A white paper regarding

A Data Repository

built on ANNTS

A discussion document and information leaflet for a workable petabyte capable database system.

The State of the Art

Data Warehousing has been a growing area of interest in the Information Communications and Technology (ICT) area since the late 1980s and early 1990s. Early systems using Terradata equipment were able to store and analyse terabytes of data, to reveal patterns for market intelligence and to improve design of supply systems.

NCR's purchase of Terradata in 1991[1] and IBM's entrance into the market signalled the emergence of data warehousing as a major industry segment. Since that time, technological advances has reduced the price of storage and processing capabilities and driven dramatic increases in performance. The effect of Moore's Law[2], which predicts that computing density doubles about every two years, has seen a dramatic increase in capability.

The increase in processing power and storage capabilities has led to the creation of computers with thousands of processors operating in parallel and supporting thousands of hard drives. These systems have generally run UNIX style operating systems and store the data in relational databases. The data is indexed when stored and accessed by invoking theses relationships.

By the late 1990s petabyte[3]-scale databases were widely anticipated, with three separate groups racing for the prestige of being the first petabyte database operator, but by 2006 petabyte databases have still not appeared. The problem is that the software proved incapable of supporting databases of such a large scale and providing reasonable access for users. The software bottleneck appeared to allow storage or access, but not both.

---

[1] NCR. (2005). About NCR – history. Available [online]: http://www.ncr.com/history/history.htm. [23 May 2006].

[2] Intel. (2006). Moore's Law. Available [online]: http://www.intel.com/technology/silicon/mooreslaw/. [23 May 2006].

[3] A petabyte is 1000 terabytes, each of which in turn is 1000 megabytes, each of which in turn is 1000 kilobytes, each of which is 1000 bytes. A byte is a single piece of information, such as an alphanumeric character. This is equivalent is 1,000,000,000,000,000 bytes.

The Bottleneck

Microsoft has highlighted the problems with massive databases. They have stated that centralised systems are not suitable for storing and accessing the petabyte-level volumes required by some customers. Instead, they have focused on multiple smaller servers combining to provide customer access to data. However, Microsoft still maintains a server-centric perspective, with the servers processing the requests for data and returning the results.

The leading edge customers include:

- Cornell University with its social science research project, aiming at archiving the Internet for research purposes. Their system is currently scheduled to be upgraded to 0.5 petabytes
- The Library of Congress has a requirement for a 3 petabyte database. They have been working with Microsoft since 1997 to approach this objective, but have not yet succeeded.
- The NSA (National Security Agency) uses massive databases containing call patterns and voice intercepts from many areas around the world. The classified nature of the NSA's activities prevents an accurate assessment of their systems, but it has been estimated that they processes databases that combine to around 3 petabytes. However, it is unlikely that any single database exceeds 1 petabyte in size.

However, massive petabyte systems have value in many areas of society. Governments in general have value in large-scale databases for the purposes of security, taxation assessment, social services and legal services. The legal industry has created a market for the archiving of corporate data and the large-scale searching of data for discovery of documents and to determine state of knowledge at a particular time. The telecommunications, passenger transport and freight industries have interest in market intelligence systems to better target customers with specific products. The airport industry shifts vast volumes of luggage and incurs high costs relating to luggage tracking and retrieval.

Current costs

A data warehousing system consists of the hardware to operate the system, the operating system and the database management system.

One supplier of a competitive system, EMC, has recently publicised their pricing for the disk array system. Their DMX- can mount form 96 to 2400 drives in a single frame, enabling up to one petabyte of storage. EMC's list pricing for their DMX-3 system is US$250,000 for a 7 terabyte 96 drive system, as of January 2006 – equating to US$35,700 (NZ$57,000) per terabyte.

The rise of Linux and open source software has a significant impact upon the computing industry, and has capped some of the costs that dominated the earlier mainframe era. Many of the large scale systems are operating a variant of Linux and some form of relational database, with proprietary modifications in order to enable the wide array of processors and storage units. The software cost for the system often exceeds US$1 million.

These systems need to be integrated into the rest of the information infrastructure. As IBM puts it, "the era of the stand-alone data warehouse is now being superseded by mission-critical integrated OLTP[4] and decision-support systems"[5]. Typically, the data warehouse system would be connected onto the corporate TCP-IP network and would be accessed through front-end systems, which would take slices of data for analysis and reporting. Users now often use microcomputer-based tools such as Applix TM1[6] to enable business intelligence systems, drawing data as required from the back systems. The front –end cost would normally exceed US$20,000 (NZ$32,000) per user.

---

[4] OLTP is an abbreviation of "OnLine Transaction Processing".

[5] IBM. (2006). DB2 Data Warehouse Edition for Linux, Unix and Windows. Available [online]: http://www-306.ibm.com/software/data/db2/dwe/. [23 May 2006].

[6] Applix. (2006). TM1 Technology. Available [online]: http://www.applix.com/solutions/tm1technology.asp. [23 May 2006.

Our Alternative

Our systems draw upon research started in the mid-1980s. We foresaw the need for massive database systems prior to the emergence of data warehousing, and concluded that the current computing paradigm was inadequate for the task for creating large-scale data repositories.

Our technology is based upon storing, searching and retrieving data in a fundamentally different way. The result, we do not have the same limitations experienced by the traditional architecture model.

Enabling Hardware

Architecture

We are not restrained to current hardware formats. Our systems assume that they form part of a single network information system, and address information in the same manner whether the information s held locally or remotely. In this manner, we have moved away from the restrictions of the original von Neumann computing paradigm of 1944 that binds the ICT industry.

We separate our hardware requirements into three main elements:

1. Library Processor. This system maintains an index of its local data and responds to requests. The system is configured in rings to provide additional capacity, with additional rings creating exponentially greater capacity.

2. Communications Processor. This system connects the different sub-systems together into a single logical network. Redundancy ensures reliable communications across the entire repository.

3. Information Access Device. The device is used by end users to access the information using an HTML interface, such as a web browser. This system can be traditional equipment such as personal computer or cellular telephone, or one of our devices with our embedded software.

Scalability

Our architecture is theoretically infinitely scalable. There are no addressing limits or overheads from the use of the system. The bottleneck restraints are not significant within the scale required to provide a globally accessible information repository.

In practice, the primary issue is access to sufficient technology.   A ten–ring system, as an example, can include 59,000 hard drives – compared to a maximum capacity of 2,400 drives on the EMC DMX-3 system referred to above – and can support 14 petabytes with access by large numbers of concurrent users. The system is scoped to support over one million transactions per day.

The system is also scoped for further growth. For example, in the event of 14 petabytes being insufficient, an eleventh ring would enable over 200,000 drives to be mounted with a capacity of over 50 petabytes. There is no theoretically limit to the number of rings that can be mounted.

In fact, the practical restraint is the supply of components. The hard drive industry manufactures some 380 million drives annually, and the creation of a single data repository would need to be managed in order to avoid the disruption of global supply patterns.

Business

Economic Advantage

A rule of thumb is that any technology or service with a 30% or greater advantage over an earlier technology will displace that technology from the marketplace.

Our hardware architecture is particularly economic to implement. We can supply data warehousing suitable hardware and software installable into standard rack type mounts for less than NZ$10,000 per terabyte. This compares favourably with the $57,000 per terabyte of the EMC system referred to above. On that basis alone, we have a source of competitive advantage.

Further, we can mount multi-petabyte databases. The current database market incumbents, such as Microsoft, Oracle and IBM, are incapable of replicating our database capabilities.

Business Demand

There are major systems planned in Europe and the United States that cannot be implemented without our technology and represent multi-billion business opportunities. The ITU needs a new system to manage the next stage of the Internet's growth. The American FBI needs a system to assist in the identification of terrorist activities. The Library of Congress needs to place its archives online to enable access throughout Government and the United States in general. The international airports are being reorganised to ensure improved tracking of baggage and passengers. The Center for Disease Control (CDC) requires a massive system to monitor pandemic outbreaks. Similar systems are required by all countries, representing a large market for large-scale data storage and retrieval.

These major systems all have stringent requirements in order to effectively continuously operate a large-scale database system; including: (1) hot swap drive capability, to enable any disk to be replaced without shutting down the system or restricting operations; (2) the use of standard interfaces, to enable external systems to access the data without extensive software development and support; (3) remote allocation of resources, with TCP-IP-style mapping of devices to enable flexible management of resources; and (4) the use of commercial off the shelf hardware, to maximise cost-efficiency and to enable easier access to spare components as required.

Business Model

We will supply access to our data repository using a fee-per-service payment model. This model enables us to control the implementation of new data systems, and thus prevent retro-orientated ICT professionals from mis-implementing the technology.

Our charges are set at a level sufficient to ensure a satisfactory return on capital and to enable rapid growth with the intention of becoming a significant supplier in this market segment.

We will leverage our existing contacts in several Governments to focus on core markets. We are positioned to deliver on national identification cards in several countries with over 100 million card requirements. A second initiative in roading supports the management of roading services in several countries. The annual movements of tens of millions of vehicles can be monitored and road charges applied based on actual vehicle usage.

Sales Channel

The systems will be sold directly through sales agents managed through our 1to1 employee support system.

We are well-positioned to establish a 1000-strong sales force around the world over the next few months. This sales force will be focussed on selling data services leveraging off the data repository.

Use of Commercial Technologies

The systems are built from components that are available in substantial volumes. The use of commercial off the shelf technologies enables access to sufficient volume of component to both meet customer demands for new systems within reasonable periods of time and to provide replacement parts for maintenance of existing systems.

First Data Repository Implementation

We are currently implementing the first three rings of a New Zealand located data repository. This system is planned to be commissioned during May 2006, and will immediately exceed the Warehouse's 4 terabyte [7] storage capacity. This system will be used to mount internal information systems and to demonstrate support for a national identification card for an international customer.

The second round of enhancements will enhance the system to over 100 terabytes, exceeding Australasia's largest data warehouse (ZDNet claims that "the largest data warehouse in Australia – with 22 terabytes to 23 terabytes of data – is believed to be run by Telstra"[8]) and will be Australia's most powerful transaction processing system based on transactions per hour capacity.

The third round of expansion will enhance the system to one petabyte, thus exceeding all other systems currently in use anyway in the world.

7 iStart. (2006). Warehouse leads the way in Australasian data warehousing. Available [online]: http://www.istart.co.nz/index/HM20/PC0/PV21906/EX224/CS22460. [23 May 2006].

[8] ZDNet. (2005). What's driving data warehousing spend. Available [online]: http://www.zdnet.com.au/news/0,39023165,39193618,00.htm. [23 May 2006].

Three Layer Software Model

Our system operates on a three-layer model for enabling transactions. The ANNTS operating system provides the software interface to available hardware and enables other software to execute. The Database technology stores and retrieves data based on pages and builds information in response to requests. The TRANSAPP transaction engine enables secure transactions.

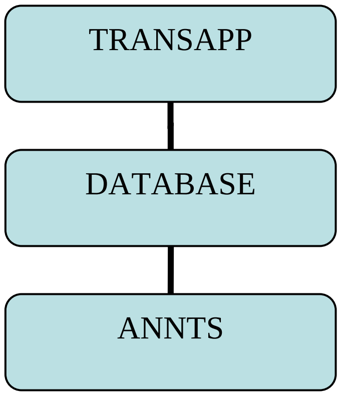

ANNTS

The fundamental capacity for the system is enabled by the ANNTS operating system, which is a proprietary technology available to Xastore.

ANNTS is hardware independent. The software can be ported onto any hardware platform with the scripting of hardware-specific drivers to enable the system to recognise the data inputs and outputs from any device.

ANNTS is a network operating system. Unlike earlier generation operating systems, ANNTS was designed on the basis of operating in a network rather than on standalone devices. The operating system can automatically share resources and tasks across a network in real-time, thereby using available hardware resources much more effectively than early generation technologies.

Software operating on ANNTS is registered before it can execute. This registration ties the software to its author and is apparent to the rest of the system. The ANNTS software model precludes the execution of malicious programs, as by definition, all programs must be registered and thus intended to execute before they can run. Thus, ANNTS has no exposure to viruses in the sense of the malicious software that historically affected early generation systems.

DATABASE

The Database system stores resources in any hardware available to the ANNTS operating system. The Database is not restrained to any one location, but instead has network storage and retrieval capability based on the requested data links to a requesting page. This page can be anywhere.

Further, the Database automatically spreads data around a large network. Data is replicated to improve data retrieval performance and to provide data security in case of the accidental loss of an individual hardware item or location.

Thus, the Database system does not require a data dictionary in the historic sense as the database operates as a single store of data across the entire network. There is no need for a dictionary linking database requests to locations or translating database elements as all data storage is standardised.

In addition, the Database technology stores data elements in original form. Information is created from data as requested by a page. There is no pre-processing of information, in the sense of fourth generation relational database systems which stored information in indexed files. Instead, the data is recompiled, thus freeing data systems from the need to accurately predict future information needs and thereby increasing the effectiveness of the system many-fold.

TRANSAPP

TRANSAPP (Transaction Applicator) is the transaction engine that is integrated into the Database system. It provides control over individual transactions and their status. TRANSAPP provides unique keys to the parties in a transaction to ensure security and enable each party to access its benefits without compromising the transaction security for the other party. The linking of keys provides transaction integrity plus roll-back capability, if required.

TRANSAPP relies on the Database and ANNTS technology to enable complete privacy protection for all parties, data integrity and a single world data view of the transaction.