

Scorpion

Simplifying the Corporate IT Infrastructure
- An IT Health Check for Business Executives

David G. Heap
Principal IT Consultant
IBM Enterprise Server Group
dheap@us.ibm.com



Executive Summary

Above the roar of the e-business marketing hype, can you hear that quiet but persistent voice whispering - "but how much did you say this is going to cost? ... And what server technology are we using for the e-business foundations?"

The very same questions were asked at the beginning of the client-server era, but were often ignored, swamped by a chorus of techno babble and a lack of comparative platform cost data.

This comparative cost data is available today in most organizations. If properly analyzed, it provides a crucial input to the next generation of IT decisions on e-business strategy. This paper shows how to do the analysis

Why is this topic important? Recent studies of very large IT infrastructures have shown some alarming results. Some organizations are 'grid-locked' and, although spending two to three times the industry average on IT, seem to deliver little or no extra value for this spend.

Most of this extra spend is being ploughed into an increasingly people-intensive IT infrastructure. There is no immediate solution to this, but a coherent, integrated IT infrastructure migration strategy is vital to proper positioning for e-business.

Some of these organizations are now spending over 50% of their IT budget on PC infrastructure, yet many of the fat clients deployed today provide little more than basic office tools and a graphical interface.

In all these cases the IT infrastructure is so expensive that the only realistic IT solution is a radical and rapid move to thin client browser based architectures, to contain the IT cost explosion. This, in itself, is a very powerful reason to move rapidly to the e-business paradigm.

The promise of e-business today is real and profound. But, as in the case of client-server, the technical IT foundations must be properly prepared.

This paper contains many real examples, key metrics and rules-of-thumb. You can use them to assess your IT organization's financial and technical readiness for implementing an effective e-business IT server infrastructure. The sequence of the approach is simple:

- 1. Assess the overall current IT spend per user, and identify the main technology expenditure areas, such as servers, network and desktop.*
- 2. Run a short efficiency and service quality health check for each of the major servers.*
- 3. Identify the real incremental costs of placing future applications on a particular server type.*
- 4. Identify a simple, but coherent IT 'vision' or blueprint, and create a realistic migration plan to get to a future IT infrastructure which is simpler, more agile and cost-effective.*

IBM has service offerings available to assist you in this health check process.



Part One
How much IT spend per user per year?

How much do you spend each year on IT infrastructure for every user? Although this is essentially a very simple question, it can be made to appear extremely complicated. A simple way to get an initial answer to this question lies in identifying the total amount of IT infrastructure spend (say 100 M\$ per year) and dividing by the number of users (say 20,000) to calculate a total cost of IT per user per year (say 5K\$ per user per year).

There is often much discussion as to what costs to include in the IT expenditure total, but most people naturally include IT service delivery people, central and distributed servers, the data network, PCs and workstations. They typically exclude voice networks and application development people costs.


Establishing the total number of users requires a little more care. For the last few years in many organizations, the number of users is more or less the same as the number of employees. This is especially true of banking, financial services and public sector organizations.

But in retail distribution and railroads, for example, there are many employees, such as truck drivers or shelf stackers, who do not need access to a screen. In many organizations there are part-time employees who effectively share IT screens. In this situation, count 'full-time equivalent' users. In yet other organizations, such as insurance or travel services, the primary users are external agents. In almost all cases it is possible to establish a reasonable initial value for the total number of 'full-time equivalent' users of the IT infrastructure.

There are many examples of organizations spending 5, 10, 15, 20 K\$ or even more per user on 'standard' IT infrastructure providing basic operational IT applications, email, etc. These organizations are not providing very specialized facilities such as air traffic control, point of sale or cash machines. So, why the big variation? Consider these four real examples:

Example "A" delivers an IT service at a total cost of 4 K\$ per user per year.

A: Thin Client, Fat Network, Two Hub Centers




This enterprise has a relatively thin client desktop environment, a relatively broad bandwidth network, and central mainframe and UNIX servers concentrated in two major data centres. There is still opportunity to improve the efficiency of various parts of this infrastructure, but the overall IT investment and total cost per user looks very reasonable.

A rough, but reasonable guide to an effective split between the technologies seems to be; 15-20% for the desktop and distributed Intel infrastructure; 20-30% for the data network including all external line and call charges; and 50-60% for the central hub and infrastructure servers, including all central mainframes, UNIX and Intel servers.

Example "B" shows a very different picture. In this case, 53% of IT expenditure on hub servers looks, and is, reasonable. But the proportion of IT spend on the desktop and network is now the complete reverse of

B: Fat Client, Thin Network, 1 Hub Center



example "A". This organization made a major investment in an extensive multi-tier client server system, including a very 'fat' desktop. This system was not only significantly more expensive than initially estimated, but, even worse, delivers very poor (98.5%) application availability and it still takes many months to deploy new versions of desktop client function.


A series of papers on the role of advanced high-end IBM servers

In this case, the desktop and highly distributed Windows NT servers are consuming 37% of the total IT budget, whereas the bandwidth-constrained data network is only 11% of the total IT budget. Here it is appropriate to invest strategically in more network bandwidth and put significant controls on the Intel infrastructure expenditure.

It is also worth noting that the cost per user per year in example "B" is about 40% more expensive than example "A", mainly as a result of the fat PC-client approach rather than the thin browser-based client.

Example "C" is a startling, real case, and, more surprisingly, not that unusual. In this case, the desktop

C: Obese Client, Thin Network, 4 Centers




and Intel environment is out of control, with a total IT spend in this area of 58%. The network is bandwidth constrained, and the four separate UNIX centers very people inefficient. The net effect is an inefficient allocation of individual IT resources and, not surprisingly, poor user satisfaction.

It is also worth observing that, in this case, all of the key areas, desktop, network and servers were requesting more investment funding, even though the total cost per user is over 90% more expensive than example "A". It was only through this top-level, business executive IT perspective of the gross expenditure by technology tower that the conflicting IT investment requests could be resolved and appropriate major investments planned; to consolidate the four separate centers, increase network bandwidth and put extensive controls on the desktop/laptop PC explosion.

Finally, **example "D"** gives yet another perspective, this time on the tradeoffs between desktop and servers. In this case, the organization has a substantial investment in a global network with substantial bandwidth.

However a strategy of massive investment in distributed

D: Obese Client, Wide Network, Anorexic Hub



desktop computing has stifled investment in central application and database servers leading to 'anorexia' in the central servers. This is a significant contributing factor to a major loss of market share against major marketplace competitors because of the lack of exploitation of centralized customer databases and collaborative tools for product design and marketing. All of this for an IT spend over 90% higher than example "A". This is a good example of enterprise IT gridlock.

IT gridlock - a real possibility

These four examples can be very surprising. Most organizations consider they are doing a good job of managing conflicting technology demands, but these charts show that there are very different and very expensive alternative approaches to IT infrastructure. Perhaps the most insidious effect of the client server distributed model is the progressive 'arthritis' it induces into the infrastructure. The inability to change the business and IT infrastructure quickly to meet new marketplace requirements, such as e-business.

The four sample organizations described above have tens of thousands of users, and hundreds/thousands of servers. If the user PCs, Intel and UNIX server environments are not rigorously controlled, IT management costs can increase substantially. But, it is IT management complexity that is ultimately the killer. As the number of servers or instances of data get too large, quality of service drops, deployment time increases and integration testing becomes exponentially time consuming. This is IT gridlock.

Check your IT budget profile

Step one in avoiding gridlock is to create these top level 'pie charts' for **your** enterprise. The intelligent interpretation of total IT spend, number of users, and the technology 'pie-slices' is a valuable initial diagnostic indicator of your overall IT health.

Total IT Spend: M\$ pa Total Users

Total IT Spend/User/Year K\$ / user / year

Technology Tower	You	Suggested
Desktop		15-10%
Distributed Servers		10-5%
Network		20-25%
Central Servers		50-60%
Total	100%	100%

In practice this data can be assembled in a few days, but must be carefully scrubbed and validated to obtain the real cost picture and to develop meaningful insights into the current state.

Another driver towards e-business

These numbers highlight a major reason why many large enterprises are embracing the rush to e-business - it can save money! More specifically, moving to browser based access to systems can radically change the total IT cost profile. With senior business unit executives moving to reduce in house administrative costs, by pushing administrative work on to internet users, IT executives can simultaneously plan their move away from expensive distributed client-server systems towards lower cost, more automated server systems accessed via browsers. The key watchwords are now simplification and automation, as severe skill shortages ultimately preclude the exploitation of advanced, but people intensive technologies.


**Part Two
Assessing today's IT Infrastructure**

These 'top level' pie charts are useful, but need to be supplemented to provide further diagnosis. Each major server 'architecture', typically mainframe, UNIX, and Intel servers, can now be evaluated to check whether they are efficiently used, deliver a high quality service and are reasonably cost-effective. The objective of the IT server health check is to produce a single chart that profiles current IT service provision for each server architecture. Broad server health check metrics include:

- People efficiency; broadly measured by 'S/390 mips per person' or 'UNIX servers per support person'
- Systems efficiency; broadly measured by 'average processor utilization' or 'disk utilization'
- Quality of Service; broadly measured by quarterly scheduled and unscheduled outage time and impact
- Service delivery cost; broadly measured by total cost and the incremental cost of additional useable capacity

In the case of UNIX and Intel servers, it is very important to obtain an inventory of the servers, including their function, location, capacity and number of users. The following sections take realistic case study data to illustrate how you can develop a useful IT scorecard for each of your major server architectures.

E: People Efficiency - S/390 Operations



A: Service Delivery - People Efficiency

S/390 Servers - People Efficiency

A series of papers on the role of advanced high-end IBM servers

The metric of 'service delivery people per S/390 mip' is a quick, broad measure of S/390 service delivery efficiency. This chart shows an enterprise, indicated by a blue star, positioned relative to other major enterprises in terms of its S/390 service delivery efficiency. The horizontal axis shows the S/390 capacity installed. The vertical axis shows the number of full-time people delivering the S/390 service.


The lines on the graph show the average number of people required to run a S/390 service of a particular size in a particular year. These lines show how the number of people required to support a given amount of mainframe capacity has dropped fourfold in the last 7 years. This has been achieved through a substantial reduction in the number of data centers and the availability of automation tools (eg tape robots and clustering software). This is powerful evidence that consolidation of S/390 centers has enabled much smaller numbers of people to run much more powerful clustered S/390 machines.

Notice also that once an enterprise is on the 'current line', the incremental support cost of additional capacity is **very** low. Doubling machine capacity installed will typically only require few additional staff. **Example "E"**, with 150 service delivery people running 8000 mips, is a good example of a leading-edge, highly efficient S/390 organization.

Where would your blue star be placed on this graph?

UNIX and Intel Servers - People Efficiency

This chart requires a lot of background explanation.



'Servers supported per UNIX/Intel delivery support person' is a useful people efficiency measure in the UNIX/Intel server world. This world is characterized by three major types of server environments.

The first environment, is an **uncloned** environment. Each server may come from a different vendor, have a unique hardware configuration, different operating

systems software version and parameters, and different application version. It will typically require individual problem determination, problem resolution, performance tuning, capacity management, backup, and recovery management processes.

An average people efficiency ratio for this environment is 2.5 servers per service delivery support person. 'Best of breed' centers, using tools such as Tivoli to automate many of the management processes, achieve people efficiency ratios of around 5 servers per person. This is the upper 'green' shaded area on the chart.

The second environment is **super cloned**, where there may be hundreds or thousands of retail store and bank branch servers. Each server has an identical hardware configuration, identical operating system and application version, and identical, automated management processes for backup, failover and recovery. If a server breaks, it fails over to an identical backup server, and the failed server is repaired or replaced.

This degree of cloning and automated process is essential to managing the thousands of servers that may be involved in these systems. However, high productivity, of 30-50 servers per support person, is achievable. This is the lower 'white' area on the chart.

Finally there is the **cloned** environment. The RS/6000 SP is an excellent example of this. Individual servers or nodes are cloned wherever possible. Typically all the servers are from the same vendor, and have identical hardware configurations, operating system and application versions. They also have automated systems management and software distribution.

Also, although individual nodes may run a variety of different applications or functions, each function is cloned whenever possible. This has a dramatic effect on productivity. Here, an average people efficiency ratio is about 5 servers per person, with a best of breed ratio around 10-12 servers per person. This is double the productivity of uncloned servers, and is represented by the middle 'blue' shaded area on the chart.

Example "F" has 150 UNIX servers, of which 102 (68%) are cloned, performing functions such as firewalls, systems management and network management. The remaining 48 are uncloned UNIX servers handling functions such as application and database serving, as well as development and test serving. As a cloned environment should be twice as productive as an uncloned environment, we can estimate UNIX people efficiency by calculating, in this case, an


'equivalent' uncloned environment of 99 (102/2 + 48)
equivalent uncloned servers.

A series of papers on the role of advanced high-end IBM servers

There are 45 people providing service delivery to these 99 'equivalent' UNIX servers. This is '**below average**' UNIX people efficiency, the yellow star on the graph. 20 people should be able to provide a best-of-breed, highly automated service for these 99 'equivalent' UNIX servers.

Example "G" has 550 Intel servers, of which 300 (55%) are cloned, performing functions such as file/print, network gateways and systems management. The remaining 250 are uncloned, providing a major groupware service and individual applications and database servers. This means there are the equivalent of 400 (300/2 + 250) uncloned Intel servers. There are 120 people providing service delivery for these 400 Intel servers. This is '**very good**' Intel people efficiency, the red star on the graph.

So, where is your organization in terms of UNIX or Intel people productivity?. Where would your yellow star and red star be placed on these charts?



B: System Efficiency - Server Utilization

In addition to people efficiency, it is important to look at server efficiency - are we driving the IT capital assets as efficiently as possible?

S/390 Servers - Processor Utilization


S/390 processors run multiple applications and many virtual servers, and typically run at very high average system utilization. There are several slightly different metrics used to measure S/390 utilization.

For capacity planning purposes, it is always important to measure peak hour server utilization, so that the appropriate time to increase capacity can be estimated. Sustained peak hour utilization on S/390 is typically in the range of 85% to 100%. S/390 mainframes are regularly driven at 100% utilization for several hours per day. As they get overloaded with work, they are designed to degrade 'gracefully' with response times increasing but with no catastrophic service failure.

For the purposes of this health check we use a different but also readily available measure. The graph shown has measurements every 30 minutes over a typical 24 hour period, starting at midnight, of the utilization of the S/390 servers. This graph is very useful to show the characteristics of the work. Is it 'office hours' or a 24 hour service. How spiky is the work, measured in terms of the peak: average ratio? What is the peak hour utilization, average prime shift utilization and 24-hour utilization?

Example "H" has average prime shift utilization of 57%, 24-hour average utilization of 50%, and peak utilization ranges between 81% and 96%. This is

H: Server Utilization - S/390 Mainframes



slightly below the industry average; our findings show that the average for S/390 prime shift utilization is about 70%, 24-hour utilization is about 60%, and peak hour utilization ranges between 85% and 100%. Example "H" looks consistently about 10% below average utilization, but this is not surprising because this enterprise recently installed additional capacity to handle a significant expected load increase in 2000. So what are the corresponding measures for your mainframes?

Prime shift S/390 utilization is ...%

UNIX and Intel Servers - Processor Utilization

UNIX and Intel servers usually run one single application or function. Consequently the question of utilization has to be addressed by obtaining measurements of many servers and combining them.


The value of the '24-hour' graph shown earlier is that it shows the peaks and troughs, and that it can also help determine how often measurements are needed. A lightly utilized gateway may only need to be measured once per quarter, but a heavily utilized database server may need to be measured every day.

A series of papers on the role of advanced high-end IBM servers

Examples "I" and "J" show a variety of UNIX workloads, some 'daytime' OLTP, some gateways, some data warehouses. These are very unusual profiles to

Aggregate prime shift UNIX utilization is ...%

I: Server Utilization - Unix Production



someone familiar with mainframe OLTP profiles. It is very important to understand peak, prime shift and 24-hour utilization, and individual server workload growth rates. Unfortunately, this usually means substantial manual effort to collect and analyze the measurement data.

Aggregate (total average) prime shift utilization across all UNIX servers is typically in the range of 10-15%. This is frequently a major surprise to IT executives. But then they realize that many systems are low utilization firewalls or gateway infrastructure, and that many UNIX services are typically daytime only. The major exceptions to this are heavily used application servers, database servers and data warehouses.

However, even very heavily utilized servers tend to have prime shift average utilization of only 30-50% (see chart), and, since most have an unused backup server, average prime shift utilization for heavy servers alone is around 20%. A 24-hour average is typically half this - often around 10% at best. So, although peak hour utilization of the most heavily loaded server is often 50-70%, the aggregate prime shift utilization across all 150 servers in the above examples is nearer 10-15%, and over 24 hours less than 10%

Intel server utilization is a more difficult question to answer. Few measurements are taken on a systematic basis, although Intel servers often represent millions of dollars of the total IT spend (12% in example B and 18% in example C).

Ad hoc measurements of the most heavily used servers show about 30% peak hour utilization. Occasional measurements indicate an aggregate prime shift utilization of between 5% and 10%. 24 hour utilization is between 2% and 5%. So, what is your aggregate prime shift Intel server utilization?

Aggregate prime shift Intel utilization is ...%

C: Quality of Service - Server Availability

Quality of service usually comprises a number of sub topics; scheduled outages, unscheduled outages, serious incidents, and response time. These are important and closely related, because sudden spikes in transaction traffic can overload less robust systems, and cause them to fail suddenly.


In general, most mainframe sites have substantial automated measurement and tracking systems, UNIX servers have few, and Intel have less. Expectations are also different; mainframe (typically on-line transaction based) systems are expected to deliver at least 99.9% availability nearly 24 hours every day, whereas many of the systems implemented on many UNIX and Windows NT servers so far have required only 'good enough' availability for largely 'office hours' computing.

S/390 Servers - Availability

You can assess your mainframe service quality as poor, average, good, or excellent against the following major characteristics:

S/390 scheduled outages: Typical scheduled outages for a well-run S/390 data centre vary significantly. One type of centre may have about 160 hours scheduled outage per year by taking '3 hours every Sunday morning' as the window for all mainframe hardware and software changes. Such scheduled availability is **average** for data centers not having a requirement for

J: Server Utilisation - Unix Production



A series of papers on the role of advanced high-end IBM servers

continuous availability. This type of computing could usefully be called '24 x 6.9' or perhaps '24 x 358' computing.


Leading-edge, **excellent** data centers are driving scheduled outage time down to around 12 hours per year, using sysplex cloning for the introduction of new releases and changes. This style of computing for continuous availability could be usefully called '24 x 6.99' or even better, '24 x 364.5'. It demands leading-edge mainframe clustering technology.

S/390 unscheduled outages: This is the acid test of the capabilities of a server platform (processor and disk hardware, operating system, database and middleware, application software and systems management processes.) Individual components may fail but it is the capability of the system to recover seamlessly from any failures and continue running that is a key requirement for major OLTP applications.

A good rule of thumb for **average** mainframe service, and a frequent management target, is a minimum of 99.9% application availability throughout committed service hours, often 24 hours a day. This is regarded as a minimum acceptable level of service.

Notice that 99.9% availability during normal working office hours represents an outage of only 2.4 minutes per week or 2 hours per year. Users get extremely dissatisfied if availability drops below 99.7%. Although these outage times appear small, an application outage disrupts many users, requires transactions to be restarted and real, revenue-generating customers get **very** dissatisfied at the point of contact.

K: S/390 Mission-Critical Application Availability



The major components of unscheduled outages are operator error (which has to be solved by automation) and application failure (which has to be solved by more robust application design). However, the underlying platform must also be reliable and have robust recovery processes. **Example "K"** shows an **excellent** data center achieving 99.95% application availability measured over a whole quarter or year on a regular

basis. In the mainframe world the expectation is that response time to 'normal' transactions is subsecond, even during periods even when server utilization is 70% to 100% busy during peak hours.

Transaction volumes vary hugely but it is not unusual to observe 5-10 million daily CICS/DB2 transactions, and annual transaction growth rates in the range of 20-50%.

S/390 Prime Shift Availability is 99. ... %

UNIX and Intel Servers - Availability

One of the challenges of comparing server platforms is the lack of real UNIX and Intel availability measurements. There are several reasons for this. The first is the general lack of automated tools to provide a meaningful base of measurement data. More importantly, the sheer number of servers that need to be measured. Compared with 4-6 mainframes, measured daily, in many different ways, there are frequently 100-200 UNIX servers and 500-1000 Windows NT servers which have few measurements.

This highlights the third reason. The usual practice in the UNIX and Intel world is to acquire servers with plenty of capacity and only to measure those providing critical services or those with known response time problems. This has been aided by the fact that until recently most of the UNIX and Intel servers have been running relatively non-critical services.

An inventory of servers is a vital starting point. Divide this inventory into two categories, firstly, application, database, mail and development servers, and, secondly, IT infrastructure servers, such as gateways, firewalls, domain, name, security, and file/print servers. Typically about 30% of the servers are in the first category, 70% in the second. Infrastructure servers are often relatively lightly loaded and often have a backup cluster server if really critical.

UNIX and Intel scheduled outages: Many of these servers run mainly during office hours. With such servers, it is normal to schedule a planned outage whenever required, typically at a weekend or Sunday morning. There is no particular concept of an overall scheduled outage and consequently it is not frequently measured; "we just schedule time to apply updates or implement hardware changes when necessary". Many systems are said to be 24 x 7 but in practice are often in the range of 12 x 5, 16 x 6, and occasionally 22 x 6.

A series of papers on the role of advanced high-end IBM servers

However, as businesses move towards real 24 x 7, or 24 x 6.9, the change window becomes increasingly critical. Many servers are now being upgraded to clusters for availability reasons, and also face the tough challenges of rapidly shrinking change windows, which simultaneously drive the need for rapid, automated, reliable backup, recovery and software distribution methods.

An **average** UNIX service delivery organization has typical weekday service hours of extended office hours (08:00-20:00 Monday to Friday) for the bulk of the application and database servers.

A useful classification of service levels based on scheduled outages is:

Category	Basic	Average	Best
Daytime window	12 x 5 108 hrs/week	16 x 5 88 hrs/week	18 x 6 60 hrs/week
Overnight window	22 x 6 36 hrs/week	24 x 6 24 hrs/week	24 x 6.5 12 hrs/week
Enterprise window	24 x 6.8 5 hrs/week	24 x 6.9 2.4 hrs/week	24 x 6.99 15 mins/week

In this chart, the first number gives the hours and days that the committed service will be provided (eg 22 hours for 6 days per week). The second number shows the resultant scheduled outage 'window' (168 hours minus committed service hours) available for change activity.

Enterprise-class computing provides committed service hours of 24 hours per day, but, some time will be taken, usually on a Sunday, for scheduled change activity. Clearly as the window shrinks it gets much harder to sustain availability without resilient automation tools and processes.

Place your principal UNIX and Intel services in the above table and calculate the total tolerable scheduled outage hours for each major type of service.

In which cells are your UNIX/Intel services today?

UNIX and Intel unscheduled outages: Very few organizations have effective measurement of the unscheduled outage time of their UNIX and Intel servers. Many servers that are not business critical are simply rebooted and the service resumes. So, it is no surprise to find the peak time for help desk requests for such servers is around 08:30 Monday morning as office users come into work after the weekend. Since there are usually no committed service hours for such servers over the weekend, there is often no record of the outage


at all, and certainly no good measure of the outage duration.

However, many UNIX and Intel servers provide key business services and consequently there are some measurements and observations available via management reports and anecdotes.

Example "L" was only able to achieve 98.5% application availability of a business critical multi-tier client-server system after spending millions of dollars ruggedizing the system and putting new systems management processes in place. The problem was only partly hardware availability. It was also the complexity of data recovery and data synchronization during recovery from failure, the fragility of certain combinations of versions of hardware, microcode, operating system, database and applications, and the lack of automated application recovery processes.

Example "M" has UNIX application availability measurements over a period of 24 months and finds an average prime shift UNIX application availability of 99.66%. Typical UNIX cluster application availability is 99.6%, compared with mainframes at 99.99%. Including the effects of application outages and operator error reduce this number. A range of 99.5-99.6% is a realistic expectation for prime shift application availability in a well-run UNIX environment

Finally, **example "N"** is a very well run IT organization delivering an advanced Domino/Windows NT groupware service via 150 Windows NT servers, including 70 mail servers, to 10,000 users worldwide is able to sustain 99.8% application function availability. But because a typical user must have a mail server, a database server and a workflow server all available at the same time to handle key business functions, total availability delivered to the users is actually only 99.6%.



Aggregate UNIX prime shift availability is %

Aggregate Intel prime shift availability is %

Cost of one hour of prime time outage M\$

3. Total Cost and Incremental Cost

Cost analysis is an important but somewhat arcane sport. It is vital to an understanding of the efficiency and affordability of various IT services. The pie-chart examples at the beginning of this paper showed instances of different expenditure patterns in various organizations - and we expressed a view of suggested expenditure ranges on desktop, distributed servers, data network and central servers in a large enterprise.

However, although we saw in an earlier section organizations spending 50% of their total IT budget on central servers, we made no observation at that stage on the appropriate mix of spend on Intel, UNIX and Mainframe servers.

In most organizations, total **current** spend is a direct result of infrastructure decisions made 2, 5 or even 10 years ago. Saying that XYZ enterprise annually spends 50M\$ on CICS/DB2 mainframe OLTP services, 20M\$ on Windows NT and Novell services, and 10 M\$ on UNIX services tells you little about the enterprise's future IT strategic direction, since much of today's total expenditure is required to sustain historic, but important IT services.

So, a good understanding of today's actual IT expenditure is very important in determining actual, achievable people efficiency and systems utilization factors that are central to a future cost comparison.

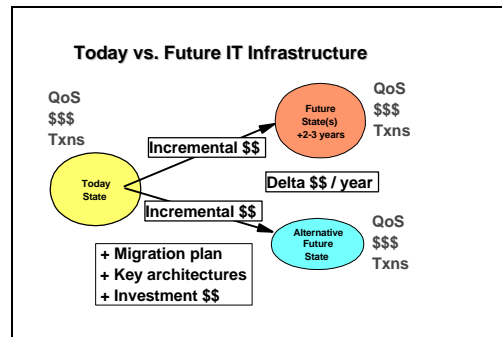
However, if you are told that future 'incremental' expenditure on new applications will be 10M\$ on UNIX and zero on Mainframe and Windows NT this gives you a **much better** indication of that enterprise's **future** technology priorities. Consequently, we need both full costs and incremental costs to assess the IT investment priorities.

The purpose of this section is to explain in outline how this can be achieved, largely using the IT health check data already collected earlier in this report. (More detail

on incremental cost analysis is planned for a companion paper in this series.)

It is important to distinguish between three or more 'states', the today state, and two or more future states. Today's state comprises the current technical IT infrastructure, (server platforms, applications, transactions, utilization factors, network, etc), its full cost, and delivered service quality. These have been the principal content of earlier sections of this paper.

For each future state we need similar information. Each future state description will need to include a full technical IT infrastructure, a total cost, and estimated service quality. This can be simplified into the **incremental** changes that are needed to get from today to each future state - incremental changes in technical configuration, and incremental changes in people costs, software costs, hardware costs and service quality.



In practice, we normally want to evaluate a number of future states for a particular service; for example based on mainframe servers, UNIX servers and Intel servers. This is the basis of incremental cost analysis.

Incremental Cost Analysis

Let's take a real example, to illustrate the importance of this analysis. Because it is actually quite hard to explain the whole approach, it is easier to show the 'answer', and then explain in outline how we get there. The detail will be covered in the companion paper.

The workload used for this example is a new, OLTP-like workload, but this analysis can also be done for e-mail or ERP or any specific application or function. It is assumed that all these incremental scenarios are properly configured to deliver the same transaction volumes, response time and target service quality, such as 22 x 6.9 and 99.8% availability.


A series of papers on the role of advanced high-end IBM servers

Note that all these costs are reduced to annualized operating costs, and include hardware, software and people costs. The actual results were as follows:

Scenario A: This technical configuration assumes the new application runs on the latest, most powerful Intel servers and Windows 2000. This scenario assumes no significant improvement in service delivery people efficiency or in server system utilization. The incremental cost of this scenario is dominated by ongoing service delivery people cost. Scenario A costs 100 units per year (eg 5 M\$ / year)

Scenario B: This technical configuration assumes the new application runs on the latest, most powerful Intel

Incremental Cost Comparison - Annualized



servers and Windows 2000. This scenario assumes a doubling service delivery people efficiency or server system utilization over a 2-year timeframe. This incremental cost is still dominated by service delivery people cost. It is in the range of 54 - 66 units of cost per year (eg 2.7-3.3 M\$ / year).

Scenario C: This technical configuration assumes the new application runs on the latest, most powerful UNIX servers and AIX. This scenario assumes no significant improvement in service delivery people efficiency or in server system utilization. This incremental cost is also dominated by ongoing service delivery people cost. It is about 71 units of cost per year (eg 3.6 M\$ / year).

Scenario D: This technical configuration assumes the new application runs on the latest, most powerful UNIX servers and AIX. This scenario assumes doubling service delivery people efficiency or server system utilization. The incremental cost of this scenario is still dominated by ongoing service delivery people cost. It is in the range of 46 - 51 units of cost per year (eg 2.3-2.6 M\$ / year).

Scenario E: This technical configuration assumes the new application runs on the latest, most powerful S/390 mainframe and OS/390. This scenario assumes no significant improvement in the current excellent service

delivery people efficiency or in server system utilization. This incremental cost is dominated by incremental IBM hardware and ISV software cost. In this specific case, the total incremental cost is in the range of 27 - 31 units of cost per year (eg 1.4-1.6 M\$ / year).

Observations on Incremental Cost

This information is new and relatively surprising. Notice that substantial, 30-45%, expenditure savings can be made by consolidating from current Intel servers (scenario A) to more powerful Intel servers (scenario B), or from current UNIX servers (scenario C) to more powerful UNIX servers (scenario D).

These savings can be achieved by consolidating many 'cloned' server functions, such as mail, Web, or file/print serving into a smaller number of higher utilized centralized servers. The main effect is to ride the scale and price/performance curve, although displacing many recently acquired servers can lead to significant book-value write-off issues. By reducing the number of physical servers, and through automation, the total people cost can be reduced.

However, in view of the typically low current utilization and low service quality of many current server farms, scenarios B and D require substantial validation with IT service infrastructure people to ensure that the claimed utilization and people savings are really achievable.

But also note the very substantial (60-70%) operational savings achievable by considering S/390 (scenario E) compared with the current state. This approach can be considered not only for functions such as Domino, file/print, and Web serving but also for new OLTP, e-business TP, and ERP applications and database servers.

Clearly actual incremental costs will vary considerably from situation to situation, and, depending on the workload type and service quality requirements, S/390 may not always be the lowest cost answer. Nevertheless it is worth considering all these server platform alternatives in major application platform decisions.

What are the key factors in this analysis?

- A good technical analysis and outline design of the future state is an essential starting point. Enough detail, or a set of realistic assumptions, is required to identify the main transaction volumes, main software options and likely service quality. Typically three future configurations will be needed, Intel, UNIX and S/390, each of them capable of handling, say, 1000 units of the application load.
- The difficult next step is to ensure comparability of the performance and throughput of the Intel, UNIX and S/390 configurations. In most cases industry standard benchmarks, (such as TPC-C, SpecInt, SD steps/hour), actual benchmark data, or actual measurements can be used.
- The future assumed utilization factors, need to be realistic and based on current values achieved. In the S/390 world most people are comfortable assuming 70-80% load factors, as S/390's regularly run at 90-100% utilization during prime shift. This means that 1000 units of workload will require an additional $1000/0.75 = 1333$ units of S/390 capacity.

UNIX requires separate servers for separate workloads and the need for empty standby backup servers. Very busy servers may have an actual peak load of 70%, but this is typically reduced to a prime shift average of 40-50%, and further reduced to 20%-25% when the backup server is included. An assumption of a prime shift average of 20-25% is appropriate for very busy servers, noting that measured aggregate UNIX server utilization is typically 10-15%. 1000 units of UNIX workload can realistically require 4000 to 5000 units of actual UNIX server capacity.

A similar rationale applies to Windows NT servers. Peaks on individual Windows NT servers are often in the range of 30-40%, compared with UNIX at 70% and S/390 at 100%. An Windows NT server with prime shift average utilization of 20% is reasonable, but adding an empty backup server takes the server average to 10%. It is easy to see how current Windows NT aggregate prime shift utilization's of less than 10% arise. 1000 units of Windows NT workload can realistically require 5000 to 6666 units of actual Windows NT capacity.

- The final key assumption is the number of IT support people to deliver the S/390, UNIX or Windows NT service. Clearly today's actual experience relative to industry averages is a very important starting point. There are many measurement points over many years in the S/390 arena. With UNIX and Windows NT, people support costs dominate the platform cost. In the case of UNIX, we find total people costs usually account for 35-45% of the total IT spend on the platform. This can be compared with 45-55% in the case of Windows NT, and 20-25% for S/390. It is worth noting that the heaviest people costs are incurred in supporting uncloned application and database servers, which require people-intensive monitoring, problem determination, tuning, backup, and recovery processes.

A proper financial analysis usually subjects all these assumptions to major 'what-if' analysis. It is surprising how relatively insensitive the model is to hardware costs. On S/390 systems most attention is required on the software contract clipping levels. On UNIX and Windows NT, the key drivers are achievable utilization factors and people. Scenarios B and D above contain relatively aggressive assumptions over the 2-year timeframe.

4. Conclusion

A clear perspective on today's actual IT infrastructure cost, by technology platform (eg Windows NT, UNIX, Mainframe) is extremely important in determining future IT strategy and investment decisions. This report has shown a practical method for determining these costs in your enterprise today. We recommend the following four action points:

1. Establish a realistic estimate of the actual IT cost per user per year, and check to see that it is in the range 5-10 K\$ per user per year. Anything higher should be considered a symptom of gridlock unless **really** exceptional IT circumstances prevail. Check the balance of spend by desktop, network and servers to see if the current spend is unbalanced, as in examples C and D.
2. Conduct an IT health check of all major server platforms. Establish actual people efficiency, systems utilization, quality of service and total cost. This report contains charts and rules of thumb that will help position where you are. If it too hard to do a comprehensive study, assume you are average and make sample checks to validate the numbers of people and utilization factors. You should be aiming for

'best-of-breed' people and system efficiency ratios and high service quality on all platforms.

Hasta la vista.

David G Heap, Somers, USA

3. Build a simple financial model to estimate the incremental cost of each server platform. This is best if based on a specific project or application but can be cross-checked against the generic model described in this report. Make realistic conservative, verifiable assumptions about achievable gains in productivity and utilization, which can then be measured on implementation. These comparative incremental costs will be a major factor in strategic IT decision making.

4. Avoid gridlock. Ensure that PC and distributed server proliferation and data fragmentation are strongly controlled. Many organizations are using browser-based computing as a major reason to move very rapidly away from expensive client-server computing. But, if executed in a piecemeal manner, current IT gridlock will prevent the effective rationalization of a currently bloated IT infrastructure and seriously slow e-business deployment. An effective IT strategy has three major components;

- A credible, simple, constantly updated IT vision and blueprint of the target IT infrastructure linked to business needs. This IT infrastructure blueprint needs to describe selected strategic technical components, such as the primary server platforms (eg S/390, RS/6000, NetFinity,) database (eg DB/2), core middleware (eg CICS, MQ), development paradigms (eg Java, WebSphere, Linux) and the data network (eg TCP-IP).
- Committed investment funding (eg 10M\$ total investment over 3 years) in server platforms and services, to enable the simplification of the IT infrastructure itself and make the transition to a robust 'e-infrastructure'.
- Third, and most difficult, is a realistic technical migration plan from today to the future state. This typically is a phased multi-year project, which migrates selected business applications and functions to target platforms. Often this cannot be achieved directly and involves the selection of technical 'stepping stones' as the locus of interim strategic service.

These items form the basis of a successful e-business IT infrastructure. We plan to journey with you in a series of companion papers to this one which will address many of the central issues and decisions in building a viable cost-effective e-business IT infrastructure. Your first step is to complete this IT health check to establish where you are today.



© Copyright IBM Corporation 2000

IBM Corporation
Marketing Communications, Enterprise Systems Group
Route 100
Somers, NY 10589

Printed in the United States of America, 06/00
All Rights Reserved

This publication was produced in the United States.
IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

You can find additional information via IBM's World Wide Web server at <http://www.ibm.com>.

IBM Hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

By "Year 2000 ready", IBM means the capability of a product, when used in accordance with its associated documentation, to correctly process, provide and/or receive date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange date data with it. IBM's statements concerning the Year 2000 readiness of its products are Year 2000 Readiness Disclosures under the Year 2000 Readiness Information and Disclosure Act of 1998.

® IBM, IBM logo, AIX, CICS, DB2, NetFinity, OS/390, RS/6000, S/390, and WebSphere are registered trademarks of the International Business Machines Corporation.

Domino is a trademark owned by Lotus Development Corporation.

Java is a trademark of Sun Microsystems Inc.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Windows is a registered trademark of Microsoft Corporation.

LINUX is a registered trademark of LinusTorvalds.

All other registered trademarks and trademarks are the properties of their respective companies.

All customer examples cited represent how some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

GF22-5168-00