written piece of software, tailored specially according to their needs. However, it is also possible to purchase billing software from other ISPs.

Knowledge Base

As previously mentioned, the second part of the project involves the continuation of the Knowledge base, which was started by the BTech project team in 2002. The objective of the Knowledge base is to provide a central base in which to store and manage relevant information and resources for use by IT professionals and their clients. The idea behind this is to allow IT consultants within the company to have access to resources that will enable them to make informed judgements and recommend the most appropriate technologies or solutions to their clients.



The Knowledge Base consists of a MySQL based back-end database combined with an Apache based front-end web site with functionality added through the use of PHP.


The Knowledge base consists of a MySQL based back-end database, currently using MySQL version 3.23.51, where data relating to users, projects and the company is stored. This is also where the technology reports we created in the first part of the project are to be stored on completion of the knowledge base.

The front-end is an Apache based web site, using SSL to provide secure communication between the client and server. The Knowledge base is currently implemented using Apache version 2.0.39. Interaction between the Apache Web server and the MySQL database is provided through the use of PHP, in this particular implementation, PHP version 4.2.2.

All the files reside on an Intel Celeron 300MHz machine, which is located at the ISP. The server machine is running the open source Linux Red Hat operating system version 7.3 with Kernel version 2.4.18. The original intention was to make this a completely open source project, and it has so far been achieved.

Hardware specifications of server:

- Intel Celeron 300A
- 128MB PC133 SDRAM
- Asus P2B-F
- Matrox G200 8MB
- Seagate U6 20GB
- Accton EN1207D-TX



The current implementation uses a 3-tier, client/ server style architecture, with the clients accessing the web server through their own web browsers using HTTP. As previously mentioned, the functionality to the front end is provided through the use of PHP code and the web server accesses the Relational DB using MySQL.


After an initial evaluation and meeting with one of the former project team members, we found the initial status of the Knowledge base to be the following;

Database	Prototype complete
Site Front-end design	Prototype complete
PHP Coding	Incomplete – 50% remaining
Documentation	Incomplete – 70% remaining
Knowledge/Data	Not Started

The implementation of the MySQL database and the Knowledge base front-end design was practically complete. We were given the opportunity to redesign the UI, however we found that the current interface was sufficient and did not require much modification. However, it has been stated that the initial implementation may require some modifications to accommodate the information intended for the Knowledge base.

Documentation was another of the main areas that needed to be addressed. Initially only the MySQL database was sufficiently documented. The PHP code has been commented appropriately, however further work was required in documenting how the PHP pages have been coded, such that it could be easily understood and modified by other programmers as required. Documentation and help pages also needed to be created to provide first time users with information on using the Knowledge base.

We started by looking at the supervisor level and decided we would need to change the structure of the actual KB to better accommodate knowledge base articles, and to better define the access for the three levels. Also by starting from the supervisor level we were able to reuse the code, with some slight modifications, for the other 2 access levels later in the implementation. The initial KB used a file called dataclasses.inc, which contained all the necessary classes, and was included in each subsequent page in the KB. This modular approach provided a good way of keeping the code easily usable and modifiable, so we continued in this manner for consistency and simplicity.




One of the main issues regarding the Knowledge base, was how to structure the Reports section so that the information will be displayed in an appropriate format, one that is helpful to the user. One of the major changes we implemented involved the restructuring of the Project and Technology sections. These 2 areas were largely incomplete and needed quite a bit of work. Another of the changes made was to link the Reports section back to Project Reports, this was to provide better access to the stored knowledge. We also found that a lot of functions within the page had not yet been implemented, and many functions such as the search functions and update functions were not always working correctly.


After we had decided what needed to be done, we broke it down further and divided the work according to sections, allocating the work amongst the 3 of us. I was allocated the Project Management section and its related subsection. Chris was assigned the Staff & Client management section and to implement the News section. Vincent was assigned Supplier, Technology & Reports sections to work on.

Working predominantly from home, we spent about 10 hrs each on our assigned areas per week, reporting back to Edmond as required. We also had regular meetings with Edmond where we collaborated and combined our code on the central server located in Helensville. During this time we made use of a Gantt chart to document and monitor our performance, this was discussed and updated at each meeting to give the team a clear indication of our progress.

CNZ Knowledge Base Diagram v2.2




The following diagrams show the structure of the Client and Staff access level interfaces and their functionalities. Pages that were implemented are shown in the orange brackets and functions that were implemented or changes that we decided upon are displayed in orange text. Also included is a brief discussion of the changes we decided to make to the original KB during the second part of the project.




For the Client access level we revised the original structure so that all functions relating to Client support will be implemented in one page. The Site Map and Help pages were yet not implemented from last year, so we were required to implement this for all three user-groups. We also added functions for displaying and searching client news, and modifying the way in which Staff details are currently displayed. Initially the Staff Details page allows a user with Client access level to modify details, this needed to be changed so that modifications cannot be made.

Overall the Client access level was mostly complete and generally only required coding to make pages consistent and easier to use, for example adding 'Back' buttons for easier navigation.




For the Staff access level interface we need to implement the Site map, News and Help pages as with the Client access level. Our aim was to continue the modular approach taken by the previous BTech group so that the code can be reused throughout the 3 separate user interfaces. We also decided to restructure the Projects page in order for the navigation to be more logical. A number of extra functionalities relating to adding and updating reports also needed to be implemented. During the coding we found that we needed to create another page for Supplier services to provide further separation between the different services offered.

Knowledge Base - Project Section



Main Projects page

For my part of the coding I was assigned the Project Section of the Knowledge base to complete. This turned out to be quite a large section, as it required a number of pages and functions to be implemented. The following gives a brief overview of some of the pages and changes that were implemented during the coding and screen shots of the final pages.




Project Details page

The Project section consisted of a list of projects that gives the option of viewing all projects or only the projects undertaken by the current staff member (as shown above). Clicking on one of the project hyperlinks on the Project page takes the user to the Project Details page, in which


more details relating to the project can be found. The Project page and the main Project Details page was initially complete with little or no extra work required. However the sub-categories, Project Reports, Project Staff & Project Invoices, on the Project Details page still required implementation.

Project Report pages




The Report page gives a list of reports available for the selected project. There are 2 different types of reports that can be stored about a project, Status and Research. The Research reports has been modified to hold the reports we have created throughout the R&D part of the project. The Status reports are less detailed, giving a description relating to the status of the project. Clicking on one of the report hyperlinks displays the full details of the information contained in the report, as retrieved from the MySQL database.

For the reports page I was required to implement the functions to enable the user to add a new Status or Research report, as well as update any current reports. To do this I was able to reuse existing classes from the `dataclasses.inc`, as there were adequate similarities between the format for the Status report and Research reports.




Add new Status report



Status report successfully added to database


Project Staff pages




Project Staff page

For the Project Staff pages I was required to implement functionality to allow users to add or remove staff from a selected project. After trying various options it was decided the best way to implement this would be using check boxes, this would also allow multiple selections. Processing multiple selections proved to be a fairly challenging task and required some further research both within the group and on the Internet.

To solve this we created a new class in the dataclasses.inc, which extended the existing Staff class. This enabled us to easily add Project Staff with a single call in the code, with all the details to update the database accordingly encapsulated in the data class.




Add Project Staff page




Staff successfully added to the current project

Once the add Staff function had been completed, the principles behind the remove function was fairly similar, although the actual implementation of the remove function did require some time due to difficulties in updating the database. This was easily resolved by inserting a statement to check the access level was sufficient to modify project details.




Remove Project Staff page



Selected Staff successfully removed from the project


Project Invoices pages



Project Invoices page


The Project Invoices pages were yet to be implemented, as we discovered after the initial evaluation of the KB. The Project Invoice pages needed to hold information relating to any transactions that were associated with the project. To do this I first had to create a new table in the database to accommodate this information. This was called the Project_Supplier_Invoices table and resolved the many-to-many relationship between the Project and Supplier tables. Once this was created, I created a new data class in dataclasses.inc to define all the necessary database functions. The SELECT function proved to be the most difficult as the identification for the Project_Supplier_Invoices table was defined by 3 primary key values. As there was often the need to search for an invoice using various criteria, we implemented 3 different methods of selecting Project Invoice records.

To create the actual pages I used the same layout and structure as the rest of the pages, this was done for consistency and ease of use. Clicking on the description hyperlink shows the full details of Invoice.




Project Details page

In creating the Project_Supplier_Invoices table we set the Supplier GST Number and Date Paid fields such that they could be null. This meant that an invoice could be added when it was received and then updated with further information once it had been paid. The update function was a bit more difficult to implement than the others, again due to the 3 primary key values.



Add Project Invoice page



Project Invoice successfully added to current project

Some extra coding was required to complete the KB so that it kept a consistent and easy-to-use interface, such as the addition of Back buttons.

Knowledge Base final status

Database	Prototype complete
Site Front-end design	Prototype complete
PHP Coding	Completed
Documentation	50% remaining
Knowledge/Data	Completed

Through liaising with each other and Edmond we have managed to complete the Knowledge base to a stage where it is ready for use. There is some further work that will be required before it can be fully implemented for the company's use, such as security features and thorough testing. In terms of documentation, the KB has not yet been fully documented, however we did try to ensure that our code was properly documented so that future users can benefit. For the Knowledge/Data requirement we have also created many reports, between the 3 of us, on various topics to provide useful information to populate the KB. We have also modified the KB so that it will accommodate this information sufficiently.

Overall, this project has given us the opportunity to apply programming skills and concepts learnt throughout our degree. We also managed to keep to one of the original aims of the project, which was to create a completely open source implementation. As we were using open source software, a lot of research was both on the Internet and through advice from other programmers. We also learnt a lot from each other, firstly by attempting to solve the problems on our own, then referring back to the each other and using the most efficient methods.

Conclusion

We have now reached the end of our BTech project in which we were acting as part of the Research and Development team for CNZ, collecting & collating information and writing discussion papers for use within the company structure. The deliverables produced, as outlined in the Project brief, include various research reports and Industry Channel presentation slides. I have also produced reports based on our initial evaluation of the Knowledge base.

The Research & Development section has allowed me to apply my research skills learnt throughout the course of my degree and to further refine these skills. As the majority of the information on my assigned topics was sourced from the Internet, I was encouraged to try and determine the reliability of the source and material and also its relevance to the topic. As the Internet has no structured governing body, much of the information found within it can be biased or inaccurate. I feel that through this process I have gained a better understanding of how to find and classify reliable sources for writing reports. Writing the reports also proved to be an interesting challenge as there was often so much information available that I needed to be able to collate and summarise it all into one report.

For the Linux for Small Business topic, I have produced a number of reports for the Knowledge base, these included:

- Red Hat Linux 9.0 Product Details
- Red Hat Enterprise Product Details
- Red Hat Linux 9.0 System Requirements Table
- Red Hat Linux 9.0 Installation Guide
- Linux Emulation software v0.2
- Linux Accounting Software v0.4
- Hardware Compatibility List
- Linux Kernel History v0.1
- Linux Implementation Guide
- Open Source File Systems
- Linux Networking Report v0.4
- Linux Network Administrator requirements

From this I have learnt a great deal on the Linux environment as well as learnt about other open source tools available and their benefits. From the extensive research I have done on the Linux operating system and its implementation, I feel I have gained the theoretical knowledge that would allow me to pursue a role as a Linux Network Administrator.

The second phase of the project focused mainly on the continued implementation and documentation of the CNZ Knowledge base structure. This part of the project allowed us to apply the programming theory we had learnt during our Btech degree. We had each been assigned an area of the Knowledge Base to work on, for myself it was the Project section. I found that it took me a little while to understand the code written by the previous team, but once I had refreshed my knowledge of PHP I was able to get into the coding. The code was quite well structured and documented so it was easy to follow after that. Combining our PHP, MySQL, Apache and Linux understanding, together with the help of Edmond and some external resources, we were able to complete the Knowledge Base.

The project has also given us a good respect for teamwork and time management. Liasing with Edmond and the rest of the team has given me a good understanding of group dynamics and how to co-ordinate members within a project. I also learnt how to track and schedule our progress with the knowledge base through the use of the Gantt charts.

Knowledge and skills I have gained from University and have been able to apply to the project include the following;

- COMPSCI 230 - OOA and OOD design, Software architectures, Encapsulation, documentation, Software design/development
- COMPSCI 210, 314 - Network architecture, Unix platform, Linux concepts
- COMPSCI 334, 335 - PHP, Apache, Internet based applications, HTML, Remote/Socket programming, database connections
- COMPSCI 708 - Web page design and theory, Usability, Development methodologies, Web Testing methods, reliability of the Internet
- INFOSYS 222 - Database structure, MySQL commands, ERD diagrams
- INFOSYS 224, 329, 339 - Network devices & connectivity

APPENDICES

Knowledge Base Initial Assessment Report

Meeting: Thursday 1st May, 6:00pm
Place: The University of Auckland, City Campus – Rm 303.499
Attendees: Edmond Chan, Vincent Yee, Chris Chou, Anisha Ramji

Objectives

- To inspect and evaluate the current state of the Knowledge Base infrastructure
- Plan a course of necessary actions for continued development of the Knowledge Base structure.

Results

After meeting with Edmond, we learned how to implement the Knowledge Base from the files given on the CD. The step-by-step details of this implementation procedure are to be written up in a separate document at a later date.

The Knowledge Base has been divided into three categories; Client, Staff & Supervisor. With each category offering a differing level of access and modification rights. The Client side seems to be completed, with a few areas in the Staff and Supervisor levels still requiring implementation or re-checking of the code. As discussed with Edmond, some of the code will be repetitive and once written can be applied to other areas with a bit of careful renaming.

The MySQL database and user interface side seems to be completed to a high standard and it does not appear that any further work would be required in that area.

Documentation appears to be the area that will need the most work. This will include the checking and possibly modifying the current ERD diagrams and supporting documentation, as well as the checking and commenting of the source code. The structure of the files has been done so that the main code is contained in one PHP file and is referred to in all subsequent files. This will definitely be helpful in the learning and documenting process.

As previously discussed by Edmond, the structure of the Knowledge Base may need to be changed slightly to accommodate our reports. At present one issue is how to display documents that spread over more than one page. This will need to be addressed before any of the reports can be entered into the KB, depending on the way in which we decide to link the documents.

From the initial meeting we discovered the following areas that will we need to come back to;

- Site Map
- Search function
- Uploading reports function
- Updating reports
- Forum

Strategies

We need to go through all the links within the KB and thoroughly test all aspects to ascertain specifically what areas need to be modified, and what is required to do this.

Forum needs to be re-implemented, this will require downloading the necessary file packages from the PHP BB website (www.phpbb.com)

Some changes to the PHP coding may be necessary, this will first require us to read and fully understand the existing code. Some PHP files still require comments to be added as well.

Documentation will need to be created and modifications made to existing documentation.

Conclusion

Overall we found the Knowledge Base so far to be quite impressive and well implemented. As previously discussed by Edmond, the Knowledge Base overall is fairly complete, with a few areas still to be implemented and modified. We have decided to each try to implement the KB on our own machines and meet again at later date, when we can further discuss with Edmond any questions we may have. Also to determine what specific areas need to be addressed and what action is required.

[illegible]

Microsoft Greenwich for Enterprise Instant Messaging

Classification

Greenwich is the latest RTC software from Microsoft, designed to provide secure and manageable real-time communication, including instant messaging, video conferencing & voice communication, with a strong emphasis on presence.

Research Objectives

To gain an understanding of:

- What the technology is about (functional specification & technology standards)
- How it fits in with other complementary technologies
- How it compares to other competing technologies
- Evidence of its promise

The Technology

Greenwich is the latest in Real Time Communications (RTC) software from Microsoft, due for release mid 2003. It was designed as a central means to manage variety of Web-based communications, including instant messaging, video conferencing and voice communication, within a business. It aims to incorporate a new era of communications with the emphasis on presence, being able to ascertain whether a contact is online and available, for time sensitive issues.

Greenwich aims to bring four additional key elements to the instant messaging scene: security, manageability, standards based architecture and extensibility. These elements are addressed below:

Security

Security is one of the primary concerns that businesses have with employees using free public Instant messaging services, as confidential company information and details may be compromised during online collaborations. Greenwich addresses this issue with integration with Microsoft's IT security management structure. Microsoft are also currently developing technology to protect unencrypted text as it travels across the internet and provide standards support for existing technologies such as SSL encryption, Digest and NTML/Kerberos encryption.

Manageability

In recent studies, 84% of companies reported that their information workers use instant messaging, most which are private applications which are not endorsed by the company's IT department. Greenwich aims to incorporate manageability and security through integration with the Windows server family, allowing system administrators to deploy an Instant messaging solution which is managed using the same tools as the existing Windows infrastructure.

Standards Based Architecture

Microsoft released its first enterprise instant messaging product as part of Exchange 2000, since then they have looked at improving this by making the platform extensible

and allowing it to be used in other applications, other than just the core application. To do this they have based the new product on an industry standard protocol called SIP, Session Initiation Protocol, (and its related protocol SIMPLE – SIP for Instant Messaging and Presence Leveraging Extensions) which is approved by the Internet Engineering Task Force (IETF). This provided the base for presence, one of the major features in Greenwich, as well as the ability to support API's and real-time content including text, speech, video and files.

Extensibility

Extensibility is another main feature in Greenwich, as it was designed to be a robust communications base, which can be further developed to integrate with other services. Microsoft predict that Line of business integration will be one those vital points of evolution, especially in areas such as Customer Relationship Management and other such systems that rely on real-time communications as their main information delivery means.

Complementary Technologies

Microsoft chose to implement the Greenwich platform directly on the Windows platform, as they believe this will allow it to be widely exploited and familiar to a large user base. Part of their RTC implementation process included the release of MSN Messenger Connect for Enterprises, this allows businesses to extend MSN Messenger and Windows Messenger beyond the corporate firewall to customers and consumers, and allows administrators to manage the solution within the familiar Windows infrastructure. MSN Messenger Connect enables the management and logging of instant messenger conversations between the enterprise user and customer to help protect company privacy and intellectual property, also for regulatory compliance in certain industries. Greenwich will interoperate in a standards based environment with MSN Messenger via MSN Messenger Connect, Windows Messenger and any other communications applications. Microsoft aims to make real-time communication and collaboration part of every application and intend to integrate that aspect into all the Microsoft Office applications. Being standards based, Microsoft hopes to create partnerships with other Systems integrators and developers who can provide their own solutions by building on top of the Greenwich capabilities. To date, areas where there has been strong interest in developing include: integration with telephony applications, line of business applications, information worker productivity applications and conferencing applications.

Another of the Microsoft RTC platforms due for release soon is Microsoft's Windows Messenger 5, which targets enterprises and is based on the beta version of Greenwich. Windows Messenger 5 includes simultaneous support for the Internet based .NET Passport, SIP Communications Service and Exchange Instant Messaging.

Competing Technologies

Other major instant messaging networks, ICQ, AOL Instant Messenger (AIM) and Yahoo Messenger, which currently do not interoperate with one another, are also looking into moving towards a SIP based approach. Microsoft believes that the SIP/SIMPLE based approach sets Greenwich apart from the rest as their competitors are using gateways, these pose a performance, compatibility and manageability issue for IT managers.

Since the announcement of Microsoft Greenwich, Yahoo has also announced the release of its new instant messaging client, Yahoo Messenger Enterprise Edition, due out in the first quarter of 2003. AOL has also announced its Enterprise RTC Client, 'Enterprise AIM', which will incorporate VeriSign security credentials and end-to-end encryption. It is hoped that this

competition will bring about the development of standards in Enterprise Instant messaging, particularly in relation to logging, auditing and encryption.

Evidence of Promise

Greenwich is currently still in the beta testing stages, but Microsoft reports that the feedback and figures from Microsoft Exchange enterprise instant messaging client and similar products such as Windows Messenger & MSN Messenger, have been favourable. According to recent studies by IDC and Ostermann Research there are 200 million instant messaging users worldwide, this number is predicted to increase to 500 million by 2006. Of those, there are currently 75 million using MSN Messenger for their real-time communication and collaboration requirements.

Large international company, Reuters, in alliance with FaceTime, are currently using Reuters Messaging developed with the help of Microsoft and based on the Greenwich technology. Targeting the intense security and reliability needs of the finance industry, where the capability to log, track and audit all electronic forms of communication is a critical requirement. Reuters are currently implementing the solution as a fully hosted service, but envisage that it will eventually become a federated network of presence collaboration.

References

Microsoft Website – www.microsoft.com

- *This website provided much in-depth technical information on the product and was very useful, I would rate this resource as excellent. However, in terms of opinions, some of the information may be slightly biased as it is the website of the creator of the product.*

eWeek Website – www.eweek.com

- *Information found on this website was useful in gaining the consumer opinion on the product from a non-biased point of view, I would also rate this resource as very good.*

InfoWorld Website – www.inforworld.com

- *Provided useful information relating to real world application and opinion of the product, a good reliable source*

Instant Messaging Planet.com – www.instantmessagingplanet.com

- *A good resource providing much detail on the product and similar products available. I would rate this resource as very good, as it provides much information from an informed view.*

Open Source Jabber IM XML Protocol for P2P data exchange

Classification

Jabber is an open source XML protocol for the real time exchange of messages and presence between any two points on the Internet. The XML platform is the basis for peer-to-peer collaboration.


Research Objectives

To gain an understanding of:

- What the technology is about (functional specification & technology standards)
- How it fits in with other complementary technologies
- How it compares to other competing technologies
- Evidence of its promise

The Technology

Jabber is an open source XML protocol for the real time exchange of messages and presence between any two points on the Internet. It consists of an extensible instant messaging platform and an instant messaging network that offers functionalities similar to that of AIM (AOL Instant Messenger), ICQ, MSN and Yahoo, but with several advantages. The first is that it is open-source; the Jabber protocol is free, publicly available and easy to understand. There are currently many open source implementations of the Jabber servers and clients, as well as development libraries to aid users in extending the protocol to suit their needs. The protocol is highly extensible through the use of simple, XML namespaces, allowing the user to easily add custom functionality and interoperability. Another advantage is that the technology is decentralized, enabling individuals or organizations to run their own Jabber servers. This control over the server adds to the security of the technology, as the server is able to be isolated from the public Jabber network. The Jabber client also supports uni-directional or bi-directional subscription to presence to allow users greater privacy and security. These factors along with the interoperability of Jabber have made it well suited to the Peer-to-Peer (P2P) scenario.



The XML platform is the basis for peer-to-peer collaboration, creating the interface for sharing enterprise computing resources in the P2P environment. This is done through the combination of XML and Simple API for XML (SAX) to create a way in which to pass data and call remote processes. Initiated on the server and client, it then enables real-time collaboration and integrated workflow. When implemented in this manner, servers can also be treated as Peers, allowing collaboration between them, just as with clients, and increasing the scope of peer-to-peer communications. When combined with other technologies such as Rich Site Summary (RSS), communication & collaboration process can be further enhanced.

As interoperability is a key factor, Jabber is currently working with the Internet Engineering Task Force (IETF), along with a consortium of other professional groups, to develop a standard that addresses issues such as security and internationalisation.

Currently Jabber supports SSL between client and server communications and SASL (Simple Authentication and Security Layer) as the preferred authentication method, although this protocol is still in draft status and has not yet been fully implemented by all Jabber server implementations. During the sign in process the Jabber client integrates with various network directory and authentication services such as Active Directory and LDAP.

Jabber clients use simple TCP sockets to exchange XML documents with the Jabber servers. The Jabber servers can be extended to act as gateways to other Instant Messaging services as well as to e-mail and telnet services. As it based on the XML platform and TCP sockets, any device that is able to access a TCP socket has the potential for Jabber access.

Complementary Technologies

The main advantage of Jabber is its interoperability between other IM systems, such as AIM, MSN, ICQ and Yahoo. This is achieved by the Jabber protocol 'masquerading' as the foreign protocol, acting as a translator between the two systems to provide seamless integration from the users perspective. Although it is not currently fully compatible with all other IM systems, its XML base allows for the protocol to be extended in order to accommodate other such systems and applications.

Jabber client are currently available for many platforms including Linux/Unix, Windows, Macintosh and several others. Web clients that include other technologies such as Java applets, Mozilla integration and HTML & JavaScript are currently in progress, with the intention of enabling Jabber communications over HTTP. Future expansions include Jabber access for handheld devices, mobile phones and other embedded systems.

Competing Technologies

Since the early days of Instant messaging, beginning with ICQ, there have been many companies entering the IM market, most of which have implemented their own proprietary servers and clients that only operate with like-clients. It has been reported that the main IM corporations – AOL Time Warner, Yahoo and Microsoft, currently divide the majority of instant messaging users between them. The Jabber protocols' current interoperability between these legacy IM systems along with its XML platform for easy extensibility sets it apart from its competitors.

Evidence of Promise

Since the Jabber project was created by Jeremie Miller in 1998, the protocol has continually been evolving. The Jabber Open Source project now has over 150,000 server implementations, Jabber libraries for a wide variety of languages and Jabber clients for a vast range of platforms. Jabber services have been deployed at thousands of domains on the Internet and on private intranets, and it is believed to have over three million Jabber instant messaging users worldwide.

The rate at which Jabber services are being adopted has been compared to the trajectory of email services, and with the integration of new technologies, such as the integration into the P2P computing scene, this growth is expected to continue.

In the corporate world, Jabber has been deployed in the telecommunications, enterprise and software development markets by major international customers that include: Hewlett-Packard, BellSouth, AT&T and Landmark Graphics.



Left: Jabber Messenger Client

References

Jabber Inc. – www.jabber.com

- *The proprietary website provided much information on the protocol and its application. I would rate the reliability of this resource as excellent.*

Jabber Software Foundation – www.jabber.org

- *This website was useful in gathering information on the product from a non-commercial point of view, I would also rate this resource as very reliable.*

OpenP2P.com – www.openp2p.com

- *A good resource providing a good level of detail on the product and its extensions. I would rate this resource as very reliable, as it provides much information and is from a well-established and reputable organization.*

internetnews.com – www.internetnews.com

- *Provided useful information relating to real world application and opinion of the product, a good reliable source.*

Microsoft Desktop Site for Windows XP Tools

Classification

The Microsoft Desktop Deployment Centre site contains a collection of references, tools & services, as well as outlining best practises and training for migration to the Windows XP operating system.

Research Objectives

To gain an understanding of:

- What the technology is about
- How it fits in with other complementary technologies
- How it compares to other competing technologies
- Evidence of its promise

The Technology

The Desktop Deployment Resource Centre was created by Microsoft to help IT professionals evaluate and deploy the Windows XP Operating system and its latest release of the Office applications, Office XP. The aim was to provide the decision makers and IT specialist within the business with resources to aid in evaluating, deploying and maintaining the Windows XP and Office XP. It includes resources such as TechNet, white papers, guides, services offerings, knowledge base articles, third-party tools and best practises & advice from those that have already implemented the system. These resources are indexed by product, company size and stage in the migration process to put the most relevant information at the users fingertips, during whichever stage of implementation.

Some of the Tools available in the Desktop Site include:

Microsoft Application Compatibility Analyzer 1.0

The Microsoft Application Compatibility Analyzer uses a command line tool called the Collector to gather information regarding applications currently installed on the machine. It then creates application inventories, organizing the information, generating reports, and analyzing application compatibility data. This allows the organization to prioritise in-house software testing.


Windows Application Verifier 2.5

Used during the application development and testing process, the Windows Application Verifier is designed to identify potential application compatibility issues. Included are tests for compliance with various requirements of the Designed for Windows XP and Windows Server 2003 Logo Programs, as well as tests for issues with heap & memory management and common security issues.

Compatibility Administrator

The Compatibility Administrator is used for selecting and applying compatibility fixes to applications that have been found to be incompatible. It allows the fixes to be grouped together into a compatibility database, which can then be deployed to other computers.

The Desktop Centre site also includes a new version of Microsoft's Baseline Security Analyzer, which scans Windows systems looking for common system misconfigurations and missing security updates.



Complementary Technologies

In addition to the online resources Microsoft have also been offering courses and free online Web seminars to provide project guidance, technical training and outline resources available to assist in upgrading. Aimed at complementing the information available on the Desktop Centre site, the Desktop Deployment Web Seminars consist of weekly, one hour, interactive events that cover deployment technologies, tools and best practises from Microsoft professionals. The courses, held throughout the US, address key tasks during the evaluation and deployment stages, various tools & technologies available, and areas for further study.

The Desktop Centre site also provides links back to the main Windows XP and Office XP sites for users to download the latest updates and security fixes.

Competing Technologies

One of the main issues that most businesses have had concerns over is that old applications would be incompatible with the new platform. Microsoft reassures that around 95% of applications will still work fine with the new operating system and the Desktop Deployment site helps in assessing which of the current applications may cause problems. For the small percentage that are not compatible the site offers multiple tools, including generic fixes, to help overcome these issues.

Evidence of Promise

One of the aims of the site was to drive a greater adoption of the Windows XP operating system and with the number of Windows XP users worldwide, it seems this has been successful. Although it is not clear how many hits the site has received, many business and organizations have reported it be a very useful resource aiding in the process of upgrading to Windows XP and Office XP.

References

Microsoft Desktop Deployment Centre –

<http://www.microsoft.com/resources/desktop/deployment/default.asp>

-The website in discussion, it provided very detailed, useful information as well as downloads and links to extra resources. As it provided by the creator of the product, I would rate this resource as excellent in terms of reliability.

Microsoft Website – www.microsoft.com

- This website provided much in-depth technical information on the products and was very useful, I would rate this resource as excellent. However, in terms of opinions, some of the information may be slightly biased as it is the website of the creator of the product.*

C Net News.com – news.com.com/2100-1001-985820.html

- Provided useful information relating to real world use and opinion of the product, I would also rate this resource as very reliable.*