Does My Infrastructure Look Big In This?

A review of Fujitsu's Infrastructure Services

Author: Robin Bloor
Published: April 2002

Bloor Research

Contents

Management Summary

Bloor Research was commissioned to produce a white paper, which investigates the whole area of technology consolidation and particularly infrastructure consolidation and reports on the benefits available and the nature of the service that Fujitsu provides. The paper looks at many aspects of consolidation. The following is a summary of its observations and its findings.

- We characterise the problems of IT infrastructure as having the following causes:

    1. **The Skills Problem:** The general shortage of IT skills
    2. **The Quick Fix Problem:** The tendency to "add another box to deliver another application"
    3. **The Absence of Network Architecture:** The lack of architectural design in the growth of IT infrastructure
    4. **The Lack of Central Control:** The over-distribution of computing power in IT infrastructure
    5. **The Lack of Metrics:** For measuring the actual service levels provided by an IT infrastructure.

- The Fujitsu infrastructure services are based on a comprehensive methodology and associated generic architecture, which goes by the name of ***Enterprise Infrastructure Integration***. This provides a framework within which infrastructure consolidation projects can proceed. The methodology targets the organisation of IT staff and operational processes to manage the IT resources. The architecture provides a solution for the consolidation of data management, the consolidation of enterprise servers and the consolidation of the end user layer. It targets the creation of a flexible environment within which the whole IT infrastructure can be optimised. In our view, it is a well thought out and excellent basis for infrastructure consolidation.

- The architecture and methodology are flexible enough to be able to be implemented in stages, allowing organisations to focus on small areas in a controlled manner, providing the potential for a fast return on the investment.

- Fujitsu is the third largest IT services company in the world and in respect of consolidation and IT infrastructure work, it brings the asset of experience that is both wide and deep. The services provided are, however, technology independent with the technology employed in consolidation projects being chosen on a best-of-breed basis.

- Having considered the economic benefits of consolidation we consider the following "rules of thumb" to be reasonable for most organisations:
  - If you consolidate a widely distributed environment then you can expect IT operational and support staff costs to fall by anything up to 80 percent.
  - If you consolidate a fragmented, but largely centralised, environment you can expect IT operational and support staff costs to fall by anything up to 50 percent.
  - In both instances, you can expect the service provided by the network to improve and the IT infrastructure to be more flexible to change.
- The same "rules of thumb" can be applied to data consolidation on its own except that it only covers staff involved in data management. In our view, the consolidation of data management provides the most likely area for cost savings to arise.
- In our view, there are reasons to take a strategic view of infrastructure consolidation as it provides many long term benefits. In particular, it puts an organisation in a position of greater flexibility for future enhancements to the IT infrastructure, puts it in a better position to consider outsourcing and provides it with the ability to monitor service levels accurately on an ongoing basis.

The Great Convergence

The Big Waves

Information technology has seen many dramatic changes in the 50 or so years of its existence and it continues to evolve. There are many driving forces behind this, which make the future evolution of IT difficult to predict. However, there are also two driving forces which are reasonably well known and can help us gain a broad perspective on what is yet to come.

The first of these two forces is generally referred to as "Moore's Law" which derives from a comment made in 1965 by Gordon Moore of Intel when he predicted that the number of transistors per integrated circuit would double approximately every 18 months and that this trend would continue for at least ten years. In effect, he was predicting that CPU power would double every 18 months. This prediction proved to be true and has continued right up to the present time.

It has also had a much wider area of application than just the CPU. In fact, currently, all of the following technology areas are increasing in power at roughly the same rate:

- Switch speeds
- Bus speeds
- Memory speed (and size)
- Disk density and access speed
- Offline magnetic tape storage
- Optical fibre and communications bandwidth
- Wireless bandwidth

In summary, all of the components of computing are accelerating in capability at this remarkable rate and the current belief is that, technically, they will continue to do so for another 15 years, assuming that there is market demand to support this.

The primary impact of Moore's Law was to distribute computer power to a wider and wider audience and create new areas of application for it on a regular basis. In corporate computing, power was distributed to a departmental level with the advent of minicomputers and then to the individual level with the advent of the PC. This wave of "the distribution of computer power" still continues and is now making its presence felt in mobile devices of various kinds.

The second major driving force behind the proliferation of computing can be referred to as Metcalfe's Law, which originated in the late 1980s with Bob Metcalfe, one of the inventors of the Ethernet communications standard for networking. Metcalfe claimed that:

"The 'value' of a computer network is proportional to the square of the number of computers in the network."

This expressed the idea that the value of sharing computing resources, and particularly applications, escalated in an exponential way as the size of networks increased. Although not as measurable as Moore's Law, there was clearly an observable effect here. Put simply, one could say that the value of a telephone network depends upon the number of people with telephones and that this is in turn is proportional to the number of calls that can be made (which, mathematically, is proportional to the square of the number of telephone users). Metcalfe was simply applying mathematics to computer networks.

The truth of Metcalfe's law became apparent with the advent of the Internet and particularly with the proliferation of email and web sites. It can best be thought of as a community or collaboration effect, with the collaboration being driven by computer-to-computer connectivity. Like Moore's Law, Metcalfe's law is still having an impact. As more and more people connect to the Internet and more and more applications are made available through the Internet, the value of the whole network increases and so do the business possibilities that it offers. This escalation in value will also continue for many years as the Internet improves in capability.

The Tactics of the Last War

The PC revolution was about individual empowerment and the networking revolution is about collective empowerment. With the advent of the PC, the IT industry was driven by the numbers. The cost of computer power dropped and those vendors that provided technology to the PC market grew rapidly in size to become industry giants as they took advantage of a rapidly growing market, which only began to move into decline in 2001.

The computing power that the PC eventually began to deliver had a huge impact on the way that businesses used computers. PCs were deployed to provide individual productivity. Control over the use of technology devolved, and in many organisations control of the corporate computing infrastructure degenerated.

The need for central control started to be re-established when the power of networking became apparent. It began with the simple sharing of printing resources and file stores (file and print) for PCs and moved on to co-operative "work group" applications, which were primarily concerned with office automation and the use of email. The Internet suddenly upped the stakes with the proliferation of web sites (internal and external) and possibilities of electronic interaction with suppliers and customers.

It was predicted many years before that when the Telecommunications industry and the IT industry converged it would spark a business revolution. This commenced in 1995 and it has been rolling ever since. Networking came of age with the linking of many small networks together into an incredibly large one. Prior to the Internet it was uncommon for more than a few hundred PCs or computer terminals to be connected together. By the beginning of 2001, the Internet was connecting about 500 million computers together of which at least 350 million were PCs.

It is said the generals will try to fight the next war using the tactics of the last one. When the weapons technology has changed, it is usually, of course, folly to do so. A similar claim could be made about the spontaneous growth of the Internet. It was assembled using the tactics, and indeed much of the technology, from the previous era of computing.

It is clear that the PC was never designed to be an Internet access device. It is also the case that PC LAN servers were not designed for duty as Internet servers. Even the UNIX servers that run most of the large web sites were not designed with that target in mind. We can visit every layer of IT infrastructure and make the same comment. None of the operating systems, the middleware, the databases, the email systems, the business applications and so forth, were designed for a world-wide network of millions of machines.

Nevertheless, the business case for exploiting the network was irresistible and most companies pursued the business opportunities using the technology available and using the tactics that they had previously used to expand their corporate networks. The IT industry itself was not prepared for what began to happen. It too was pursuing the tactics of a previous era.

The Swing of the Pendulum

The typical commercial model for PC oriented IT companies was based on volume sales. The PC market was numbered in the hundreds of millions, so a product that sold for $100 could, in theory, garner revenues in the $billions. Upgrades to products were treated as if they were new products. For many involved, it followed the pattern of an FMCG (Fast Moving Consumer Goods) market. Many IT vendors, including true giants like Intel, Microsoft and Dell had business models based upon these revenue patterns. The business market for information technology and the consumer market had become deeply entwined.

With the advent of networking, this trend began to reverse and with the advent of the Internet it swung back at great speed. However, it was not possible for IT vendors that were caught up in the status quo to suddenly change their products, pricing strategies and behaviour pattern, even if they saw the need. And they did not need to change immediately anyway, because many business users of computing continued to buy and implement IT in the way that they had become accustomed to.

Nevertheless, a great convergence was in progress. The browser became the new user interface, superseding Windows. Most new systems that were created were built with this in mind and many old systems were refaced to

allow usage through the browser. The need for hugely scalable applications made itself felt almost immediately. Applications had only catered for small numbers of users, rarely above the hundreds. Now they needed to cater for, possibly, hundreds of thousands of users. It suddenly became necessary to run applications for 24 hours per day every day. The whole game changed. The logic of the market now favoured large scale computing; big servers, big databases, big applications and high availability.

And it was clear also that the network would continue to expand, not just growing in size, but embracing a range of wireless devices and ultimately providing connections to embedded processors in everything from washing machines to cars. Ultimately an information utility would emerge rather like a national power network, but far more complex and far more varied in its usage.

Moore's Law Revisited

Metcalfe's Law was in full swing, but Moore's Law had not been repealed. It simply began to act in a different manner. Having enabled the distribution of computer power it now began to enable the sharing of computer capability. Increases in CPU speed were still welcome, but now they contributed to the power of large clustered configurations of computers. Faster data access speeds enabled the running of large databases and the emergence of SANs (Storage Area Networks). It was also beginning to enable a newly emerging collection of network-attached storage subsystems. Faster switching speeds made Internet routing less of a problem. The bandwidth increases that Moore's Law delivered enabled the building of high capacity information highways.

Moore's Law is now delivering economies of scale. The annual increase in computer power, for example, cuts the cost of executing transactions, on average, at the rate of about 20 percent per year. In other words, the latest server will usually be capable of a 25 percent greater workload than one bought a year ago for the same price. This is an impressive statistic, but it is not so easy to turn this theoretical figure into genuine cost reductions.

The major reason for this is that the costs of running a network are not directly related to the costs of hardware; they are far more dependent on the cost of maintenance and management. For most organisations, achieving economies of scale in corporate computing means addressing the legacy of the past. It means consolidation of IT infrastructure.

Riding The Consolidation Wave

The Pattern of the Future

As the worlds of telecommunications and computing merge, a convergence of technologies is gradually occurring that will allow different components of technology to co-exist on a pervasive world-wide network. The hope is that, as time passes, the resources that constitute the Internet will provide well defined and predictable levels of service. Databases will be able to exchange data seamlessly across the network. Applications will be able to collaborate in the execution of complex transactions. A dependable security service will protect applications and data. Software upgrades will proliferate automatically. Back-up and recovery will be guaranteed. Indeed everything will happen in a well ordered manner.

This is an idealistic technology vision, and thus a distant possibility. However, there is a definite movement in progress that aims to turn the Internet into an information utility – to provide a service that is as reliable as fixed line telephones. It is interesting to note, for example, that the British Government, through its Office of Science and Technology, is building a "National Grid" for collaborative scientific research. This is to serve as a test-bed for deploying "e-utility computing" - the delivery of computing resources - including bandwidth, applications, storage - as a utility-like service over the Internet.

There is also a market pull that is in operation. Like the world's telephone network, the Internet is not a single entity, but the aggregation of a whole series of communication lines and computing resources, each of which varies in terms of who owns it, the cost of using it and the service that it provides. Within this, there are already private virtual networks that do, within their domain, provide a guaranteed high level service. There are also data centres running applications for multiple companies that share resources under a common management structure and also provide high service levels. Many companies in the Internet sector already provide e-services of many kinds including e-mail, data storage, collaborative capabilities, business applications, security services and so forth. The information utility is gradually coming into existence by a combination of market evolution and government initiative.

Guaranteed service levels across the whole network are necessary if businesses are to be able to use the Internet confidently as an information utility. Products that fail to meet the necessary criteria to work well in such an environment will eventually be discarded in favour of those that do. IT vendors will thus gradually conform to provide the necessary robustness, resilience and compatibility that this environment will come to demand.

The whole point of any utility is to provide a low cost service based on economies of scale. You can generate your own electricity if you wish, but it is not economic. As an information utility develops the same logic will apply. You can run your own systems if you like, but it will be less economic.

The natural question to ask, then, is who will be providing the resources that make up the Grid? The answer to this is the obvious one – it will be those companies that are already in the business of providing large networks and running very large data centres. It will be the telecommunications companies (BT, Cable & Wireless, Energis, etc.), the big ISPs (AOL, Microsoft, etc.), big computer companies (IBM, Fujitsu, Hewlett-Packard, Sun Microsystems, etc.) big outsourcing organisations (EDS, Fujitsu, CSC, etc.) and big organisations with high IT spend (banks, utilities, and possibly, government). They have the skills, the facilities and the inclination to get involved, and they will partner with each other, in "co-opetition", to build the future Internet.

The Pattern of the Past

Computer networks everywhere suffer from the legacy of the past. They have been assembled from different generations of technology, with each incremental change being a tactical solution to cater for the pressing needs of the time. Corporate networks have grown pragmatically, but usually chaotically. At the hardware level, they are a mixture of mainframes, minicomputers, PCs, PC LAN servers, Intranet servers, Internet servers, data servers, application servers, mail servers, domain servers, firewalls and so on. They contain a variety of operating systems, database products, system and network management products, middleware and a plethora of applications.

In most instances, the cost of technology itself was never an issue and neither was the cost of deploying it. The systems and applications that were implemented delivered benefits that outweighed the upfront costs. However, the high costs of maintaining a poorly integrated network were rarely appreciated, and as the networks grew, these costs escalated. And they continue to do so.

The major causes of poorly integrated networks are:

- **The Skills Problem:** There is a general shortage of IT skills throughout the industry. On its own, this generates problems in retaining and recruiting skilled staff. This problem is exacerbated by rapid changes in technology, which regularly de-skill IT professionals. It has been further exacerbated by the major switch from the distribution of computer power, which characterised the PC era, to the sharing of resources, which characterises the Internet era, and demands a wholly different approach to deploying technology.

- **The Quick Fix Problem:** This is the tendency simply to add the next business solution into the network in a "stand-alone" manner for the sake of expediency. It often manifests as "adding another box to deliver another application" without any proper consideration given to the ongoing management and support costs of doing so.

- **The Absence of Network Architecture:** This can best be thought of as having no technical or management strategy for optimising the corporate network to extract the best value for money from it as a whole. Implementing a quick fix may be the most pragmatic action to take at times, but failing to integrate it later is costly.

- **The Lack of Central Control:** The distribution of computing power in many companies led to the loss of strong central authority over the network. This, in turn, made it impossible or difficult to impose common policies and standards.

- **The Lack of Metrics:** In the era of mainframe computing, IT departments kept detailed metrics of the cost and usage of computer resources. This made it possible to know the actual costs of providing any application or service accurately. The gathering of such information for whole networks is rarely known accurately and often the true costs of any application are not known.

The move to Internet computing has created many challenges to corporate computing. The major issue is that many applications need to be available for 24 hours in every day, leaving no time window for backup or maintenance operations. At the same time, service levels must be kept as high as possible, with response time and throughput being key measures. To exacerbate the situation, large parts of an application can be beyond the control of the corporate network, subject to the vaguaries of Internet performance or possibly run in a remote environment by an ISP or other outsourcer.

The more hardware and software components that exist within the network, the greater the risk that something is going to fail. In very complex environments, simply understanding the relationships between components and the effects of even minor variations in behaviour becomes a problem. For this reason, if for no other, centralisation and consolidation of the corporate network is desirable.

Where there are problems, there are usually opportunities and this situation is no exception. Centralisation, standardisation and consolidation bring cost savings. They also create greater flexibility allowing the corporation to take advantage of new hardware, software and communications opportunities that emerge. The corporate computer network needs to be treated as a single resource space; to be built as such, managed as such and expanded in an orderly fashion. Where this is done, the likelihood of being able to take advantage of the power of networking and the power of the Internet is increased and the cost of computing declines.

Fujitsu Consolidation Services

The Forms of Consolidation

While the picture varies from company to company, the general state of corporate computing is that networks have grown chaotically and there are opportunities to save costs through various forms of consolidation. There is also a strategic necessity to do this to take advantage of the global computing trends. Fujitsu identified this need some time ago and developed a set of services to assist organisations in making sense of their networks, and they have been tested in the field at many customer sites.

These services range from complete outsourcing through to consultancy projects, which deal with all forms of consolidation. Before going on to describe the Fujitsu services, it is worth considering the different opportunities for consolidation, one by one.

Vertical Application Consolidation

In some areas of application deployment, the "add another box" tactic has led to a proliferation of servers and, in extremis, to the creation of server farms. This has particularly been the case with email systems and web sites/intranet sites. It has also happened with the proliferation of PC LAN servers at a departmental level. The opportunity here is to consolidate many servers, which have similar workload characteristics, into a single expansible server configuration.

When similar or identical applications are distributed over a multitude of servers the utilisation of each server is generally very low, often much lower than 50 percent and in some cases as low as 20 percent. On top of this, the level of service provided usually leaves much to be desired in terms of resilience and recovery from failure. In general, the cost of server management is proportional to the number of servers. The ability to standardise increases as the number of servers decreases. The opportunity here is both to reduce the cost of management and to improve levels of service.

General Server Consolidation

In other areas, different applications may be deployed in the same "add another box" style, which can also be consolidated. The consolidation of such servers requires a greater amount of investigation and capacity planning work, and there may be limitations on what can be achieved. It may be possible to consolidate on a small number of strategic operating systems (Solaris and Windows NT, for example) running on large servers, but, it may not be possible or cost effective to switch some applications from one operating environment to another, because of technical factors or even because of the licensing implications. However, in most large corporate environments there are opportunities for this type of consolidation. As before the opportunity is both to reduce the cost of management and to improve service levels.

Data Management Consolidation

The development of storage area networks (SANs) gave birth to opportunities for data management consolidation. The principle was to cater for back-up, recovery and all other aspects of data management under a single umbrella. Because of the fragmented nature of most corporate networks, many organisations have opportunities to consolidate in this area, including some that have already employed such technology but have not deployed it across the whole network.

Operational Consolidation

The fragmented nature of corporate networks has led to uneven service levels in terms of performance, application and data availability, the cost of upgrades and maintenance, security and support. Operational consolidation is the consolidation of service methodologies across the network to provide a robust and dependable service to all network users. The outcome of such a consolidation may be a reduction in costs or an improved service, but it will also deliver greater flexibility and upgradability for the future.

Enterprise Infrastructure Integration, A Methodology for Consolidation

To pursue the goal of consolidation, there is a need for both a methodology and an architecture. It may appear attractive, for example, to fix on an obvious target such as consolidating a distributed email system, and proceed with it in haste. It is not good sense to do this without considering the big picture. Even if there is an apparent cost justification, it may be a sub-optimisation in terms of the network as a whole and there will be implications for the management of the network. For that reason, consolidation needs to be driven by a methodology and architectural principles, which optimise the whole.

Fujitsu employs a service methodology, called ***Enterprise Infrastructure Integration***, for optimizing a network infrastructure and its management. The methodology and the generic architecture that it employs target the following specific benefits.

- To reduce the total cost of ownership
- To deliver high levels of reliability and availability (up to 99.999% for specific cases)
- To deliver efficient data management facilities, including total data integrity and fast backup and recovery
- To deliver high levels of scalability and flexibility across all applications
- To provide ease of extensibility in small increments and upgrade without the need for further re-engineering
- To deliver high levels of security
- To provide integrated and automated systems management, both local and remote
- To provide efficient support and maintenance services
- To coexist and integrate with existing systems, in particular, with existing mainframe systems

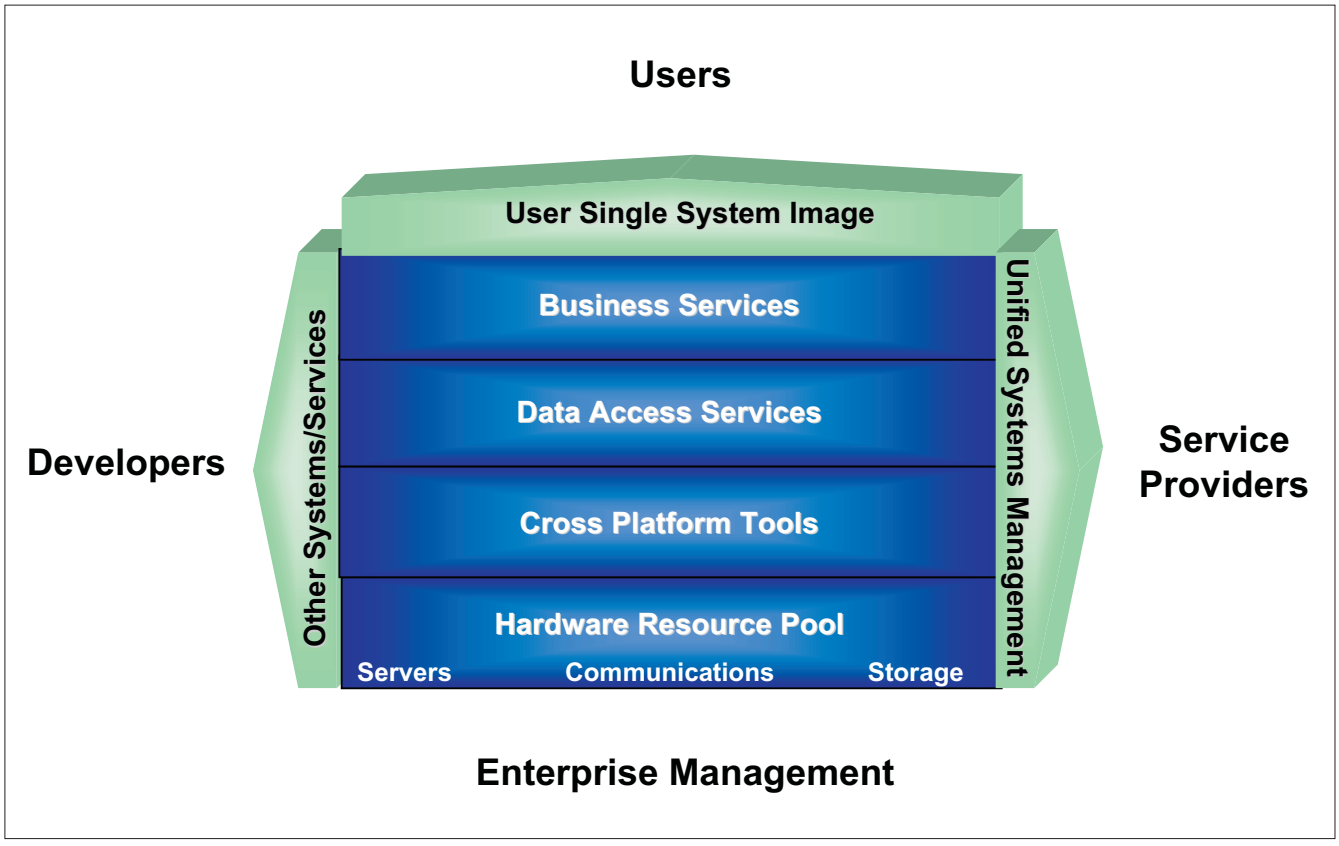

Fig. 1: Enterprise Infrastructure Integration Framework

The ***Enterprise Infrastructure Integration*** architectural framework views the building of IT infrastructure as comprising three ingredients:

1. Products: The selected technologies, such as hardware and software components, from which infrastructure solutions are built. (Fujitsu focuses on a core set of products and services from a carefully selected set of partners that have been identified as providing best-of-breed products.)

2. People: The individuals who bring the skills and knowledge required to build solutions from the selected components, deliver the solutions to the customer, and manage the commercial aspects involved.

3. Processes: Formal proven processes, which cover the commercial and engineering aspects of the solution.

These resources are orchestrated to build a solution that meets the required capabilities. A simple diagram of the overall framework is provided as Figure 1.

The importance of getting the management structure and organisation right should not be underestimated. Research by other analysts and vendors has established that products are responsible for only about 20% of the problems

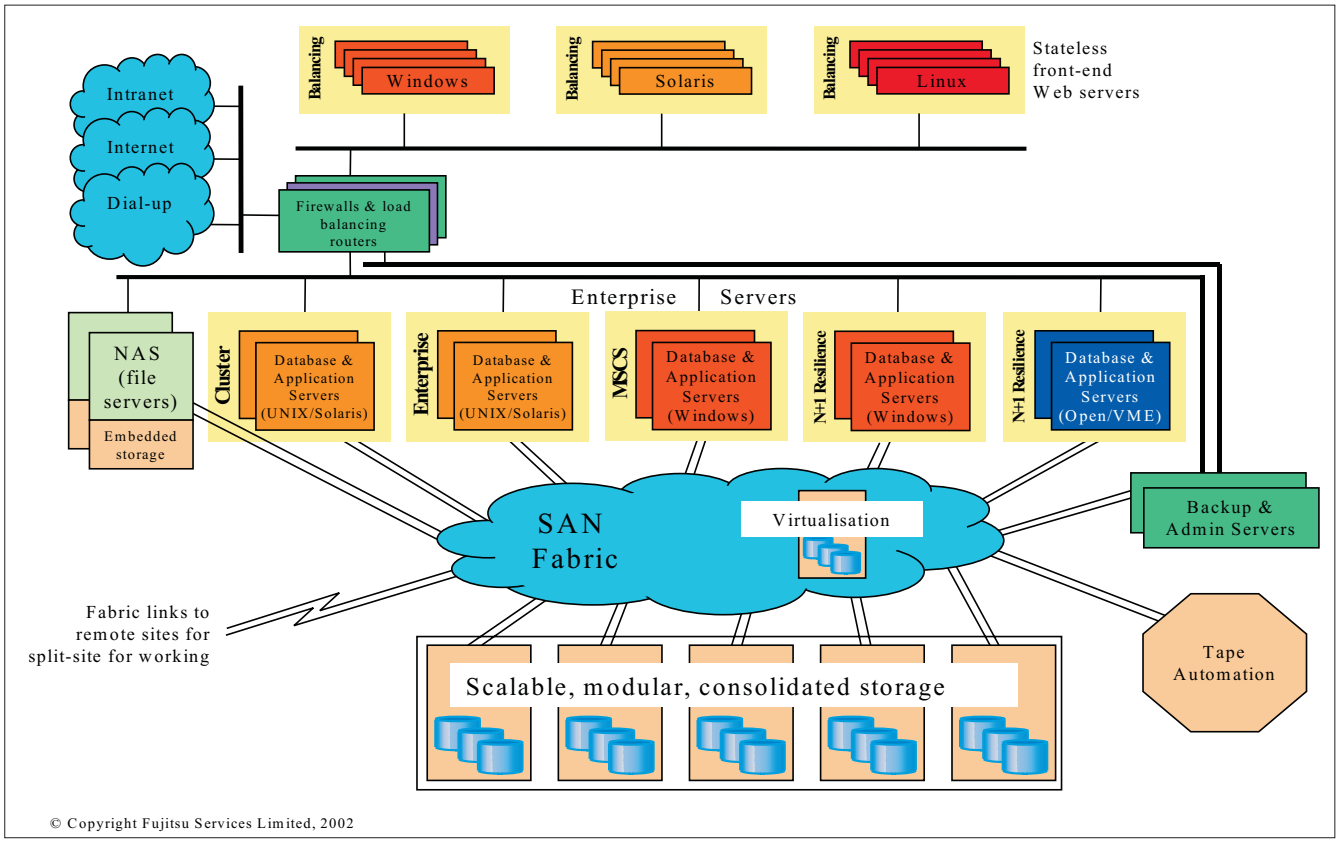

Fig. 2: The Generic Architecture

that arise in an enterprise infrastructure. The remaining 80% is split equally between the people and process elements. The ***Enterprise Infrastructure Integration*** methodology employs the Fujitsu management framework specifically to address this reality. The whole problem – the 100% is addressed by combining skills, proven processes and appropriate best-of-breed technology.

The Generic Architecture

The generic architecture that enterprise infrastructure integration employs is illustrated in Figure 2. It can be thought of as consisting of three layers:

- a data management layer, which is based on the network storage technologies; Storage Area Networks (SANs) and Network Attached Storage (NAS)
- an enterprise server layer, which consists of consolidated servers (usually clusters) that drive "enterprise systems"
- a front-end layer based on web servers

Put simply, the way that this architecture works is that the SAN/NAS technology handles all aspects of data management for all data. This provides data to applications that run in the enterprise server layer and also

acts as a back-up resource for all departmental servers in the front-end layer, and possibly also for remote servers in other geographical locations.

This may be linked for disaster recovery purposes to a mirrored site. The back-up and administration of this is handled from a single point and a tape library, linked to the SAN, can be included.

The enterprise server layer is consolidated, as far as possible, into a minimum number of highly configured servers running the major applications of the enterprise. This embodies the system management capability, which provides operational control of the whole configuration. It may also include some network attached storage (NAS) for simple file serving, to act as a consolidation point for file serving on behalf of distributed PCs. Taken together, this server layer and the data management layer make up the organisation's data centre, connected together on a high speed LAN, confined to the data centre and protected by a firewall.

The front-end layer consists of dispersed servers deployed at different locations throughout the organisation. These are "stateless" so that all aspects of recovery can be handled within the data centre.

In practice, Fujitsu can bring a whole series of specific technologies and the associated expertise to bear in order to generate a solution that is appropriate to the requirements and budget of the customer. At the moment this includes:

Operating Systems: Windows NT and 2000, Solaris, OpenVME and IBM MVS.

Enterprise Servers: Aside from mainframe support, Windows servers running Microsoft Exchange and SQL Server applications and Solaris Servers running Oracle applications are catered for.

Front-end Servers: These include Microsoft terminal server and web front-end servers running Windows or Linux and general infrastructure servers such as domain controllers.

System Management Technology: These embrace the major frameworks such as CA Unicenter and Tivoli.

Back-up Technology: This includes products from Legato, VERITAS, CA, providing for all types of back-up option through the SAN or otherwise.

Security Technology: This covers everything from authentication through to virus checking and employs a whole range of security products.

Fibre I/O Networks: Switched fabric fibre optic networks are employed. These cover both IP/Ethernet for general data communications, gigabit Ethernet and Fibre Channel for SAN connectivity and resilience, disaster recovery and local clustering.

Storage Technology: Both SAN and NAS technology can be deployed providing on-line and near line hierarchical storage management systems. The storage arrays used support the attachment of multiple heterogeneous hosts to single shared arrays. Off-line and near line storage is based on shared robotic tape library systems.

Resilience and Failover: Server replication and shared storage failover techniques are employed to provide availability and resilience. This includes, where appropriate, the use of Fujitsu's N+1 High Availability methodology (patents pending in Australia, Europe, Japan and USA) that caters for Windows and UNIX system failures and processor failure in OpenVME mainframe systems.

Network Connectivity: The data centre systems are connected using high speed Ethernet LANs linked through high speed routers that enable load balancing. Network gateways for Internet access, dial-up, fax, etc. are also catered for.

In summary, the methodology and framework are intended to provide availability, performance, security, manageability and flexibility via the use of best of breed products.

Consolidation Projects

It is rare that an organisation embarks on a consolidation project off its own bat. Perhaps it ought not to be the case, but it generally is. What usually prompts an organisation to consolidate its IT infrastructure is the appearance of unexpected costs or the occurrence of an expensive failure, or possibly the impending implementation of an important new application. Whatever the circumstance, it will have focused attention on the inadequacy of the existing IT infrastructure.

In reality, most existing IT infrastructures and their associated management structures are inadequate for the requirements of Internet computing. Even if the infrastructure was very well designed, it is highly unlikely that it was designed for extremely high availability (24 x 7) with a very high level of scalability. Similarly, the operational management is unlikely to be appropriate.

The ***Enterprise Infrastructure Integration*** generic architecture described above is targeted at providing exactly that kind of service and to do so for a very large operation. Similarly, most networks are too distributed, because they were built for the era of distributed computing, which has now passed. The architecture also seeks to rectify this where centralisation is feasible and demonstrably more economic.

However, the consolidation projects in which Fujitsu engages do not necessarily involve a full strategic makeover of the corporate IT infrastructure. The target may be less comprehensive, perhaps the implementation of a SAN or the consolidation of several applications and their server hardware on which they run.

The ***Enterprise Infrastructure Integration*** methodology and architecture is applicable to such situations. Its scope embraces operational management, the design of all the infrastructure aspects of data centre solutions, the provision of integrated management features to the network, the selection and provision of tools and services and, ultimately, the provision of support for any of the solutions that it is used to build. It also includes proof of concept services to test possible solutions and integration options that may never have been tried before.

After the provision of terms of reference, consolidation projects commence with a data gathering and study of the IT infrastructure that is to be enhanced. The production of a proposed solution normally takes no more than a few weeks, unless specific proof of concept work is required, and solutions can be implemented in a few months, if the circumstances allow.

The Fujitsu Pedigree

The consolidation projects carried out by Fujitsu involve an application of skills and know-how that is uncommon even among large consultancies. Although some may be unaware of it, Fujitsu can claim to be the third largest IT services company in the world, with revenues of over $46 billion. Roughly one third of this derives from the sale of software and services.

Outside the Far East, Fujitsu Services and Fujitsu Consulting are responsible for the bulk of the services business, generating revenues of $4.6 billion. This derives from a variety of activities including application integration, consultancy, training and a variety of IT Infrastructure services embracing outsourcing, managed services, project services, data centre, network and end user services.

It is arguable that the computer vendors who focus on the services market are best placed to offer the kind of consolidation service described above. Not only do they run many data centres within their outsourcing businesses, but they have a deeper understanding of technology developments and technical skills in considerable depth. Fujitsu's spend on R&D alone is $3 billion and covers all areas of the computing and networking world.

Fujitsu is unusual amongst such vendors in that the solutions it provides do not have a bias to its own technology. Indeed, although it has its own range of Intel based servers, in its consolidation services it deploys more Intel servers from other vendors than it does of its own. Its solutions depend upon selected best of breed products and proven combinations of technology irrespective of origin

The Bottom Line

The Trend to Consolidation

The trend to consolidation could have begun several years ago when it became clear that networked computing and the Internet would be the dominating technology driver for the foreseeable future. It did not happen at that time for several reasons, principal of which was that many IT vendors were not geared up to provide the appropriate technology or to change their business models to align with the emerging trend. Such things take time and it also takes time for IT users to adjust to the change in direction.

Nevertheless, the trend to consolidation is clearly underway and will continue as technology converges to satisfy the demands of Internet computing. Even in the era of distributed computing that has passed, surveys regularly reported that the cost of running centralised applications was always significantly less than the cost of running the same applications in a distributed environment. There always have been economies of scale associated with the use of more powerful and manageable server configurations.

The motivation for distributed computing was primarily to provide applications that were impractical or simply unavailable in centralised computing environments. They were mainly office applications, graphical applications and, later on, workgroup applications. It was a different style of computing and it never replaced the role of the mainframe in running mission critical applications. Telecommunications costs also played a hand in how computer power was distributed. These have now fallen, altering the cost equations.

The trend to consolidation is being supported with the delivery of more sophisticated and scalable products at every level of hardware, software and communications. This will continue while the trend persists and it will lead to a gradual reduction in the numbers of server configurations attached to the world wide web. Remarkable as it may seem, there are over 100,000,000 such servers, a number that could be reduced considerably by consolidation.

Economic Motivation

As we have already stated at various points in this paper, there is a host of operational motivations for consolidation, which can be summarised by the words ***flexibility***, ***efficiency*** and ***service***. It is worth examining the economic aspects of this in a little depth. Let us first consider the costs of storage.

The Storage Situation

In a report produced in 2001, the Gartner Group reported that the costs of data management were somewhere between 5 and 7 times greater than the costs of procuring storage devices. This is a rather alarming figure. The data management problem has exploded for a variety of reasons, most of which relate to the need to record new information in a digitised form. According to a report from U/C Berkeley Research, the world produces 1.5 million terabytes of new information per year – which amounts to about 250 megabytes per living person. The report claims that 93 percent of this data is in digitised form and an alarming 55 percent is stored on PC hard drives –

which means that it is probably inadequately backed up and inefficiently managed.

Some of this data, consisting of video, voice and graphic images, has little interest for most businesses, but much of it does; emails with attachments, web site content, web logs, electronic documents, presentations and so forth. And, inevitably, it will eventually be necessary for businesses to store and manage video and voice data. Unless it is brought under control, the problem of data management will escalate of its own accord.

There are economies to be achieved in the organisation of the data itself. Aside from the fact that a SAN can ultilise around 85 percent of the storage resource available to it without any degradation of service and, whereas distributed servers rarely achieve 50 percent utilisation, distribution also leads to significant levels of duplication (with 15-20 percent of data being duplicated many times) and the retention of a great deal of data that could easily be archived (often as high as 45 percent). In a managed data environment such inefficiencies can be removed.

A report from the Sosinksky Group this year, which classified IT infrastructure according to architecture, estimated that:

- In distributed networks, 400 gigabytes of data demanded one individual to manage it
- In a centralised network, 800 gigabytes could be managed by one person
- In a centralised and consolidated network, 3000 gigabytes could be managed by one person.

Thus, for data storage, a clear pattern emerges of lower staffing costs associated with greater consolidation. In a consolidated network with a central SAN and a single data management regime, maximum consolidation is achieved.

The Server Situation

The apparent paradox, that mainframe computing has always provided the least cost of ownership, is explained by the fact that, in recent times, the most compelling applications, which delivered office automation and Internet applications, did not run on mainframes. The consequence of this was the proliferation of Novell NetWare servers, Windows NT servers and UNIX servers. The consolidation strategy is, as with a SAN, to bring the problem under a single managed environment to achieve economies of scale.

As a rough rule of thumb, which is sometimes referred to as Fink's Fifth Law, complexity in a system is *proportional to the square of the number of components*. More accurately, in a server environment, we could claim that complexity is proportional to the square of the number of managed environments. If we have, for example, 8 servers that are independently managed and add a ninth then we will, at the very least, incur an incremental management cost associated with the extra server. If the server interacts with

the other 8 then the complexity this creates may increase that management cost further.

For small servers, the ***annual*** cost of managing the server is usually much greater than the cost of the server itself. In some sites we are aware of, the multiple is a factor of 5 or more at the low end. It is interesting, but not surprising, that this figure is similar to the Gartner figure for the costs of data management against the costs of storage devices. In the server space, the management costs derive from providing back-up, recovery, security, software and hardware upgrades and all aspects of support. There are clearly major economies in providing a consolidated management regime in the data centre, which ease all these tasks. We can expect that the same economies will emerge as those that emerge with consolidated data management. In effect, you may be able to halve the management costs by moving from a distributed to a centralised environment and then cut the costs by 70 percent again in moving from a centralised to a consolidated regime.

Thin Client Devices

In the user area, the most common strategy for consolidation is to deliver applications to users through a browser interface. This is normally achieved in a two-pronged strategy. First, all applications that can be are converted to browser operation. Secondly, Citrix software is deployed on Windows NT servers with Microsoft Terminal Server so that office applications no longer run locally on the PC, but on the server. This consolidates the PC management problem, although it does not resolve all management problems, as laptops will still do local processing under such a regime.

This enables the organisation to be able to deploy thin clients if required, but if not will still delay the upgrade cycle of PCs as the client PC no longer needs regular upgrade. The popularity of Citrix attests to the economies that this strategy can deliver. Users of such a solution report high levels of saving in the area of support. A reduction of such costs by 50 percent is not unusual.

Overall Considerations

We have considered the economies in three layers; data, applications and users, but need also to consider the IT infrastructure as a whole. Any consolidation that improves the service provided to system users is likely to diminish the cost of supporting those users and also reduce the amount of work interruption experienced by users because of hardware or software failures in some area of the operation. Thus there can be a genuine productivity gain across the organisation that derives from consolidation.

Similarly, there may be some other economies that emerge from a consolidation project as a whole. Consolidated servers occupy less space and place fewer environmental demands. Space may be saved within the data centre or across the whole organisation. If the number of operational staff is reduced then this will deliver economies in staffing costs from recruitment through to training.

There is no simple, automatic formula for determining how much cost saving can be gained from consolidation because each context is different and poses different problems. Fujitsu, when tasked with designing and planning a

consolidation, can provide some indication of the likely savings, but much depends on context and the goals of the organisation involved.

However, we suggest the following "rule of thumb":

- If you consolidate a widely distributed environment then you can expect operational IT staff costs to fall by anything up to 80 percent.
- If you consolidate a fragmented but largely centralised environment you can expect operational IT staff costs to fall by anything up to 50 percent
- In both instances, you can expect the service provided by the network to improve and the IT infrastructure to be more flexible to change.

The Second Value Proposition

In an IT world that is now driven by the trend to consolidation, service levels are going to become increasingly important. Because service levels are rarely accurately measured in most organisations, it is difficult for an organisation to calculate the value of increasing service levels. There are admittedly, some examples of financial loss from poor IT service – for example, the massive collapse in the share price of E*TRADE when its service went out of action in spring of 2000. This, however, is an extreme example and, in most instances, poor service levels gradually erode customer confidence and have a cumulative negative impact on business.

From a general perspective, service levels can be broken down under three headings:

Performance: This is the actual performance of applications and particularly the performance required to satisfy customers and not inhibit the productivity of staff. Poor performance leads to inefficiency within the organisations and dissatisfaction amongst customers. It is measured simply by application response times. Related to this is application scalability, the aim of which is to not affect performance when customer numbers increase.

Availability: This covers when applications are available as scheduled (since not all applications need to be available all the time) and how frequently they are unavailable in such periods. It should be noted that providing very high availability, for example the 99.999 percent availability that is often bandied about, can be very expensive and is rarely justified. 99.999 percent availability implies a loss of only 5 minutes in a year. 99.99 percent implies about 50 minutes in a year. For most business the difference will be negligible but the cost of 99.999 percent will be significantly greater. Nevertheless, when applications are unavailable staff productivity is reduced and customer business can be lost. Availability is a product of many factors from resilience through to security.

Support: The third aspect of service levels is support. This is measured by gathering statistics on the occurrence of service problems from software anomalies and errors through to the failure of equipment and the speed of resolution. Support may also include elements of education as the support

team can sometimes simply advise users how to achieve things that they are not sure how to do. The cost of support diminishes significantly when fewer failures occur and hence is strongly dependent on the robustness of applications and the environments in which they run.

Unless an organisation knows the service levels that its applications and the IT infrastructure as a whole delivers, it has no means of placing a value on improving these service levels. It is thus a sensible goal for a consolidation project to implement mechanisms for the measurement of service levels and all the factors that contribute to them. In the projects that employ the Fujitsu ***Enterprise Infrastructure Integration*** methodology, such measurement is provided by the implementation of a comprehensive system management framework.

There is thus a second value proposition on offer to organisations that implement a consolidation. It will provide them with an improved ability to understand the actual costs of the IT infrastructure and its impact on the business.

The Outsourcing Proposition

The knowledge and understanding required to build sophisticated, multi-purpose IT networks has increased quite considerably in recent years as a combination of the appearance of new technologies and the advent of the Internet. Organisations may well conclude that it is more cost effective to completely outsource the operation and management of their IT infrastructure, especially given the dramatic reduction in the costs of communications bandwidth.

Historically, the business of outsourcing has involved the transfer of an IT department (or at least the operational elements of it) to the outsourcing company, along with the associated problems, under an agreed contract that covered costs and service. For the outsourcer, the business was to provide operational, data centre and organisational expertise that the customer never had for a suitable reward. Contracts were usually long term because of the major expenses involved at the beginning.

The business of outsourcing is changing. In particular, it is no longer an "all or nothing affair". Companies may choose to outsource a large email system or outsource the running of web sites and intranets, and some do. The dynamics of consolidation plays a part in this. Large ISPs, for example, can consolidate a large resource for web serving and provide a service with a reasonably high service level more economically than the customer, through simple economies of scale.

Fujitsu, and other companies in the outsourcing market, now provide a range of services including the outsourcing of specific systems, the direct provision of support all the way to complete outsourcing. Full outsourcing, if the circumstances permit, is now likely to begin with a consolidation project to rationalise the IT infrastructure prior to taking over full management of it.

Conclusion

In this paper, we have examined the industry trend towards infrastructure consolidation in some depth while reviewing the extensive service that Fujitsu provides to assist companies towards such consolidation. Our primary conclusion is that there is now a visible trend towards consolidation, which organisations need to take note of for their own financial benefit and efficient operation.

We note that consolidation projects are not simple undertakings and that both a methodology and architectural considerations need to be applied to achieve a satisfactory outcome. In particular, it is clear that:

- A successful consolidation project needs to address the organisational skills, the processes and the technology that is used, because an efficient IT infrastructure depends critically upon all these factors.

- The application of a generic architecture, such as that embodied in Fujitsu's ***Enterprise Infrastructure Integration*** methodology, will assist in arriving rapidly at a solution that is effective and exploits the latest available technology.

- Few companies have the requisite expertise, experience and resources to carry out such projects for complex IT environments. Fujitsu, however, is one.

We note that most organisations do not usually initiate such projects proactively, but tend to become involved in them because the infrastructure problems suddenly surface due to other factors. However, we believe that it will be worth while for most organisations to consider the opportunities that consolidation presents – and perhaps consider having some form of audit.

We believe that the benefits for most organisations that will emerge from such a problem will be a reduction in the number of operational and support staff or an increase in the service levels provided by the IT infrastructure, or even both.

We expect the trend to consolidation to be a long term trend. This does not mean, in our view, that all IT infrastructures will become fully centralised - as much depends upon context and business requirements. However, we believe that the management of such networks will, of necessity, become centralised and that a significant level of consolidation will be involved in that process.

Ultimately, we expect that the Internet itself will gradually become a robust network that offers, in part at least, dependable service levels. Consequently, we believe that organisations will increasingly exploit the services it provides. Within this movement, we expect to see a continuance of the trend to the outsourcing of IT infrastructure together with the external hosting of business applications.

Copyright & Disclaimer

Suite 6, Challenge House, Sherwood Drive,
Bletchley, Milton Keynes, MK3 6DP, United Kingdom

Tel: +44 (0) 1908 625100 - Fax: +44 (0) 1908 625124
Web: www.bloor-research.com - email: info@bloor-research.com