

## The Dinosaur Myth — 2004 update



With all the recent announcement activity, and the anticipated push from IBM to extol the virtues of the mainframe once again, we can expect many users to re-visit the platform choice issue. To help them in this task we have updated our 2002 *Dinosaur Myth* paper <sup>(1)</sup>, which concentrates on the total costs of the different platforms, and we have added some information on the less financially-focused aspects of platform choice.

We have always maintained that, in comparing the TCO of the competitive platforms, what really matters is the number of users that a computer can support, performing whatever functions are necessary to the organization with a reasonable level of service. So the key yardstick of a computer's cost-effectiveness is not TCO (total cost of ownership) but TCU (total cost per user), measured over a reasonable time-span to even out any high up-front costs. Five years is a reasonable period for this purpose.

This distinction between TCO and TCU is significant, as a mainframe can support many concurrent users performing different tasks and changing between tasks on a real-time basis. To help them achieve this multiple-application capability, mainframes have evolved mechanisms for the efficient sharing of resources among large numbers of concurrent users. In particular, they have multiple 'interrupt levels', permitting them to switch from task to task without losing track. This means that a task waiting for an external event (a transaction from a terminal, or a data transfer from a disk drive, for example) can be suspended and returned to later, while other tasks are attended to in the meantime. They also have very sophisticated resource management capabilities which allow the users to have their work completed on a priority basis, such that even when fully loaded the key applications get the capacity needed to perform the task in hand. These resources can be reallocated literally second-by-second to achieve this goal. Unix systems and PCs do not have such sophisticated mechanisms, although from some vendor claims users could be forgiven for believing differently.

(1) The *Dinosaur Myth* has been regularly updated in recent years as part of the Mainframe Market Information Service. The *Myth* was first published by Xephon, and is now maintained by Arcati ([www.arcati.com](http://www.arcati.com)) in association with the author Barry Graham.

This ability for all work to be completed on one system is crucial to minimizing the TCU. For example, let's say that ten applications are each used by all of the staff. If each application required up to 10 MIPS of capacity, but in total no more than, say, 20 MIPS were needed at peak load, then on a mainframe 20 MIPS would suffice, whereas in the Unix or PC case a minimum of 10 systems of 10 MIPS each would be required, assuming they could run effectively at 100% utilization. This equates to five times more capacity in total than on the mainframe.

The assumption that all platforms can run at 100% utilization is totally unrealistic, as neither PC servers nor Unix systems can run effectively at anywhere near that level. Indeed, at anything above 50% utilization response times suffer and system failures occur. So in the previous example, we must double the required PC/Unix capacity, making a total of ten times the equivalent mainframe power. This also increases the storage and support requirements and by default lowers the availability, as the more complex the environment the more likely the system is to fail. This may seem to be an exaggeration, but within my client base an overall effective utilization of 10% or less for the non-mainframe platforms is quite normal.

This all negates the simple price comparison of alternate solutions, as cost per MIPS does not reflect the real-life utilization differences between the platforms.

## **The true costs of computing**

The true costs of computing fall into three main categories:

1. the cost of the hardware (including terminals, printers, and other peripheral devices) and the basic operating software, over a reasonable period. This figure should include the cost of maintaining the hardware over that period, and incidental costs like office space, electrical power, special cooling requirements, etc.
2. the cost of the application software — the off-the-shelf packages or customized programs that allow the computer to perform useful work.
3. the personnel (manpower) costs associated with operating the hardware and software and handling support. To this should be added the cost of any time wasted waiting for the computer system to respond.

There are, in addition, other costs that can be directly attributed to computer systems, which may not be so readily quantified but should also be considered. These will be touched on later.

## **Hardware and basic software costs**

The following comparisons are based on a representative selection of systems from among our clients performing the same or similar tasks:

- \* various mainframe configurations supporting thousands of users
- \* several Unix servers from different vendors supporting similar numbers of users
- \* a selection of PC servers from different vendors supporting similar numbers of users.

We calculated the basic hardware, software, and maintenance costs over five years for these systems (excluding the cost of finance, and ignoring inflation). Our estimates per end-user were:

* Mainframes:	\$4,500
* Unix minis:	\$5,400
* PC servers:	\$8,000

Already it can be seen that the alternative platforms do not have any advantage over mainframes even at this level. This is largely due to the additional capacity required, as outlined above.

These figures may surprise some readers, who have perceived the mainframe (and mainframe software in particular) as being exceptionally expensive. But if you need ten times less capacity to perform the same work, the perceived hardware and software price disadvantage soon evaporates. The software cost in particular for the mainframe is very visible, with large quarterly payments to IBM (and increasingly ESSO/ELA type long-term contracts, which focus the user even more on the total IBM software cost) and relatively frequent large upgrade charges from the ISVs. In contrast, the software costs for the other platforms are more distributed and/or bundled into the hardware, and are consequently far less visible. However from our studies the mainframe software costs are typically much lower than those of other platforms on a per-user basis, while contributing more to the business.

The Unix and PC server cost figures used here are higher than those that would be typical for very small numbers of users, as it has become apparent that neither Unix nor PC server systems are truly scalable at the same level of cost. By that we mean that, as the number of users increases, the cost per user also increases. Our own estimate, based on extensive research, is that for a doubling of the number of users the costs increase by close to 125% on a non-mainframe platform but by only 90% on a mainframe.

In our earlier TCU comparisons we had to add something for the cost of floor-space and special environmental requirements. This was typically higher for the mainframe than for Unix systems, with no costs under this heading for PCs, since they occupied much the same space as the terminals for mainframe and minicomputer systems. For mainframes, we also needed to add the cost of the network hardware and software, whilst the equivalent costs for interconnecting minicomputer systems and PCs were very variable, and any figure we proposed would have been open to dispute.

But today all of these costs are similar on each platform, with the mainframe if anything proving the cheaper when compared with the server farms that have evolved for other platforms.

## **Application software**

The application software required will obviously vary widely for different organizations. However with most mainstream packages available across all platforms and most platforms requiring similar levels of 'tailoring' of applications, the per-user costs today are similar across the board. A figure of around \$750 per user per year is the average for the clients studied to date. Adding these costs into our overall TCU comparison, we get the following approximate figures:

*	Mainframes:	\$5,250
*	Unix minis:	\$6,150
*	PC servers:	\$8,750

## **Personnel costs**

All computer systems require some human supervision, ranging in complexity from loading the printer with paper to diagnosing and fixing hardware or software faults. End users may be able to handle some of this work themselves, but even the most independent will occasionally require the assistance of specialist staff. At the other extreme, end users supported by mainframes are largely shielded from both the complexities and the chores involved in tending to the computer's needs — instead, full-time specialists are employed. Unix systems fall somewhere between these two extremes.

The staff costs of running mainframes are very visible: operators and technical support staff do nothing but minister to the mainframe, and their salaries and employment costs are easily identified. Current mainframes on average require one technician (systems programmer or operator) for every 250 end users, which, at an average employment cost of \$75,000, amounts to \$1,500 per end user over a five-year period.

Two points are worth making here. First, the number of operators and systems programmers required per mainframe MIPS has fallen ten-fold in the past seven years, and is expected to at least halve again in the next five years. Second, the estimate we've adopted assumes multi-shift 24x7 operation, which means that the batch work typically carried out overnight is included in the cost, even though we have made no allowance for it in our cost-per-end-user calculations.

For Unix systems, fewer technical staff are required to tend the system because they do not normally operate 24 hours a day, and our research puts the level at close to one person per 500 users. At the typical operator employment cost of \$75,000 per head per annum discussed earlier, we are looking at around \$750 per end user over five years. In addition, it is generally reckoned that, on average, one full-time support specialist is required for every 100 end users in a typical Unix environment, which is close to the three-to-one ratio reported by our clients. If that specialist also costs \$75,000 a year to employ, the five-year cost per end user will amount to \$3,750. Putting the operational and support needs together, we get a figure of \$4,500 per user over five years.

In many cases in the PC environment, the end user is the operator. It's his or her responsibility to take back-ups, copy files, put the appropriate paper in the printer, look up error messages in the manual, and so on. It is estimated that the average PC user spends one hour a week, or 12 minutes a day, either tending to the system or waiting for a response from it. This seems like a conservative figure to us, but it nonetheless equates to 2.5% of the end user's time. How much that might cost will of course vary depending on the end users' jobs, but if we assume a minimum annual end-user employment cost of \$36,000, 2.5% amounts to \$900 a year, or \$4,500 over five years.

To this must be added the cost of specialist support staff or local help within the user group to help out when users are unable to solve a problem themselves. A common guideline is that PC-based systems require the equivalent of one support person for every 50 PC users today — which may seem a high figure, but is in fact only the equivalent of 2% of each end user's time. At \$45,000 a year (an average between the cost of the end user and the technical employee), that costs another \$4,500 per end user over a five-year period.

Adding these estimates to the running totals for TCU gives the following results:

*	Mainframes:	\$6,750
*	Unix minis:	\$10,650
*	PC servers:	\$17,750

Arguably we also need to include an element in the manpower costs for time wastage while waiting for the system to respond. The Unix applications on which these figures are based are generally slower than their mainframe equivalents, offering response times in the 2-4 second range. An average extra delay of two to three seconds for every interaction (which we've assumed will take place on average every 45 seconds) equates to a 5% overhead. It's unlikely that end users will be able to do any useful work during that time, so in effect the Unix solution can levy a hidden cost equivalent to 5% of all end users' time (their salary plus other employment costs). At \$36,000 per person per year this adds a minimum figure of \$1,800 a year to the Unix system costs, or \$9,000 over the five years.

PC server-based systems are typically no better in this respect, and indeed often far worse, but for the sake of this comparison we will assume a 'hidden' cost of \$9,000 over five years for these systems as well.

If you accept the argument that a fair cost comparison should take account of time wasted because a computing system is slower to respond, we now have the following estimated five-year costs per end user:

\* Mainframes: \$6,750  
 \* Unix minis: \$19,650  
 \* PC servers: \$26,750

For information, the following are the UK sterling figures for the same cost areas:

	Mainframe	Unix	PC
HW/SW	£3,000	£3,600	£5,250
Application SW	£750	£750	£750
Support/overheads	£1,200	£3,600	£7,200
Response time	£0	£7,200	£7,200
Total	£4,950	£15,150	£20,400

### Other analysts' figures

To see how accurate the cost split between the more visible hardware and software costs and the support/productivity costs were, we have compared the above with an IDC study for Microsoft (which will be discussed further later). This shows that, compared with the IDC average figures, we may actually have underestimated the non- hardware/software costs, although the IDC split for Web-based applications is almost identical to our split.

#### Cost split

	Unix	PC	IDC PC/Linux average	IDC web-based average
Hardware/software	31%	33%	15%	30%
Other	69%	67%	85%	70%

The dollar figures are compared below with data published in 2001 by analyst ITG, in a management brief entitled *The Cost Implications of Platform Choice*:

	Mainframe Market Monitor	ITG	Ratio
Mainframes	\$6,750	\$14,000	2.07
Unix	\$19,650	\$39,440	2.01
PC	\$26,750	\$45,000	1.68

From this it would seem that our figures are on the low side — around half of the ITG total. But in terms of the relative costs they show a very similar picture, with the mainframe by far the most cost-effective platform. The reason for the figures being lower in our case is that they are based on very large organizations, where the cost-per-user is lower through simple economies of scale. These large users are also more efficient in their use of all platforms than the average user represented in the ITG case.

In effect, the PC solution, and to a lesser extent the Unix solution, move much of the personnel cost out from the data center to the end user. That has two consequences: first, it tends to hide the costs, which are absorbed into other budgets; and, second, it increases the total costs, because more people are involved in identical 'housekeeping' activities. For example, whereas a mainframe system will back up the data of thousands of end users either automatically or with minimal operator intervention, with the PC solution each individual, or perhaps each workgroup, has to initiate the process.

We should emphasize that all our cost estimates ignore inflation and the cost of money, and assume a 'green-field' site. They cannot be compared with the budgeted costs of existing installations. For mainframes in particular, given that prices are falling all the time, systems installed some time in the past will be correspondingly more expensive, and our figures are also based on complete systems rather than upgrades, which tend to cost considerably more for equivalent performance.

Many organizations also write off capital costs more quickly than over the five years we've allowed. And as noted earlier, the figures don't take account of batch applications. If these and other salient factors are taken into account, the budget costs over a five-year period for all solutions will be higher than the figures we've quoted here. However, the relative costs would certainly not change as the figures we have calculated are conservative for the non-mainframe solutions.

## **New mainframe benefits**

There are also some other significant advantages materializing for the mainframe world as users move further into e-commerce. For example, in the e-commerce world most transactions must be handled at a secure level (the SSL or Secure Sockets Layer) when the transaction actually takes place. But non-mainframe systems are between 15 and 100 times slower in this mode. You will almost certainly have experienced this whilst waiting for an on-line purchase to be approved. On a mainframe there is no noticeable degradation of service at this time.

Possibly most significant of all, non-mainframe solutions do not provide scalability with linear cost increases, as we stated earlier in this report. For example, large Sun or HP systems typically cost 125% more per user than the smaller ones. We believe from our research that all non-mainframe servers will exhibit this same tendency, with the actual cost per user increasing as the number of users increases. This means that, in practice, all of the Unix and PC costs cited in this report should be increased substantially if it is your intention to support more than a few thousand users.

The problem with IBM mainframes in the past was that a relatively small number of models covered a very wide power spectrum, from perhaps 50 end users up to 25,000. Unless its workload was growing very fast, an organization could be faced with a much larger, and more expensive, upgrade than it really needed in order to add another application with a small number of end-users. In that situation, the cost of the upgrade could seem exorbitant compared with the cost of a Unix or PC server system capable of handling the new application. The temptation was, of course, to implement the new application separately on a smaller free-standing system. When the need for another new application arose, the same logic applied. And so it went on — in time, the organization had both a mainframe and a number of separate smaller systems running individual applications, all of which could far more economically be accommodated on a larger mainframe.

In many cases, removing small applications from the mainframe does not reduce the required mainframe capacity anyway! The reason is simple: in most organizations 20% of the applications take 80% of the capacity, and 100% when peaking. The smaller applications use up 'idle' time between peaks. Removal of such applications therefore has little or no impact on the overall capacity needed.

Now, however, with IBM's Parallel Sysplex architecture and Workload License Charging (WLC) fully implemented, users are able to add quite small increments of processing power to their mainframe relatively cheaply. This is one important development that has occurred recently, which helps to reduce the single significant drawback of mainframes that we identified in the early 1990s.

## **The best of both worlds?**

Another major new opportunity for mainframe users is the availability of Linux on the zSeries. This solution is a half-way house, as it brings many of the mainframe's advantages to the 'open' world. It allows users to run multiple Unix applications on a single system; take advantage of many mainframe resource management capabilities to direct all of the capacity to where it is needed; and allow literally hundreds or thousands of simultaneous servers to be accommodated on one system. In this mode it eliminates many of the support and management limitations of the massive Unix and NT server farms that have materialized in many organizations recently.

But what does Linux cost compared with a traditional mainframe-based system? In some environments, for example supporting many different Web applications on a single system, it could be as low as half the 'standard' mainframe cost. This would give it a six- to-one price advantage over a comparable Unix or PC server system, which is in line with other studies on this topic. However, if we're comparing the costs for a traditional application with some e-commerce on the side, the figure is likely to be close to the normal mainframe cost.

The main advantage of Linux is that it allows the traditional mainframe user to add new applications to the existing environment at very low incremental cost. Combined with WLC (Workload Licence Charge), this eliminates the sole problem that we found with the mainframe in the past — the cost of adding a single small application to an existing system.

Never content with one solution to a problem, however, IBM has added yet another alternative with z/OS-e, which will enable users to implement e-commerce applications using the latest middleware, with full z/OS control of the resources — better even than Linux in many cases but at an even lower cost than traditional mainframe solutions.

## **Future cost trends**

Even if the mainframe is the most cost-effective alternative at present, what about the future? Everyone believes that the price/performance of PCs is falling far faster than that of Unix systems, which in turn is falling faster than the price/performance of mainframes. It therefore may be argued that sooner or later — and probably sooner — any advantage the mainframe may have will disappear.

In fact, the opposite seems likely in the future, just as it has proved in the past, because of another less widely publicized trend. PC servers and Unix servers still lack much of the functionality that mainframe users take for granted. These missing functions, along with the performance overhead that they impose, are being added in successive software releases.

Meanwhile, the cost per MIPS of the mainframe is falling steadily at 25% to 40% a year. And the emphasis in system software development is less on adding functionality (much of which it already has) and more on improving performance, in particular taking advantage of new hardware features such as 64-bit storage. As a result, the same system running a later release of the operating system can show significant performance gains over a system running an earlier release (for example, z/OS is said to yield a gain in throughput over OS/390 of up to 10%).

One more or less certain trend is that staff costs will continue to rise in real terms. Bearing in mind that staff-related expenditure currently accounts for around 68% of the total costs for both PC and Unix server systems, compared with around 21% for mainframes, the relative effect on PC and Unix server costs will be over twice as great as on mainframe costs. And, as we remarked above, the established trend is for mainframes to require fewer technical staff each year.

Taking all these factors into account, our predicted average five-year cost per end user in 2010 will be as follows:

- \* Mainframes: \$6,250
- \* Unix minis: \$19,000
- \* PC servers: \$24,000

## **Conclusion**

We believe that our cost estimates are realistic, and if anything understate the financial advantages of the mainframe. For example, our choice of a five-year period is very flattering to the other platforms, which rarely last so long. If any kind of intercommunication or data sharing between systems is required, then more powerful Unix or PC servers, or more of them, would be required — and we have taken no account of this in our costings. Nonetheless, we would not claim that our figures are universally applicable. They should instead be viewed as a checklist of the costs to take into account in making a meaningful comparison between different systems. If such a comparison is carried out without bias, we believe that the mainframe will prove to be the cheapest option for all but the smallest multi-user systems (typically those supporting 50 users or less if new hardware and the latest software are required, otherwise 25 users or less).

With all of the recent changes, including those noted above, the mainframe has begun a new life, with existing users growing faster than ever and new users joining the fold every week. As a result IBM has grown its mainframe processor revenues consistently for the past few years, while the competition has seen its server revenues almost halve in the same period. As a direct result (and because of early retirements), mainframe skills now command a salary premium, with employers fearing a shortage of qualified staff, according to a front page article in Computerworld (March 4 2002). Not only did the mainframe not die, but it has re-invented itself and is now set to dominate the market for the next decade.