MAINFRAMES AS ENTRY SYSTEMS

David C. Thewlis

6 November 1996

I INTRODUCTION AND SUMMARY 1

II RENEWED INTEREST IN THE MAINFRAME 1

III THE MAINFRAME TODAY 3

IV THE MAINFRAME AS AN ENTRY SYSTEM 4

V THE MAINFRAME AS A BUSINESS PROPOSITION 5
- a. Traditional Strengths 6
  - Scalability 6
  - Economies of Scale 6
  - Reliability, Availability and Serviceability (RAS) 7
  - Management 7
  - Multipurpose computing 8
  - Capacity (Bandwidth and Multiple Processes) 8
- b. New Strengths 8
  - Size 8
  - Network computing 9
  - Applications 9
  - Flexibility 10
- c. Cost of Computing 10

VI THINKING AHEAD 11

VII MAINFRAMES IN EMERGING MARKETS 13

REFERENCES AND BIBLIOGRAPHY 15

MAINFRAMES AS ENTRY SYSTEMS

I INTRODUCTION AND SUMMARY

This report discusses the System/390 platform, architecture and capabilities with particular focus on the potential customer who has never before acquired a "mainframe" computer. Such customers may be entering a new business area; contemplating the expansion of existing business areas; or interested in consolidating existing work to achieve economies of scale and greater management control. Some particular focus is paid to customers in emerging marketplaces and countries and the appropriateness of the S/390 solution in those environments.

II RENEWED INTEREST IN THE MAINFRAME

Anyone who has been involved in Information Technology for very long has been exposed to the conventional wisdom that mainframes are dinosaurs, and well on their way to extinction. The advent of the Personal Computer -- or the workstation -- or LANs, or UNIX, or client/server -- would inevitably spell or had spelled their doom. Articles throughout the trade press, in airline and mainstream magazines and even books predicted the demise of mainframes, generally sooner than later.

The new technology which was supposed to render mainframes redundant varied from time to time and pundit to pundit, but the bottom line was generally that new technologies were substantially cheaper, far better, and that applications in these environments could be developed faster, better meet users' needs, and better promote business productivity. In perhaps the most dramatic assertion of the sort, Stewart Alsop, editor-in-chief of *InfoWorld*, predicted in 1991 that "the last mainframe would be turned off in 1996," a prediction which he described as "rash" in his last *InfoWorld* article before leaving the magazine in 1996.

Of course, mainframes haven't disappeared; in fact, rather the reverse. In 1994, sales of mainframe MIPS started to increase dramatically (40% increase over 1993). White papers began to appear challenging what had become the conventional wisdom pertaining to cost of computing, conversion costs to new computing models such as client/server, and the alleged productivity gains available only in the new environments. Evaluations of mainframe purchases revealed that the downturn in purchases in the early 1990s had little or nothing to do with their replacement by other technologies; according to data from Compass America drawn from over 2,700 installations, actual growth was disguised by an improvement in the efficiency of mainframe computing resulting in a compound decrease cost of 23% between 1989 and 1993; only in 1994 did workload growth become fast enough to start driving actual purchase numbers up.

By 1995, the same publications which had predicted the demise of mainframes were starting to research what was actually going on in the world of information technology, and publishing stories about major client/server failures, the actual cost of computing in PC, LAN, distributed, and centralized computing models, and about the resurgence of the mainframe. Meantime, IBM's sales of S/390 MIPS continued to rise; 1995 was 60% greater than 1994, and 1996 is continuing the trend. International Data Corporation (IDC) projects a yearly growth in mainframe computing of 55% to 60% through 1998, and that the level of investment in centralized resources will be 40% greater in some industry sectors such as banking. META Group projects that large systems will experience a 50% CAGR in demand through 1997 and that application growth will represent over 75% of this demand.

The causes for this growth are not hard to understand. Regardless of industry hype, most organizations continue to invest in and grow their core business applications, many of which are resident on mainframes. A recent Forrester survey of IT executives from fifty Fortune 1000 companies shows that over 80% of those companies have 50% to 75% of their critical applications on S/390 mainframes and the applications are not being moved to other platforms. Mainframe alternatives have not developed as fast as expected, have not provided the function required and have been more expensive for complex applications with large numbers of users. Where large-scale distributed computing implementations have been successful, they have frequently contributed disproportionately to mainframe growth because of demand for mainframe-resident data and resources.

It has been demonstrated by several studies that PC/LAN implementations, long cited as much cheaper than mainframe alternatives, are actually significantly more expensive when viewed on a cost-per-user basis. For example, the Information Technology Group, Mountain View, California, cites per-user costs of $2,282 per user per year for mainframe computing, versus $6,445 per user per year for PC/LAN computing, or about 1:3. Other estimates for mainframe versus midrange systems, Unix servers, and PC/LANs, range from 1:2 up to 1:5 and above (Xenon Corporation). These experiences have led to a number of cases of data and application recentralization or server reintegration or consolidation.

Organizations which have implemented distributed computing models for major applications have encountered significant problems in scaling up the applications for growth or new business. Distributed data can result in an environment where the business information needed for planning and decisions is not available in one place and cannot readily be located, much less assembled. Major client/server implementations have resulted in "farms" of servers, with an aggregate cost exceeding a mainframe of equivalent power, and far larger support costs.

Worst of all, companies implementing departmental systems, LANS, and distributed applications are encountering serious problems in achieving high levels of availability and security for a variety of reasons, ranging from the inability to acquire (or afford) sufficient trained personnel, to the lack of mature products to actually support a highly robust and available environment.

None of this is to say that the goals that developers and users intended for new computing models, such as better user interfaces, improved productivity, and greater flexibility, were inappropriate. Unfortunately, these promises were not always met, or were met with concomitant problems in cost, management, architectural consistency and the ability to meet new challenges in a cost-effective fashion. In this climate, it is natural that users are taking a fresh look at the capabilities of mainframes, either with the intent of keeping or growing existing systems, or using specialized mainframe systems for new purposes.

Companies who have mainframes are questioning the wisdom of very expensive replacements for technologies which may not offer equivalent levels of service and robustness, and are becoming interested in ways of leveraging their existing technology to obtain the real benefits of the new advances (such as graphical user interfaces and application blending) without having to do wholesale conversions. Companies who have never had a mainframe are beginning to reject simple slogans such as "the mainframe is dead" and instead are making considered business judgements about the systems they wish to implement.

III THE MAINFRAME TODAY

The traditional strengths and properties of the mainframe, centralized computing environment are well-known:

- Mainframes run the vast majority of the world's work and the majority of the mission-critical applications in the industrialized countries
- The mainframe has unique qualities of capacity and data bandwidth which enable it to provide vastly more capability than any alternative
- The mainframe is scalable across a wider range (600x) than any other computing alternative available
- The mainframe offers unparalleled reliability, availability, and serviceability, with a robust and mature set of management and security tools unavailable elsewhere
- More applications are available for mainframes than any alternative platform, and they are in general more mature, stable, offer more function and have fewer problems

In the last few years, IBM has undertaken a massive effort to extend the hardware and software architecture of the S/390 to better address the emerging needs of today's marketplace. New technologies, such as CMOS, have enabled the repackaging of the hardware to achieve significant reductions in footprint, power, air conditioning and other environmental requirements. OS/390, the premier operating system for S/390, has been repackaged to significantly ease installation and maintenance and reduce the human resource requirements for support, at the same time as significant new software capabilities such as UNIX 95 branding and support for WindowsNT applications have been added. The results are dramatic:

- The mainframe, always a multipurpose computing system, today is an industrial-strength enterprise server for data, applications and services

- Today's mainframes are up to 95% smaller for the same capacity than those of a few years ago, with equivalent savings in environmental support. They offer up to a 75% price performance improvement over the last two years; are substantially less expensive to purchase, own and operate; and are far easier to use and support
- The mainframe of today is open and interoperable, supporting Unix applications and running Windows/NT applications in addition to MVS applications, while providing new applications the same security, management and continuous availability already implemented and expected for MVS applications

These points do not necessarily mean that a mainframe is the right choice for all purposes. They do provide a strong argument that a mainframe-based solution must be considered by organizations who wish to do a serious, value-based analysis of the alternatives and make the best possible acquisition decision, for all but small, departmental systems.

In environments where high availability, strong management capabilities, mature applications and great flexibility for the future are required, today's mainframe offers an affordable solution with unique strengths and cost containment properties.

IV THE MAINFRAME AS AN ENTRY SYSTEM

The primary focus of this report is to examine the mainframe as an entry system for a customer. By "entry system" we mean the selection of a S/390 as the initial platform on which a customer will implement a new application, business or service. This might be a traditional, vertically-integrated application or application suite, perhaps with transactional and batch processing, or it might be the use of the mainframe as a data, application or Internet server. In deciding upon a mainframe, or indeed any platform alternative, the customer must decide:

- Is this the best choice for the intended application, business or service, and why?
- How does the choice fit with the existing information technology architecture?
- What future considerations may affect this choice?

These issues are not significantly different than the choice of a computer system or technology at any point in the evolution of a business. However, the initial choice of a technology for a system is likely to condition future decisions, due to the subsequent cost of conversion to new platforms. Therefore, it is wise to consider the strengths and the weaknesses of any proposed system. If one or more of the following are applicable, then the selection of S/390 as an entry system is indicated:

- Best of breed application for the customer's purpose exists on the S/390 platform
- Cost/value equation makes the S/390 implementation immediately cost-effective compared to alternatives
- Planned business growth with accompanying growth in the use of applications will make the S/390 implementation more cost/effective even if the initial size may suggest that it starts out more expensive

- Ability to run multiple applications on the same S/390 makes both trial implementations and full rollout of new applications easier and less expensive, while allowing multiple types of workload, such as transaction processing and decision support data analysis, to operate concurrently
- Robustness and availability of S/390 are essential for mission-critical applications
- Extensive management capabilities, including data security and access control, are readily available and mature technologies
- Ability to manage and support distributed LANs and workstations can reduce support costs and allow the organization to shift resources to solving business problems
- Maturity and stability of S/390 and S/390-based applications are essential to guarantee quality and level of service in new ventures

In particular, the breadth and quality of mainframe-based applications is of particular importance to organizations and companies embarking on new business ventures or exploring new venues and locations, where availability and level of service will be the most important properties which will make the service used and exploited by the customer. Where the organization must focus on the implementation of a new line of business, or a new location or country, it must be able to rely upon the applications which it plans to implement or extend and be able to predict the requirements, costs, and results.

The first rule of attempting a new business venture, whether it be a new service, new location, or new line of business, is to reduce as far as possible the unknowns or unpredictable variables in other areas to allow maximum focus on the business challenge. In such an environment, it is clear that the best technology is the one which lets your attention be focused on the business, rather than the technology. In fact, most of the critical business of the industrialized world has been implemented on and runs on mainframe computers, and these applications are the most mature, most stable and most predictable.

V THE MAINFRAME AS A BUSINESS PROPOSITION

The S/390 has a number of characteristics, some traditional and some recently evolved, which make it very attractive when considered as a business proposition: that is, as a value versus cost equation. In some areas, such as overall availability, the mainframe is virtually unique among platform choices. These strengths will be looked at strictly from the perspective of entry system selection for the purposes of this paper.

Traditional strengths
- Scalability
- Economies of scale
- RAS (Reliability, Availability, Serviceability)
- Management (systems, network, security, business recovery, asset)
- Multipurpose computing
- Capacity (bandwidth and multiple processes)

New strengths
- Footprint
- Integration with LANs
- Vertical application integration
- Flexibility

Real Costs
- Cost per transaction
- Cost per user

a. Traditional Strengths

Scalability: Scalability is important for two reasons. First, and most obvious, scalability means that a S/390 customer can grow their system as their application requirements grow. Not only that, but the growth can be in comparatively small and inexpensive increments, rather than enormous step functions. A single S/390 can grow in capacity by nearly two orders of magnitude from the smallest uniprocessor to the largest 10-way system. In a Parallel Sysplex system, up to 32 S/390s can be combined into a computing entity with a growth potential of at least 600 times the power and capability of the smallest system. Unlike conventional clustering technology, a Parallel Sysplex system provides for real data and workload sharing not only across multiple processors in a single system, but across the entire sysplex, so major applications can have virtually unlimited growth potential.

This is particularly important because a rapidly-growing application generally represents a major and growing piece of business for a company. If such an application exceeds the capacity of the system it runs on, and no larger system is available, then generally a long, painful and expensive re-architecting process is required to split the application, redistribute it, or otherwise allow it to continue to grow. In business terms, this usually means that the business function has to be frozen for months or years just when you needed it to be able to grow rapidly, and the results can be catastrophic.

The other advantage of a very wide scalable range is the availability of applications, as an application written for S/390 will generally run on any sized S/390, meaning that one can select the best of breed application, implement it on a small system, and then grow it as needed, rather than having to face the prospect of converting from a small-company-sized application to a larger-company-sized application at some point. This scenario is more common than one would expect, since virtually all companies start small and successful companies tend to grow rapidly and outgrow their systems and tools rapidly as well, unless the systems and tools are selected with the long-term prospects of the company in mind.

Economies of Scale: Centralized or mainframe computing can always realize some economies of scale simply because it is generally cheaper to develop or produce a product or service in one location and distribute it than it is to develop or produce it in a variety of remote locations. It is much the same with central versus distributed computers, if only because one

can make better use of a single larger system running a larger workload, than of distributed individual systems where each one may be little-utilized. Sometimes the issue is not even geographical but workload location. For example, customers have implemented distributed systems where the servers were all located in a central location for management control or personnel support. In such cases, consolidating the workloads from the servers to a single larger system capable of running the multiple workloads generally can result in a significant cost decrease.

Recent reports of actual decentralization experiences have resulted in more exhaustive and careful comparisons of projected cost estimates. In one case a six-fold increase in expense was projected, leading to a decision to recentralize data and application servers. In another, a planned implementation was abandoned when an initial cost estimate of less than $100 million was discovered to really be more than $600 million in the first three years when all relevant costs were considered. In fact, a recent study shows that over three times as many MIS directors are now centralizing as are decentralizing.

Another factor in centralization is the increasing globalization of businesses and the growing recognition that their consolidated business data now only allows for improved processing efficiency, but is also a business asset in terms of its value in understanding customer behavior, sharing of information, availability, security and integrity.

Reliability, Availability and Serviceability (RAS): S/390 systems are by far the most reliable and available in terms of percent available and mean times to failure, typically offering availability levels in excess of 99.9%, at least half a percent higher than any alternative. In fact, the very concept of "availability" means different things to people from the varying types of computing. To a S/390 professional, "availability" means either mean time to failure or percent available. To a person from the Unix environment, "availability" has more to do with how fast the system can be recovered from a failure than whether or not it fails. To a LAN/server person, "availability" usually means whether a PC on the network can be found to do what somebody needs to do at the moment.

S/390 systems are not only very reliable and very available, but most types of hardware and software service can be performed while the system continues to run and the application(s) are up and available. To any company which regards the availability of its systems as critical to its success or survival, RAS must be a major deciding issue.

Management: In general, the management tools and services in the OS/390 and S/390 environment do not have parallels, let alone peers, in the Unix or LAN/server environments at this time. In many cases, management of LAN environments in areas such as software versioning and data backup is in fact done by using the tools on the mainframe connected to the corporate network. Most users in the "mainframe world" have gotten so used to having this kind of data and services management done for them that they regard it as a natural property of computing, and this is one reason for the many failures of newly-distributed

computing environments because people simply did not back up their data. Management tools for systems and performance management, data and resource access control, security and encryption, network management, business recovery services and even administrative asset management are far more developed and functional than in other environments. In some cases effective tools for these purposes exist only in OS/390.

Multipurpose computing: Mainframes used to be called "general purpose computers." An even better term might have been "multiple use computers" since S/390 computer systems are designed to run multiple dissimilar workloads. The significant factor here is "dissimilar." Most existing S/390s run a wide variety of work, including timesharing, application development, batch work, transaction processing, and typically run dozens, hundreds or thousands of simultaneous threads of these different workloads. This is relevant to the entry system customer because it is possible to put multiple different workloads on the same system, as for example in developing or implementing the first stages of a new application. Later, the application can be moved to its own system, or simply continue to participate in a somewhat larger S/390 sysplex. This flexibility contributes to the economies of scale in the mainframe environment and are virtually unique to mainframes.

Capacity (Bandwidth and Multiple Processes): One of the truly unique things about the S/390 hardware architecture is the bandwidth to data. A full-scale S/390 sysplex can achieve simultaneous I/O operations resulting in effective bandwidth in the terabytes of data per second, and can run enough processes simultaneously to actually exploit the data. While this may not be of immediate importance in considering an entry system, it rapidly becomes more relevant when considering application growth, data access, and data warehousing. As applications grow, especially transactional systems or server functions, they place an increasing load on all aspects of the computer system, and nowhere more than on the I/O subsystem. Most business applications run into capacity problems in handling the I/O before they run out of processor capacity, and typically I/O limitations are architectural rather than implementation -- that is, one can reach the limitation of the design even if larger models in the series are available. S/390 has a unique growth capacity especially in its bandwidth to data and the ability to run many multiple processes.

b. New Strengths

Size: One of the most significant achievements of the new CMOS S/390s has been the dramatic reduction in space requirements required by the new systems. An overall reduction of up to 95% in space, with corresponding reductions in power and cooling requirements, has been achieved in less than two years. Of primary significance to the entry system customer, however, is that the space, power and cooling requirements for a S/390 are in the same ballpark as small to mid-sized Unix servers and no longer require large, specially-conditioned rooms in which to work.

Network computing: The new S/390 models can attach directly to LANS and to high-speed networks such as ATM without additional, expensive controllers. This makes them much more attractive for use as enterprise servers, in addition to making the initial installation and implementation in an organization much easier and less expensive. Recent studies have shown a coming explosion both in public networking (the Internet) and in private networks and Intranets. The availability of Web browsers, and the security, network bandwidth, scalability and ability to handle large volumes of data and related applications of mainframes, make S/390 a natural choice for network computing. Customers, recognizing this, are choosing S/390 to implement their networking applications such as online reservation systems, order processing, billing, information sharing, and others requiring speed, availability, security and integrity to interface them to the net.

Of particular importance for transactional and network computing applications is the ability of the server to meet extremely high transient demands for service without compromising availability and while still providing acceptable response time to the requests. Online and transaction processing systems have long been known for high variability in transaction loads and resource demands even in well-understood environments where the systems are being used by employees with predictable schedules. With the growth of public Internetworking, companies are finding it harder to predict traffic peaks and that the peak-to-trough swings are becoming increasingly wide. The capacity, data bandwidth and capability for simultaneity inherent in S/390 make it the ideal choice for network computing servers.

Applications: The number of existing applications already built, tested and widely deployed for S/390 is tremendous. It is comparatively easy and inexpensive to install a S/390 and implement an available application to serve a new business need, without incurring the cost and support requirements of client/server implementations or experiencing the exposure to less robust and less mature applications. Additionally, applications are continually being added through the S/390 Partners in Development program. This program, previously called the S/390 Developers' Association, has developed or ported over 1,100 new and modernized applications, tools and middleware to the S/390 platform.

The new S/390 MultiPrise models will be available with vertical application suites to serve major industry areas, such as financial support, hospital management, and the like. As these existing applications continue to be enhanced, new releases are developing support for client/server environments, graphical user interfaces, and other productivity-enhancing capabilities without forcing the customer to lose the advantages of the mainframe computing platform or develop an entire application-oriented system architecture to do an initial implementation. In terms of rapid movement from requirements definition to an installed, operational system, progress can be made much faster with S/390 based applications without boxing the customer into an expensive, application-specific structure, and without limiting the customer in the future.

Flexibility: S/390 is a multi-purpose computing system able to effectively run many simultaneous dissimilar workloads, thus allowing effective use of the S/390 capacity. In addition, new capabilities which IBM has added to S/390 make it even more flexible in the ability to put new and different workloads on the mainframe, either by themselves or side-by-side with existing work. Among these capabilities are Open Edition, which allows OS/390 to run Unix applications, and Bristol Technologies' Wind/U, which allows OS/390 to run WindowsNT applications. Thus a S/390 may be used for one or more native S/390 applications and at the same time run native Windows or Unix applications. Or, since S/390 has far more capacity than either Windows/NT or Unix servers, a S/390 may be used to consolidate farms of servers or extend the size of an application without expensive rearchitecting and reimplementing.

c. Cost of Computing

It is not the purpose of this paper to examine the details of what goes into determining the actual cost of any particular implementation; interested readers are referred to the papers cited in the Bibliography. However, it has become clear that cost models which accurately reflect all of the costs associated with information technology have shown that distributed or client/server implementations are considerably more expensive than traditional mainframe or centralized implementations. Even though initial hardware and software costs may be less, the personnel costs for installing and supporting the distributed environments can be several times higher. This may be offset to some degree by higher productivity, depending upon the industry and application, but that can be very difficult to assess.

Clearly, for small implementations, a S/390 solution is likely to be more expensive than a LAN-based implementation, or even possibly a Unix server-based implementation, as far as hardware and software are concerned. However, as the size of a business and therefore the size of its systems and applications grows, there is a "crossover point" at which a centralized computing model actually costs less than a distributed model. This is largely due to the lesser costs of centralized computing as compared to distributed computing, which have been well-documented elsewhere. From that point, the costs shift rapidly in favor of the centralized model. Furthermore, this can be true regardless of whether the actual end-user interfaces are provided on dumb terminals, so-called "network computers" (actually thin-client workstations), or powerful workstations, primarily because of the personnel and support leverage which can be gained by centralized management and support.

In addition to the comparative costs associated with centralized versus distributed computing, there are other elements which make mainframe-based computing attractive. The ability to consolidate workloads on the mainframe, or implement trials and pilots on an existing system without having to invest large sums of money in new computers is clearly one of them, as is the ability to run very dissimilar types of work, such as UNIX applications. Another cost issue which tends to surface later is the need of most organizations to compile data from across their operational applications in order to build decision or management support functions. Such "data mining" capabilities require a data warehouse, or at least the ability to readily

access necessary data from across the enterprise. This is increasingly hard to do with distributed, separate applications and databases, and can in fact lead to organizations implementing entirely new applications which are nothing more than cross-system data gatherers to allow the management and decision processes to be undertaken.

In terms of overall cost, we have already cited per-user-per-year costs of 1:2 to 1:5 for mainframes versus other types of computing implementations Another perspective is that of costs per transaction, which may be considered as a specific interaction between an end user and an application or applications. International Technology Group reported three cents per transaction for mainframes over a five-year period, as opposed to forty-three cents per transaction for PC/LAN installations, or a difference of 1400%.

VI THINKING AHEAD

It is rare that any system selection is done in a vacuum, and it should be axiomatic that system decisions be made considering all issues. Obviously an organization wishes to make the best possible choice in terms of cost, capability and compatibility with existing systems. What is sometimes overlooked, however, is that each such selection affects future decisions, and may do so in very inobvious ways (the so-called "law of unintended consequences"). Examples of choices which affect future choices come readily to mind: a trivial example might be the choice of an application system running on computer "a" which requires the use of a particular network type. Once that choice is made, then future choices may require either using the same network, or a forced retrofit of a different system for the original application.

Each such choice has technical properties and business implications which thus affect not only future choices, but may affect previous decisions as well. It is very easy for an organization to find itself having adopted what might be called a "de facto architecture" in the sense that its future selections will be congruent with the "architecture" defined by its existing systems simply because of compatibility and costs reasons. The result of such a situation is to limit the alternatives which the organization can reasonably entertain. It is no secret that one property of so-called "proprietary" systems and architectures was to limit the existing customer to new alternatives within the proprietary architecture. Nor is it a secret that this property was quite intentional on the part of the vendor. Even in today's "open" computing world, it is remarkably easy to make choices which then constrain a company to the products from a single vendor in the future, or at least products which fit a single architecture.

In a sense, the legacy of decisions made by a company as it designs, buys, and builds it's computing infrastructure can be regarded as its IT history. Like all histories, the past in part predicts the future, since where you can go depends very largely on where you are today. It also depends on the amount of momentum a company has built up in a given direction in terms of ongoing projects, training, familiarity with a given environment, and all of those things which actively seem to work against a change in direction. The effect is to make an attempted change a problem in vectors -- it is much easier to go on a tangent than to reverse course.

At one time, it was considered essential for organizations to develop fairly detailed IT architectures which, by defining how an organization was going to go about providing information technology, dictated the internal standards which would be used in making decision upon hardware and software. In this model, the architecture provides structure in which to make design choices, while the standards provide limits which allow the organization to confine the number of different types of systems and vendors which must be understood, implemented and supported.

At the core of the problem for any organization is how to manage and develop a coherent structure and information technology architecture, which can be supported as effectively and efficiently as possible, while still retaining the flexibility to make new decisions according to changed business needs, or to choose a "best of breed" application which may not fit with its existing environment. Or, to use our metaphor above, we want the paradoxical properties of strong momentum for our existing environment with the possibility of radical changes in direction when needed.

S/390 mainframes provide this essential flexibility while offering the desired properties of management and cost constraint. The ability to run multiple applications and workloads on a single system by growing it, or by extending into a sysplex, allow what is clearly the least-expensive direction in terms of environmental complexity, systems management and human resources. When we consider the cost of acquiring, installing and implementing a different type of hardware and new software, integrating it with an existing environment, and supporting it, even the capitalizable cost of new hardware and software becomes a greater issue.

To add to the flexibility, as we have discussed, UNIX and Windows/NT server applications can be run on a S/390 as part of a mixed workload. This allows the implementation of a new application in a trial or pilot mode, or indeed the full production rollout of such an application, without incurring the costs of acquiring or supporting dissimilar hardware and software. Additionally the systems management and support tools and functions already in place can support the new application. And all of this can be achieved readily within the architectural and structural stance the organization has taken, without requiring vast amounts of money or resources, necessitating a redirection of organizational focus from business needs to technology.

The scalability of S/390, from its entry level systems with pre-packaged hardware, software and services, to the virtually unlimited capacity of a Parallel Sysplex system, should be perceived as a major benefit to an emerging, high growth potential business. The business can grow to world class economic proportions, with all the appropriate information technology support on the same platform if it chooses, and without disruptive, costly interruptions.

VII MAINFRAMES IN EMERGING MARKETS

The selection of a mainframe, or a mainframe-based application system or architecture, is particularly interesting in the context of emerging markets. Such markets are generally (but not always) in countries which have not been part of the developed world. Organizations in the emerging marketplace face issues which while not necessarily unique can seriously affect their ability to implement and run reliable, efficient computing environments in ways which can usually be ignored elsewhere. These issues may include

- Reliable environmentals such as clean power
- Reliable telecommunications to other sites or users
- Trained employees or contract personnel
- Availability of vendor personnel and resources
- Political stability and support
- Attention and focus from vendors

While many other issues are often identified which affect organizations in emerging markets who are acquiring and implementing computer systems, there are two general points which can be derived from these propositions and which are of particular concern.

First, an organization in this environment can take very little for granted, and must expend a considerable amount of its effort to overcome obstacles which are frequently overlooked or not well appreciated by others. Under these circumstances, it is overwhelmingly important that the organization can focus its remaining resources on its business goals and plans, rather than being consumed by the requirements of installing and implementing the technology which is supposed to further those business goals.

This argues for the selection of applications and systems which are mature, robust, well-understood and predictable, so that the organization can make credible projections for resource requirements, project plans and timescales, and can be as confident as possible. It also argues for the selection of systems which provide for flexibility and support for future choices and applications, rather than distracting the organization by forcing the architectural decisions discussed in an earlier section for each new system.

Second, the organization must be as certain as possible that a potential vendor and business partner actually understands and appreciates the unique circumstances facing the organization and is prepared to support the customer in those circumstances. Too often, vendors make an implicit assumption that once the sale is made and an initial implementation done it will be "business as usual" to support the customer. The customer must be very sure that the vendor has a credible and permanent business presence in the customer's country and marketplace and an established and relevant track record with other organizations in the country.

Some issues, such as the ones delineated below, are affected not only by the vendor, but also by the familiarity and availability of the technology being offered:

- Timely on-site support for hardware and software problems when necessary
- On-site training both before and after system implementation
- Availability of trained personnel to meet unusual requirements by the customer
- Knowledgeable support to assist the customer in dealing with unusual or country-specific issues which must be resolved
- Commitment to adequate service level agreements for the performance of not only the computer system but the vendor

In other words, the exigencies of installing, implementing and deploying computer systems in emerging markets requires an organization to be able to depend upon its vendor(s) even more than in areas of the world where computers are ubiquitous. This is partly because the organization's own resources are in part dedicated to other tasks more unique to the emerging environment, thus requiring alternatives to in-house capability and support, and then in turn because of the lack of alternatives to vendor-provided resources. In this climate, careful selection of both the technology and the vendor are extremely important for successful implementation, since an incapable vendor will limit the organization's ability to succeed with the technology, while an immature or insufficient technology cannot be successfully implemented without more resource that the organization or the vendor may be able to bring to bear. As we have seen, the S/390 mainframe offers great strengths in technology and capability, which translates to business value for the potential customer. Equally important, IBM offers great strengths in technology and capability, great resources, and worldwide presence to its customers, making the mainframe even more attractive.

REFERENCES AND BIBLIOGRAPHY

This is a partial list of papers and reports referenced in the text. Some brief notes are provided commenting upon a few of the papers.

1995 Large Systems In Review, International Data Corporation, 1996.

Architectural Evolution, The, Enabling Technologies Group, September, 1995.

Cost of Computing: Comparative Study of Mainframe and PC/LAN Installations, International Technology Group

> This paper provides an extensive look at costs, workloads, service levels, cost structures, and measurement criteria. While the comparative costs of mainframe versus PC/LAN computing models are shown to be fairly dramatic, the really interesting finding is that in general they can be very complementary, since mainframes can be used to reduce many of the costs associated with PC/LAN structures while providing unique value, especially in properly assessing business value of applications.

Cost of Scalability: Comparative Study of Mainframe and UNIX Server Installations, International Technology Group, September 196.

> This is a companion paper to "Cost of Computing" and provides many of the same types of insights with respect to mainframes and UNIX systems.

Enterprise Data Center Strategies, META Group, November, 1995.

Future Is Not What It Used to Be, The, Gartner Group Report, April 1996.

Mainframe as Enterprise Server: Leadership IS Strategies for Growth Companies, International Technology Group

> An examination of how mainframe based assets such as large-scale transactional systems and databases can be extended and exploited as part of an enterprise server strategy, supporting overall corporate needs such as decision and command-and-control systems in a natural evolution.

OpenEdition MVS: the System, the Strategy, the Significance, Carico, Bill and van der Zel, Willem, ACTS Corporation

> An analysis of the introduction of full UNIX capability into MVS and its follow-on Operating System, OS/390, with particular focus upon the strategy employed by IBM

and the significance of this capability to users. It also offers some fascinating and sobering reports of attempts to move or implement large workloads on Unix servers and some dramatic cost comparisons.

U.S. Retail Banking Transaction Growth and the Impact on IT,
International Data Corporation, 1996.

What's Really Happening to Mainframes? Chisholm, John, UNIX Review, November 1994.

COPYRIGHTS AND TRADEMARKS