

WebSphere Application Server for z/OS: the metamorphosis of the mainframe?

An Arcati Research Report
August 2005



Contents

The mainframe in today's IT environment	3
WebSphere Application Server and the mainframe metamorphosis	4
Technical issues	4
Cost management, zAAPs and the wider marketplace	8
Cost visibility	10
The question of image: mainframe skills and software choice	12
Bottom line: WAS for z/OS as a strategic choice	14



WebSphere Application Server for z/OS: the metamorphosis of the mainframe?

by Mark Lillycrop, Arcati Research

The mainframe remains the powerhouse behind the IT infrastructure of the world's largest commercial organizations, but it struggles to achieve appropriate recognition as the platform of choice for developing mission-critical Java-based applications for the 21st century. This paper re-examines the primary issues that surround the zSeries architecture, such as cost and marketplace viability, and considers why a z/OS-based WebSphere solution offers benefits for industry-strength applications that are unavailable on other platforms.

The last few years have witnessed a fundamental change in the way that business is conducted across the world, accompanied by an overwhelming growth in customer expectations. Not only have e-commerce and the Internet culture revolutionized the way that individuals and companies do business; they have also placed considerable pressure on commercial organizations to reinvent themselves. It's now only the most agile and imaginative businesses that survive – those that can achieve differentiation through innovation and by meeting and surpassing customer needs.

In most cases the key to this differentiation lies within the data stored inside enterprise applications. IT departments that are closely aligned with business developments and whose core systems and databases are tightly integrated can provide a powerful competitive weapon; but poorly integrated systems, islands of data, and application silos can turn a potentially leading-edge IT installation into an expensive and obstructive overhead.

Another key development that has taken place in recent years is the growing sophistication of IT governance. Large enterprises are now required to be more accountable in many aspects of IT management – to comply with a vast range of regulations (Sarbanes-Oxley, HIPAA, FSA etc) that are designed to ensure the integrity and security of their data assets and business processes. For companies running large IT operations, these two contrasting demands – to embrace rapid



Arcati Research Note

change on the one hand, and to control and manage change more conscientiously on the other hand – pose a unique problem that can only be satisfied by building the most manageable, scalable, and secure infrastructure possible.

The mainframe in today's IT environment

For many users, the mainframe already provides an unparalleled platform for supporting the flexibility and governance requirements of emerging enterprise applications. The zSeries has grown up at the heart of the enterprise, managing the corporate and transactional data assets on which most of today's web-facing applications rely. The IBM mainframe platform acts as a highly secure warehouse for more than 80% of business-critical data, and this situation is unlikely to change in the foreseeable future. Fronting this data are some 200 billion lines of COBOL code, the core applications that have driven business development in the financial services, retail, automotive and manufacturing sectors for the last 40 years.

In terms of scalability – arguably the most important consideration for companies building on-demand applications with unpredictable capacity requirements – the mainframe has no effective competition in the commercial environment. The current processor range scales from 26MIPS (Millions of Instructions per Second, the *de facto* measure of mainframe performance) in the z890-110 to more than 9000 MIPS in the z990 model 332, a figure that is likely to double at the high-end of the new z9-109 range. With processors coupled together in a Parallel Sysplex, it's now common for users to be managing tens of thousands of mainframe MIPS all within a single system image, processing billions of transactions per day.

What really sets the zSeries (and its predecessor S/390 systems) apart from other platforms is its flexible workload management capabilities. The whole mainframe architecture is designed to manage mixed commercial workloads – batch and interactive, transactional and database-oriented – in the most efficient manner possible, using the performance goals embedded within the z/OS WorkLoad Manager (WLM) technology to meet highly complex objectives for each application in the mix. Because of this unique workload management capability, and the mainframe's genuine 24x7 capability, the system actually achieves far higher levels of overall utilization than alternative distributed platforms – routinely 80 to 90% and even up to 100% during peak loads, compared with figures that typically fall within the 10 and 30% for server farms. This high utilization level marks the zSeries out from other hardware platforms, but it also makes the system notoriously difficult to compare with other servers in terms of total cost of ownership (as will be discussed later).



Arcati Research Note

WebSphere Application Server and the mainframe metamorphosis

The mainframe is arguably going through a metamorphosis. The performance and management capabilities which have kept the system at the heart of the data center for so many years are now being leveraged to support a new development and execution environment build around leading-edge tools. The zSeries is evolving to embrace two distinct roles – the back-end, highly secure warehouse for business-critical legacy applications; and the top-end server for new, high-performance development, much of it focusing on Java and J2EE.

One of the key technologies behind this new role is WebSphere Application Server for z/OS (WAS for z/OS) version 6. WAS is IBM's Java application server, and as such is the company's key offering for supporting J2EE-based applications across all platforms. As well as z/OS, WAS v 6.0 supports OS/400 and i5; Linux across the board; AIX, Solaris, and HP-UX; and Windows NT and 2000. Version 6 is a major release of the WebSphere product, which has been evolving for some eight years now.

WebSphere Application Server is already prevalent in many mainframe sites in one form or another, often running under Windows or Unix (the distributed and mainframe versions have taken some years to provide directly comparable function, but this has been achieved with version 6 – see the box on the next page). One of the decisions that users now need to make is whether to run their Java applications on the zSeries alongside their core legacy systems or on a distributed server, more loosely integrated with CICS, IMS and other mainframe resources. The three key factors that will influence this decision are technical strengths, cost, and image/support within the broader IT community. It's important to consider each of these topics separately.

Technical issues

There are a number of technical issues that make a z/OS-based Java application server particularly attractive. There are numerous performance, security and availability benefits to be obtained by placing WAS on the same physical server as the mainframe software that manages the organization's most business-critical data and transactions:

- *Data integration.* As stated earlier, companies that can closely integrate their information resources are far better positioned to achieve competitive edge




WebSphere Application Server for z/OS version 6: convergence of mainframe and distributed

The Java application server market is fiercely competitive: BEA, Sun and others offer mature alternatives to WAS, and the vendors in this space regularly vie with one another to provide the most functionally advanced product. The main advantages of WebSphere Application Server, at version 6, are its full compliance with J2EE 1.4, its close integration with WebSphere MQ and Enterprise Workload Manager, and (in the mainframe version) its support for key new technologies such as zAAP (Application Assist Processors, discussed later in this paper). WAS for z/OS version 6.0.1 has also seen a 15% performance enhancement over version 5.1 (up to 38% in some cases), with near linear scaling across multiple 16-way processor configurations.

Other enhancements in Version 6 include support for Java Server Faces; the WebSphere Rapid Deployment feature which allows deployment of applications into a server without continual recycling; and a new pure-Java JMS provider. There has also been a fundamental overhaul of WebSphere system documentation between versions 5 and 6.

For IBM, the greatest achievement with WAS version 6 is that it has brought the mainframe and distributed versions into line, with a common code base and simultaneous delivery. As can be seen from the chart below, the two versions of the product have taken some time to converge.



With versions 4 and 5, new features in the z/OS version would typically be based on modified distributed code, delivered some months later. From now on, the promise from the IBM architects is that the two versions will be locked together, with service updates provided in step and with identical content. The only differences in the z/OS version will be code supporting features unique to the mainframe, such as RACF- or Sysplex-specific functions. This common code base and synchronized delivery are absolutely essential as IBM repositions the mainframe as the server of choice for heavyweight Java deployment.



Arcati Research Note

than those trying to manage islands of data. By running WAS under z/OS, possibly in the same logical partition as ultra-high-performance CICS or IMS applications and IMS/DB2 data sources, it is possible to achieve the tightest levels of integration between new and existing corporate systems, with communication at channel speeds, unhampered by network bottlenecks. This has a number of advantages. First and foremost, the performance is unparalleled. Some customers have recorded a 50% improvement in response time by running WebSphere AS and DB2 together on the mainframe rather than on distributed systems. Secondly, since the application no longer needs to cross a TCP/IP network to retrieve data, there are security and reliability benefits. The data does not need to be encrypted, users require less regular authentication, and the whole end-to-end communication between apps remains unaffected by the prevalence of network failure. Sysplex-enabled mainframe systems consistently offer 99.999% availability and beyond; but if an application needs to interoperate with a network-attached distributed server, this level of availability can decline dramatically.

- *Security.* Security exposures attract more attention and negative publicity than any other aspect of business life. Company reputation and customer confidence dictate that IT departments maintain the highest levels of data security, and recent breaches have proved that the cost can be almost incalculable. Mainframe security remains virtually impenetrable: provided either through RACF or through third party products such as CA's Top Secret and ACF2, the zSeries is in a class of its own for authentication and authorization. This is further enhanced by high-performance SSL (Secure Sockets Layer) management and hardware-based cryptography, a major focus of the recent z9 announcements.

There are many security benefits to be had by running WAS under z/OS. In addition to the LDAP and Tivoli Access Manager, WebSphere Application Server for z/OS supports RACF, Top Secret and ACF2, which allows users' credentials to be carried all the way through the system, a particularly valuable feature for customer file access. WAS for z/OS is also tightly coupled with SSL and the JCE cryptographic features, which greatly reduces the transaction overhead imposed by these security technologies.

- *Management data.* Another distinct feature of the mainframe environment is its highly detailed recording of performance and accounting data, through well-established management features such as RMF and SMF. These tools gather



Arcati Research Note

information from deep inside the operating environment, and help the user not just to fine-tune the system's performance but also to keep a very detailed record of all data movements and transactional changes. Many customers have not needed to utilize such detailed system information in the past, but things are changing. IT governance and regulatory compliance are placing increasingly stringent demands on users to track the movement of data around the enterprise, and some of these requirements can only be satisfactorily met by the mainframe at present. One of the technical advantages of WebSphere Application Server for z/OS is that, like other key mainframe software components, it can generate its own SMF records; this means that it can be used to contribute vital management data for accounting and performance monitoring purposes.

- *Transaction management.* When it comes to transaction management, WAS for z/OS supports the same standard distributed services as other versions of the product – XA (X/Open Architecture) and OTS (Object Transaction Service). However, it also supports the mainframe's highly efficient and resilient Resource Recovery Services (RRS), which allows optimized two-phase commits to be performed between WebSphere, CICS, IMS, DB2, and MQ. This greatly increases the level of coordination between the centralized applications and sub-systems that are exchanging data, and puts WAS for z/OS in a league of its own for managing the integrity of Java-based transactions.
- *Workload management.* WAS for z/OS supports the z/OS WorkLoad Manager (WLM), which means that mainframe-based WebSphere apps can be prioritized according to the same performance criteria as other mainframe applications. zWLM is a far more advanced technology than workload managers running on other platforms: it controls memory usage, CPU utilization and I/O priority queuing to support detailed performance goals and to ensure that the conflicting requirements of test and production applications are catered for appropriately. This is particularly important as J2EE applications are gradually introduced to the mission-critical environment with its highly complex workload requirements.
- *Parallel Sysplex and system availability.* Lastly, WAS for z/OS is fully Sysplex-enabled. This means that the system will keep track of WebSphere transactions and affinities across the whole cluster. This has enormous resilience benefits: in the unlikely event of a problem with one server, the remaining systems will recover the WebSphere work and system logs and continue with minimal effect on response times. Moreover, WAS can take advantage of the vast scalability offered by the Parallel Sysplex which, in one recent customer case, involved



Arcati Research Note

more than 16,000 transactions a second with multiple DB2 accesses per transaction.

In view of the numerous technical distinctions of the z/OS platform, it is scarcely surprising that many customers are concerned about portability, and about the number of application changes needed to ensure adequate performance in the z/OS environment. In fact, it has long been the case that any application written to the J2EE specification will port directly to the mainframe, and this process has been simplified considerably by the common code-base in WebSphere Application Server version 6, which allows immediate function-for-function support between source and target platforms. A problem can arise when poor design in the original application is offset by adding cheap processing power in the distributed environment; this approach is neither desirable nor cost-effective on the mainframe, and can cause temporary problems during migration. However, the solution lies in applying a more rigorous approach to application design and modification across the whole IT infrastructure (one of the underlying principles of the service-oriented architecture), and this process can be managed far more tightly on the mainframe than elsewhere.

Cost management, zAAPs and the wider marketplace

However attractive the technology may seem, the mainframe exists within a fiercely competitive marketplace. While it is now well accepted that most large and mid-sized mainframe users will retain their existing systems for many years to come – as nothing can equal their performance or manageability – the zSeries has a far tougher battle on its hands among smaller users and as a platform for new applications. The Microsoft-led Mainframe Migration Alliance, for example, says that it is targeting users at the 200 MIPS level and under, for whom Windows and Unix-based alternatives are a more viable proposition.

Even if Windows and Unix appear to offer more attractive price/performance on paper, however, users need to think very carefully about migrating mature applications from the mainframe to distributed systems. Performance and reliability levels achieved through years of fine-tuning applications under z/OS and OS/390 (not to mention the flexibility and control of mainframe operations management) are notoriously difficult to re-create on other platforms, and there are numerous cases on record of users investing millions of dollars trying to 'downsize' their mainframe infrastructure unsuccessfully.



Arcati Research Note

On a simple cost per MIPS basis, though, it is not easy to convince users that the zSeries delivers the best price/performance on the market if utilized appropriately. Mainframe costs are a highly complex issue. The average IBM software portfolio may have been acquired over a considerable number of years, and be subject to a range of different pricing schemes: one-time and monthly charges; usage- and processor-group-based costs; based on Parallel Sysplex usage or geared to capacity on demand. In recent years IBM has introduced a number of new pricing options, intended either to make pricing easier to manage (such as the five-year ESSO contract, which includes a combination of one-time charges and licence fees) or to provide better value at the low end of the range where new customers need to be attracted to the zSeries (eg NALC and zELC).

Simplifying pricing for new workloads is absolutely essential to the future success of Java-based development on the mainframe; while users will tolerate a degree of cost complexity for existing systems, there is far greater pressure on IBM and the ISVs to demonstrate direct cost savings for new applications. A substantial proportion of IT users today (even those growing very quickly) may not even consider the mainframe platform when planning a major hardware acquisition, and yet the entry-level zSeries processors now offer more flexibility and granularity than ever before. With widespread availability of sub-capacity pricing and on-demand services, the mainframe offers a real value proposition for relatively small companies, particularly those that have a mixed workload to run and that value security and high availability.

For users developing new Java applications, the recently announced zAAP (zSeries Application Assist Processor) is worth close examination, as it promises to make a very significant reduction to the cost of running Java-based code on the mainframe. Fully supported by WebSphere for z/OS version 6, zAAPs are specialized processing units that are dedicated to Java apps. When configured alongside general-purpose processors (CPs) running under z/OS, they provide a means of off-loading Java-based applications that are still tightly integrated with core mainstream data and business systems (with all the functional and performance benefits mentioned earlier). The real benefit, however, is cost. At \$125,000, the zAAP is considerably cheaper than general-purpose CPs, and IBM imposes no software charges for zAAP capacity. There is also a hardware cost dividend here: customers who have bought an IFL (Integrated Facility for Linux) with a z900 and then upgraded to a z990 will be getting the additional processing capacity for the Linux workloads for free. Similarly with the zAAP Java application capacity, if they opt for a z990 zAAP with a view to upgrading to




Arcati Research Note

the z9 in the future, their investment will be fully protected. Though customers are only just beginning to take advantage of zAAPs, they do mark a significant and very visible step towards reducing the cost of Java-based workloads on the zSeries.

Cost visibility

Despite the complexity of mainframe charges, they do have one great advantage for users – they are entirely visible. The mainframe hardware and software inventory is far easier to identify and ring-fence than any other part of the IT budget. Mainframe resources are, in most cases, acquired by one group within the organization and housed in a large room with a raised floor. This picture contrasts dramatically with the distributed server world, where systems and software are frequently acquired out of departmental budgets and located across the business. Moreover software charges, which are clearly defined on the mainframe, are often bundled in with hardware and other costs on Windows and Unix servers. Keeping track of distributed IT assets is one of the greatest challenges facing large businesses today, particularly as recent changes to legislation make it increasingly important to account for all IT purchases and to keep accurate records of usage and licensing terms. A large proportion of the distributed IT spend falls outside the remit of the central IT function in many large organizations, and this poses a real problem when it comes to comparing the overall cost of ownership of different platforms.

There are other hidden costs that make a direct comparison between the mainframe





Arcati Research Note

and other platforms extremely difficult. As explained earlier, the mainframe architecture is designed to run a mixed workload at very high utilization rates – often well above 90%. z/OS management and automation tools allow system resources to be allocated to different tasks at a highly granular level, which allows IT performance and business goals to be aligned far more effectively than in the distributed environment.

These differences can have a substantial effect on the overall efficiency and bottom-line IT cost of ownership, but they are not immediately apparent when comparing simple hardware and software costs.

A further hidden cost in the distributed world is that of user support. Manpower costs take up an ever larger part of the IT budget, as salaries rise but hardware and software costs decline, yet they are rarely given due consideration during acquisition discussions. Once again, the cost of mainframe-oriented user support is highly visible, and is greatly reduced by the level of automation in this environment. At the PC server level, support is often ill-defined and provided on a departmental basis, the costs well hidden from view.

Because of the difficulty of quantifying these hidden costs, it often falls to individual companies to make their own comparisons. *The Dinosaur Myth* report, published by Arcati and last updated in 2004, attempts to make a general cost comparison between platforms by assessing the cost per user averaged out over a five-year period. The report argues that cost per user is a more meaningful way of measuring real costs than by attempting a TCO analysis. *The Dinosaur Myth* takes into account the additional expenditure incurred in the distributed environment through factors such as lower utilization rates, less visible user support, less manageable performance goals etc, and adds those into its own estimates of direct hardware and software costs. The result, as shown in Chart 1, suggests that the mainframe can have a very clear cost advantage.

The Dinosaur Myth is sometimes accused of over-simplifying the issue of cost comparison, and indeed it does. The report simply provides a cost estimate, and sets out to highlight some of the issues that are frequently forgotten when processors are compared on a cost per MIPS basis. It is essential for finance managers (especially those who do not have a background in mainframe technology) to look at the complete picture, and to appreciate where the mainframe adds real value.

Unfortunately, the architectural advantages of the zSeries are often overlooked



Arcati Research Note

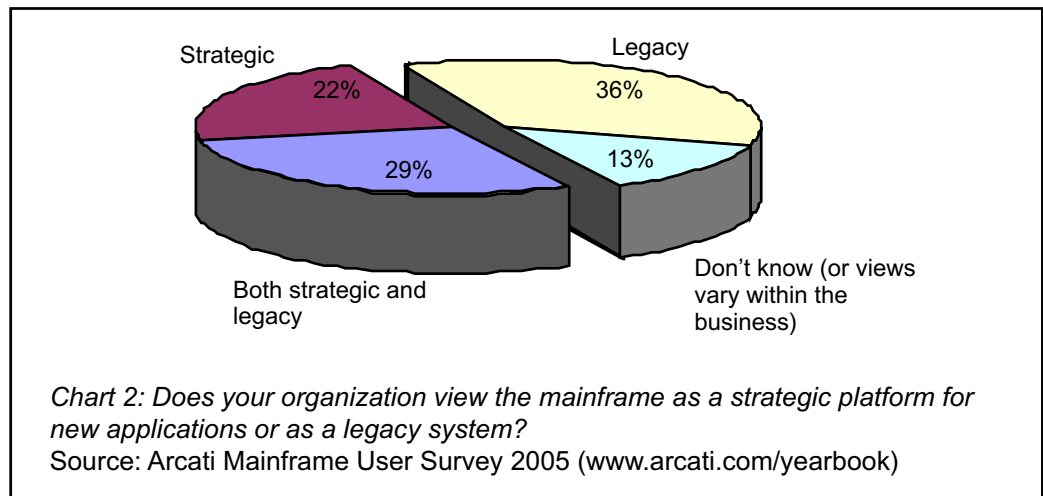
when making cost decisions about new Java-based workloads. The platform offers not just greater flexibility in setting workload priorities – a big cost consideration in its own right – but also provides dramatically better failover capabilities than distributed alternatives. This means that, instead of providing separate servers for test, production, and back-up of one application (as is quite common in the Windows environment), the same workload can often be accommodated at a fraction of the price on the mainframe, using zWLM to manage test and production images cost-effectively and taking advantage of the Sysplex environment to provide the necessary levels of redundancy.

The most expensive option for any IT department is to implement a new application on what appears to be the cheapest platform, only to discover the hidden costs when the system is up and running. However difficult it may be to make a direct TCO comparison between different servers and operating systems, it is essential to be aware of all the factors that contribute to hidden costs. In view of the high visibility of mainframe costs, and the very attractive pricing of new initiatives such as zAAPs, z/OS should now be carefully considered as an option for all WebSphere Application Server implementations.

The question of image: mainframe skills and software choice

Quite apart from cost of ownership, the mainframe does have an image problem to overcome. Mainframe capacity is growing at a very healthy rate worldwide but, as shown in Chart 2, a large proportion of users currently view the platform as a legacy overhead rather than a strategic platform for future development.

Image is a difficult thing to change. The traditional zSeries development





Arcati Research Note

environment, built around COBOL and PL/I, accounts for a huge proportion of existing commercial code, but it is perceived as increasingly irrelevant by younger IT professionals who are coming into the marketplace with J2EE, XML, and .NET skills. Managers planning a long-term IT strategy will inevitably take into the account the fact that most mainframe support staff are now in their 50s and there may be a looming shortage of younger replacements.

There have been a number of responses to this problem – IBM's *Mainframe Charter* pledges substantial training and support for mainframe users who lack suitable internal resources and, thanks to the extensive Academic Initiative and zSeries Scholars Program (<http://www.developer.ibm.com/us/en/university/scholars/products/zseries/>), a large number of universities are now offering a broader range of z/OS based courses. Recently publicized examples, such as the University of Arkansas which benefited from \$7 million of IBM investment in zSeries training resources, underline the significance of these initiatives. Furthermore, many mainframe users are now working closely with their local educational institutions to reposition large systems within the broader IT curriculum, such as the alliance between the Royal Bank of Canada and Mohawk College.

The *Computerworld* article "Dear IT graduate, just one word: Mainframes" (<http://www.computerworld.com/newsletter/0,4902,103659,00.html?nlid=PM>) highlights some of the reskilling projects that are underway. In the article analyst Mike Kahn of the Clipper Group, who has discussed the issue with a large number of Computer Science students, noted that "They were all talking about job security and getting a good job and not being laid off in three months. There is a lot of security in large systems. Mainframes is a place where you're needed." Ironically, it may be the recent downturn in the broader IT environment that encourages IT graduates to think again about the 40-year pedigree of the mainframe and adjust their areas of specialization accordingly!

Ultimately, though, the mainframe needs to be seen not just as a reliable platform for legacy applications but also as the best place to develop and run new applications. This is also essential in order to address the third key concern for mainframe users (after costs and skills): the ability to run a broad a range of leading industry software packages. Many ISVs have been reluctant to support the mainframe in the past, because there was little incentive for them to undertake a major migration effort for what appeared to be a small return on investment.

This is changing fast, however. As the zSeries evolves into a leading-edge platform



Arcati Research Note

for Java applications and Linux consolidation, it is becoming increasingly attractive to a very wide range of third-party software vendors. IBM has put in place a broad range of tools and support programs to encourage ISVs to port their products across the entire IBM server range, but focusing particularly on the zSeries. The WAS for z/OS Enablement Initiative (http://www.ibm.com/software/webservers/appserv/zos_os390/isvinit.html) is an extensive 'hand-holding' service for software vendors, designed to make the application port as smooth as possible and underline the fact that ISV applications written to the J2EE specification will run on WAS for z/OS.

In the end, it is down to customers to put pressure on their preferred ISVs to take advantage of these services and make the required port to the mainframe. If a customer has a strong price/performance incentive or other reasons for implementing new systems on the zSeries (as many have), it is essential that a reluctant ISV is not allowed to stand in the way of technical progress and cost reductions.

WebSphere Application Server is rapidly developing into a common denominator for Java-based code, allowing users to port new applications easily from one server to another, depending on changing business requirements and priorities. J2EE is a truly cross-platform specification, and is server-agnostic to an unprecedented level within the IT industry, and lack of mainstream software choice on the mainframe is rapidly disappearing.

Perhaps the most important issue for those developing and supporting mainframe applications is that they are sharing more and more common attributes with their distributed counterparts. As explained earlier, version 6 of WebSphere Application Server on the mainframe is now functionally in step with the distributed version of WAS, with the two product streams sharing a common code base. The look and feel of the products is now extremely close, to the extent that one technical team can now manage both.

Bottom line: WAS for z/OS as a strategic choice

There can be no doubt that the mainframe remains a technically attractive platform for running next-generation J2EE-based applications. In many ways WebSphere Application Server for z/OS provides the best of both worlds, combining the future-proofing of a robust industry-standard Java application server with the transactional scalability and resilience that can only be provided by a fully mainframe-enabled



Arcati Research Note

product set. For companies with a large existing investment in z/OS-based applications, there are clear and measurable benefits to be gained by running the application server on the mainframe and taking advantage of the broad range of technical synergies between the new and legacy environments.

For customers that are new to the mainframe platform, or who have concerns about the future availability of skills and applications, there is still some work to be done. However, there is a noticeable convergence underway between the technical skills required on the mainframe and those needed to support Windows and Unix systems. The WebSphere environment, in particular, now offers a common code base with consequent benefits for programmers and administrators.

As regards cost, direct comparisons between platforms remain difficult, but it is essential for companies to consider the overall cost per user of implementing new applications, and to be aware of the many factors that contribute to hidden costs in the distributed environment. IBM's zSeries pricing is becoming increasingly aggressive, and new Java-focused options such as the zAAP feature have fundamentally changed the way that users perceive the cost of developing new code.

The mainframe remains a highly cost-effective platform for running new applications, particularly for large and medium-sized companies with a mixed commercially-oriented workload and high performance requirements. WAS for z/OS is now very well positioned to allow users to leverage the mainframe's strengths for new development work.

Furthermore, for companies that need to achieve tight integration between their existing information resources and new applications, or that need a more consolidated view of their IT environment to comply with IT governance requirements, the zSeries offers some unique capabilities that deserve very close examination. As a platform for mission-critical applications, businesses will be hard pressed to match the security, reliability and performance of today's mainframe.

This paper was sponsored by IBM. The author, Mark Lillycrop, is Chief Analyst at Arcati Research.