



mainframes
in perspective
a classic going strong

Contents

1. Introduction and summary	4
2. How we got here: the road less travelled	6
2.1. Market position	6
2.2. Availability, performance and utilization	6
2.3. Software and standards	8
2.4. Cost of ownership or usage	8
2.5. The evolution of computing	8
3. How does 'legacy' mainframe fit within current IT?	10
3.1. Market position	10
3.2. Availability, performance and utilization	12
3.3. Software and standards	12
3.4. Speciality engines	13
3.5. Cost of ownership or usage	14
3.6. Environmental considerations	15
3.7. Mainframes and Cloud Computing	15
3.8. Greying of the workforce	16
4. New developments within the mainframe world	18
4.1. The IBM Mainframe Charter	18
4.2. Hybrid computing	18
4.3. Workload consolidation	18
4.4. Workload modernization	19
4.5. New Workloads	19
4.6. Mainframes and Cloud	19
4.7. Service Oriented Architecture (SOA)	20
5. The mainframe business case	22
6. Colophon	24
6.1. References	24
6.2. Acknowledgements	24
Appendix A.	
A short history of the mainframe	26
Appendix B.	
What are MIPS, MSUs and Service Units?	28

... and it keeps on going

This White Paper first appeared in 2007, to confront the then-prevalent myth that the mainframe was dead. We in Atos knew otherwise: we are constantly busy, running critical systems for prestigious customers on mainframes. That work keeps on going - and growing.

Now it is over five years later and that market, and the mainframe that serves it is increasing all the time.

So, why should we feel at all reticent about extolling the virtues of that mainframe? It is not just a case of "if it ain't broke, don't fix it": there are a lot of positive virtues that make the mainframe an extremely viable platform, now and for the future.

These days user companies are much more exposed to the outside world than they used to be: they have customers wanting to access their accounts or order goods via the internet around the clock, and hackers trying to break their way into the systems. And all that needs to be backed up by an ever-more-complex ecosystem of partners and supply chain. If their systems fail, it can be embarrassingly public.

Now more than ever, therefore, enterprises want technology that simply works - delivering 100% reliability, without fuss, around the clock - and that runs all the latest systems as well. This is what the mainframe, delivered by Atos, can do for you.

Mick Symonds, Atos

Utrecht, March 2013

1. Introduction and summary

In spite of continual rumors of their death, mainframes are alive and well, and represent a very viable and cost-effective platform for running modern IT systems. This may come as a surprise to some, especially those who moved off the platform a decade or more ago.

In their use of IT, businesses have to make a continuous balance between quality of service (QoS) and cost.

QoS:

- ▶ Improved availability of services
- ▶ Reduced time-to-market for new services
- ▶ Reduced complexity, ease of use
- ▶ Compliance to regulations.

Cost:

- ▶ Reduced purchase and running costs
- ▶ Increased resource utilization
- ▶ Reduced administrative overheads
- ▶ Contract transparency.

Some fashions seem to run in circles and come back to haunt us. For some of us, it may seem to be so with mainframes: we all used to use them years ago, but then most organizations drifted off onto other platforms as they downsized. Now those people are busy re-centralizing, and many are reminded of how things used to be on the mainframe.

But it would be a mistake to think that those people will find the mainframe platform exactly as they left it years ago. The mainframe has not stood still in that time, but has evolved alongside other platforms, while still retaining the levels of quality and reliability it always had, and the growth in mainframe sales is steady.

Mainframes, as the 'big iron' of the computing world, are relied on by many large businesses to:

- ▶ Perform thousands of transactions per second
- ▶ Support thousands of concurrent users, all accessing the same systems and resources
- ▶ Allow multiple application programs to access numerous resources, safely and securely
- ▶ Manage enormous volumes of stored data
- ▶ Transfer vast quantities of data in and out of the system
- ▶ Interact, as a truly 'open' system, with applications on other platforms.

The mainframe platform has a reputation for its extremely high level of quality but, amongst those who do not use them, for being expensive, monolithic and inflexible. Yet it can be very cost-effective compared to a proliferation of smaller systems, can be flexibly configured and re-configured. It is actually a constantly-evolving platform. Many times it is not the platform that is inflexible, but the procedures built around it over time and/or a reluctance to exploit the new mainframe capabilities.

Recent studies, undertaken by IBM themselves but endorsed by Illuminata, as well as by completely independent parties, show a very positive cost comparison for the mainframe against other platforms.

There are also concerns about it being an 'island' of computing. In the first years of its life IBM indeed invented its own de facto standards, but arguably had no option to do otherwise, and since then it has been busy supporting industry and de jure standards; these days it is extremely well-connected to other platforms. And, for the best of both worlds, the mainframe will happily run very old systems on the old standards alongside new ones built to modern standards.

There is further irony to any resistance to the use of mainframes, because everyone in IT is talking about virtualization and cloud computing: but the underlying level of Cloud, Infrastructure as a Service (IaaS), is essentially the well-established utility computing service (see the White Paper, Shaping the Cloud, Atos, November 2011). The concepts involved have been around in the mainframe world for 40 years or more. By providing virtualization, VMware is doing for Intel platforms what VM/370 did on mainframes in the 1970s. Mainframe is a utility service, and has almost always been so. Whilst claiming to be innovative, many advocates of cloud services are actually 'reinventing the mainframe' in many aspects of what they are trying to achieve.

With the introduction of 'zEnterprise', the latest generation of mainframes, it is also the first platform to deliver an out of the box, integrated hybrid computing solution.

Mainframe

LinkedIn Enterprise Systems' blog 2013: "What's the best way to explain what a mainframe is to children?"

n. 1. "It's a huge computer that never crashes but you can't play Angry Birds on it."

n. 2. A large PC peripheral

From: **The Devil's IT Dictionary, Isham Research, www.isham-research.co.uk**

A recent analysis by Gartner showed that two of the most significant obstacles to continued use of the mainframe were management perception (20%), where they cite the 'Instant experts', saying "I read the mainframe is dead in the airline magazine", and opposition from Linux support staff (18%), who don't think the mainframe is 'cool enough'.

Nevertheless, the mainframe style of computing still has a role for those whose businesses are dependent on high-quality computing environments, and what better platform to supply this need than the mainframe itself?

This paper explains this apparent conundrum and gives some outline as to how the mainframe platform can best be strategically considered.

At the European air traffic management organization Eurocontrol, located in Maastricht, The Netherlands, it's all about managing air traffic in the safest possible way. Substantial amounts of data, such as flight data, personnel deployment and training scenarios, are continuously processed in real time into useable information. A reliable and efficient IT environment is crucial to this task. To achieve such a smart data center, Eurocontrol and IBM are consolidating part of the present server environment into a virtual hybrid data center based on an IBM System zEnterprise Server with Linux as operating system. This is to result in a better view on operations, more flexibility, faster responding applications, greater capacity, higher availability of the infrastructure and considerable cost savings on supervision, licenses, floor space and energy.

Eurocontrol consolidates IT environment into virtual data centre, IBM case study, 2011

2. How we got here: the road less travelled

2.1. Market position

In 1943, Thomas Watson of IBM had famously said, "I think there is a world market for maybe five computers". But over the next two decades a market did indeed arise for a few more systems than that. In the 1960s and 70s there were a number of suppliers of large-scale computers, known colloquially as 'IBM and the seven dwarves': Burroughs, Control Data, General Electric, Honeywell, NCR, RCA and Univac.

The introduction by IBM of the System/360 in 1964 was a revolution, mainly because of the nature of the architecture. Before then, each computer produced was a unique system, made to a customer's order, with no continuity in design. And each was, essentially, a single-application system, with different systems for either commercial or scientific workloads.

The System/360 was a family of originally five machines, sharing common architecture and peripherals. Customers could upgrade to newer and larger processors without concern as to problems of compatibility.

The System/360 architecture design brief was to fulfil a number of key requirements:

- ▶ Handle very large amounts of flexibly-managed storage, with a variety of data formats
- ▶ High input and output (I/O) rates with standard interfaces
- ▶ General-purpose processing capability, with controlled supervisor functions
- ▶ Scalability, with an initial range of one to fifty.

The S/360 was the first system to use microcode (or firmware), rather than have all of its functionality hardwired. These stored micro-instructions, not accessible to normal programs other than to execute them, provide a level of flexibility between the hardware and software: if a system needs to be corrected or extended, that upgrade can take place without replacing the system hardware itself. The standard architectural interface was described in a document called the 'Principles of Operation'.

Such an architecture also meant that other suppliers (e.g. Amdahl, StorageTek, Hitachi) could build hardware, whether processors or storage, that would sit alongside that of IBM and be indistinguishable to the systems software, the applications and to end users. This development opened up the mainframe market to price competition.

Even by the 1970s, with the growth in supply options, computers were still large and expensive centralised resources, involving major capital investments and with teams of support specialists, used by large companies such as banks for bulk data processing. By that time, the competition to IBM had become 'The Bunch': Burroughs, UNIVAC, NCR, Control Data and Honeywell. This was a US-centric view, as there were other local manufacturers in Europe, such as ICL and Siemens.

The market suffered further shake-out and take-overs in the 1980s. By the early 1990s almost everyone had decided that the market for the mainframe was dying, as more and more customers 'down-sized' to other platforms.

Ironically, as a back-office system, the mainframe market was revived in the 1990s. This was partly by the 'opening up' of the platform to industry standards, and partly by the advent of e-business and the use of the internet for commercial activities. The number of back-office transactions and the size and activity level of databases grew enormously, generating a revival in the market.

These systems soon became too big and critical for a migration to any other platform to be considered, or even possible.

2.2. Availability, performance and utilisation

The factors of availability, performance and utilisation are essentially what determine what is and is not a mainframe. Mainframes are built to a level of quality. From the outset, the design of a mainframe is that it can run multiple workloads in the same system environment in the same machine. The emphasis in the mainframe world was on improving Reliability, Availability and Serviceability (RAS, as IBM called it).

Internal security and integrity was built into the original S/360 hardware architecture, with provision for a protected level of 'supervisor mode', containing system instructions for use only by the operating system.

This continued with the advent in the 1970s, of the first MVS (Multiple Virtual Storage) operating system, in which system integrity was paramount, and which completely isolated any one workload from another running on the same machine. It made very efficient use of (then scarce and expensive) real memory resources by sharing them, and doing so very securely, by the use of virtual memory, using 'dynamic address translation' (DAT) to map virtual to real memory, and hardware storage keys to completely segregate the processing of different workloads.

The security mechanism within the system, RACF (Resource Access Control Facility), was later extended to outside resources, such as storage and network, guaranteeing the workloads and their data could be secured and completely isolated.

Workloads within a mainframe environment are given the appropriate priority via a policy based workload scheduler and dispatcher. The machine can be loaded up to 100% without performance issues; a sustained peak load does not cause major performance issues.

Availability was already high but became higher in the mid 1990s with the addition of Parallel Sysplex, one of the first clustering techniques, where you can couple up to 32 mainframe LPARs (Logical Partitions) acting as one machine. The concept is similar to Unix clustering, except systems can be added or removed as needed, while applications continue to process uninterrupted. Parallel Sysplex allowed concurrent data access from any node without impacting performance or integrity.

Workloads can be cloned and distributed across the cluster (load balancing), which can be spread up to a distance of 20 km cable length. You can put one half of your servers in one location and the other half in another location, at a distance of up to 20 km in the twin data centre complex, with them still acting as one machine. This makes it a perfect clustering technique to deliver a platform for true 'continuously available' applications.

Disaster Recovery (DR) and High Availability (HA) can be implemented through GDPS (Geographically Dispersed Parallel Sysplex), which in spite of the name is not actually Parallel Sysplex, but can extend Parallel Sysplex capabilities over much greater distances (easily 100 km and more). It does depend on using the right mix of hardware and software.

With the introduction of the zEnterprise systems in 2010, mainframe made its first steps in the Hybrid Computing arena. For the first time it is possible to deploy an integrated hardware platform that brings mainframe and distributed technologies together: a system that can start to replace individual islands of computing and that can work to reduce complexity, improve security, and bring applications closer to the data that they need.

Virtualization is the creation of a logical abstraction of the physical environment. It involves the de-coupling of an interface, so that different and variable resources can be used 'beneath the surface'. It can make many appear as one, or one appear as many.

Virtualization can be applied to a myriad of resources, examples of which include:

- ▶ **Services:** de-coupling what is used from the infrastructure used to provide it
- ▶ **Applications:** using Web Services interfaces, so that multiple customers can make use of a single application process
- ▶ **Servers:** running multiple applications on the same box, whether in the same or using multiple operating systems
- ▶ **Storage:** allowing data to be stored across a range of physical storage devices
- ▶ **Networks:** both wide-area (WAN) and local (LAN): running different services over the same cables.

Virtualization can be made to work across an increasingly heterogeneous range of devices, so that devices of different types can be combined and replaced. Virtualization allows increased utilisation of these resources.

“Two roads diverged in a wood, and I— I took the one less travelled by, And that has made all the difference.”

The Road Not Taken, Robert Frost (1874-1963)

Within the zEnterprise System a new concept in IT infrastructures is being introduced: zEnterprise ensembles. A zEnterprise ensemble is a collection of highly virtualized diverse systems that can be managed as a single logical entity where diverse workloads can be deployed. Ensembles, together with the virtualization, flexibility, security, and management capabilities provided by the zEnterprise System are key to solving the problems posed by today's IT infrastructure.

Software Virtualization was introduced in 1972 with the VM operating system¹. With the advent in late 1984 of Amdahl's 'Multiple Domain Facility', the first real hardware-based virtualization was introduced.

A few years later IBM followed with a more flexible implementation: PR/SM (Processor Resource/System Manager). Both are 'type-1 Hypervisor' implementations that allow multiple logical partitions (LPARs) to share physical resources such as CPUs, memory, I/O channels and access to storage devices. The hypervisor 'kernel' was built into the system's microcode.

Logical partitions (LPARs) can be created to run completely independent workloads under separate operating systems on the same processor complex. The virtualization is not only of fractions of server capacity but also of memory and I/O channels. It is interesting to note that there are no known security breaches of the mainframe virtualization environment.

Virtualization has also been extended to other facilities, such as storage and networks. Early on, mainframes used virtual storage addresses and multiple channels to the same device to ensure flexibility of access. IBM mainframes were one of the first platforms to introduce Hierarchical Storage Management (HSM), meaning that data stored on disk could be archived off to lower-cost storage, or even tape, and recalled if accessed, all transparently to the application or user. Since the early 1990s tape robots have been configurable as Virtual Tape Libraries (VTLs): a VTL presents a storage component (a Virtual Tape Server - VTS - usually hard disk storage) acting as a tape library, connected to a real tape robot with associated library. There are many more techniques available now (storage replication, RAID, virtual or thin provisioning, de-duplication) to extent this abstraction of the logical from the physical data storage.

1. Comparing Virtualization Methods for Business - A Summary, Solitaire Interglobal, 2012

2.3. Software and standards

Mainframes retain their continuity largely because, although all of the components may change, they maintain strict backward compatibility for old software. New versions extend facilities rather than change them, and no-one has to undertake a forced migration or re-write because a new version of the hardware or software emerges.

Multiple operating systems can be run on the mainframe:

- ▶ z/OS: the flagship operating system for high-availability and high-load workloads, evolved from MVT, MVS, OS390. In z/OS you can even directly run Unix programs, because it contains a POSIX-standard Unix under its hood
- ▶ z/VM: a hypervisor, allowing multiple operating systems to be run within one complex, evolved from VM/370; z/VM is nowadays mostly used as hypervisor for 'Linux on z' systems, running hundreds of virtual Linux systems in one z/VM system environment
- ▶ z/VSE: a 'smaller' operating system, originally for lower-end systems, evolved from DOS, DOS/VSE
- ▶ Linux on z: a full Linux implementation, available as either SUSE (Novell) or Redhat distributions
- ▶ z/TPF: a real-time, large-scale and very high volume online transaction processing environment, developed as a special OS for airline reservation systems and payment card processing.

They can all be run in an LPAR, under control of a hardware hypervisor, or under z/VM.

IBM introduced the concept of Architectural Level Sets (ALS) to couple the levels of hardware and operating system support: an ALS defines the architectural features and functions that are required to physically be present on an IBM mainframe system to support a certain level of the OS. These encapsulate the fairly-long lifecycle support cycles in this environment²: typically, a new model range emerges every 2.5 years or so, but can be used for between eight to ten years.

Middleware first arose in the mainframe world as a layer of software above the operating system to facilitate building and running applications. Within the operating system you have the choice of a number of middleware environments:

- ▶ Database management systems (DB2, IMS, IDMS, Datacom, Adabas)
- ▶ Transaction processing systems (CICS, IMS, IDMS, TSO, Roscoe)
- ▶ Batch
- ▶ Third-generation languages (good old COBOL, Fortran, PL/1, Pascal, C, Assembler, JAVA)
- ▶ Fourth-generation languages (4GLs: Focus, Natural, Mark-IV, EGL)
- ▶ Web hosting (Websphere, Apache, Weblogic)
- ▶ Independent Software Vendor (ISV) off-the-shelf applications (SAP, etc.).

The vast majority of mainframe programs are still those written in Cobol, with a significant use of PL/1, Assembler, C and C++. Growing proportions are now being written in Java, especially with the growth in SOA and Web Services.

There is a large and mature set of system management tooling to automate operations, schedules, security management, tape and data management.

2.4. Cost of ownership or usage

Because of its innate scale, the mainframe has never been cheap to own, but to help compensate for that it can be very highly utilized and the various workloads can be run alongside each other with complete integrity and security. So it is very appropriate to share a mainframe between multiple customers, making the cost of usage more acceptable. Several independent studies show that when applying real TCO (Total Cost of Ownership) comparisons in many cases the mainframe easily wins over distributed systems.

2.5. The evolution of computing

Computing has evolved in a number of 'waves', each seen as a 'new paradigm' and a radical improvement on what went before:

2. The Top 10 Factors Involved in Determining the Life Cycle of Your IBM Mainframe, Gartner G00211475, March 2011

Rather than remaining stuck behind in the original model, mainframe processing has actually ridden many of these waves. Indeed, within an up-to-date mainframe environment, it is possible to implement the best aspects of each of these models of computing.

The mainframe architecture, while retaining compatibility throughout, has itself evolved over the decades:

1960s

System/360 in 1964: the original introduction.

1970s

System/370 in 1970: adding multi-processing, more instructions and virtual memory.

1980s

S/370-XA (eXtended Architecture) in 1983: bringing 31-bit addressing.

1990s

ESA/390 (Enterprise Systems Architecture) in 1990: supporting the use of multiple address spaces.

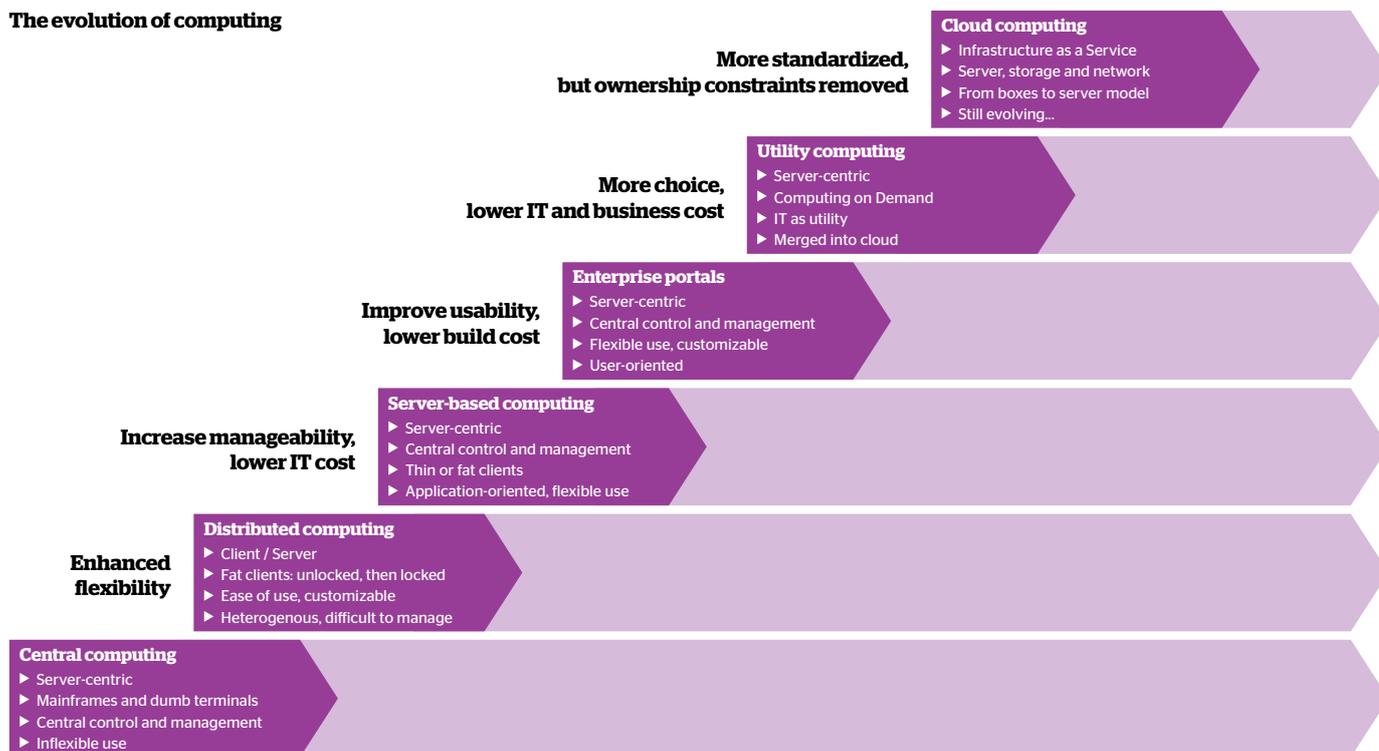
2000s

z/Architecture in 2000: bringing 64-bit addressing and dynamic channels.

2010s

zEnterprise in 2010: a step into the hybrid computing environment that includes mainframe and distributed systems; workload that spans System z and AIX, Linux and/or Windows.

The evolution of computing



3. How does 'legacy' mainframe fit within current IT?

Why is the mainframe still being built and used? They continue to provide the ability to run database, application, file server, web server, firewall, etc., all on one machine, and do so at a very high level of security and reliability.

In fact, the greatest long-term threat to the mainframe market comes not from any weakness in the platform itself, but from the continued evolution of other platforms to a level that is 'good enough' to replace it.

3.1. Market position

There used to be many platforms other than the IBM System z which were described as 'mainframes': Unisys manufactured ClearPath systems and their ES7000 was sometimes described as an 'Intel mainframe'. HP sold the ex-Tandem NonStop systems, and Groupe Bull the DPS system, and there was the DEC/VAX 9000 series. Whilst these did have some characteristics in common - such as designed-in high availability through redundancy - the IBM System z has well over 90% of market share in this field: it is 'the mainframe'.

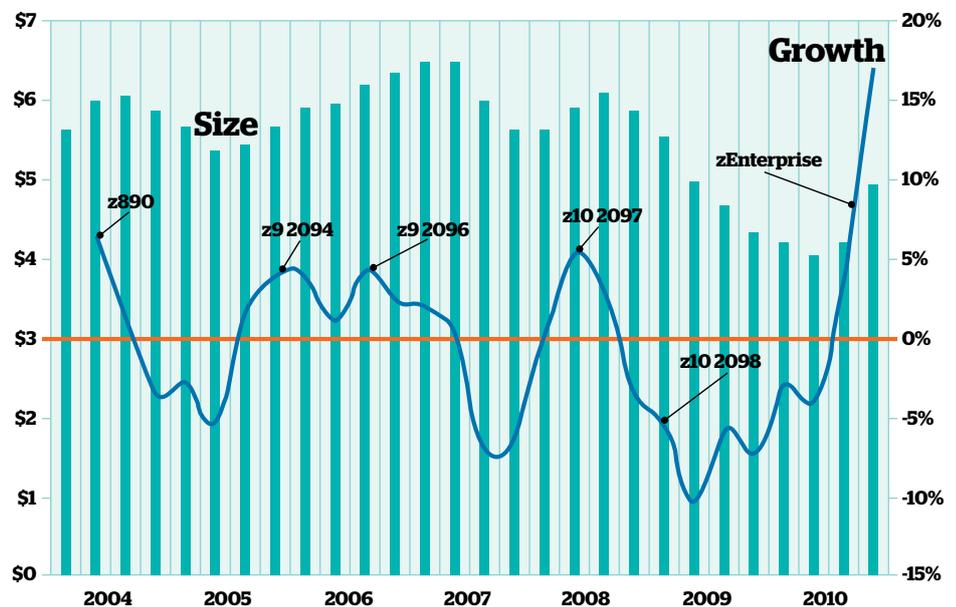
It will be seen from the accompanying graph from ITCandor that there is a continuing substantial market for mainframe systems, which has recently been experiencing significant growth. This seems to arise because those for whom the mainframe was not suitable have already moved off the platform, and those who remain are still growing their usage along with their business.

From the most recent IBM Annual Report, 'Revenues from the Systems and Technology segment totaled \$5.8 billion for the quarter. Revenues from System z mainframe server products increased 56% compared with the year-ago period'.

IBM's revenues from mainframe hardware alone are believed to be around \$4.3 Bn per year³, to which software and services can be added.

Those revenues have remained approximately level in recent years while the number of processor MIPS shipped has increased manifold, indicating that the price per MIPS has dropped significantly over that period. Indeed, that trend is shown clearly in the following graph, which also shows how it has reduced even more dramatically for special-purpose MIPS.

IBM system z revenue (\$US Billion) and revenue growth with major product introductions - 2004-2010 - rolling 4Q Analysis

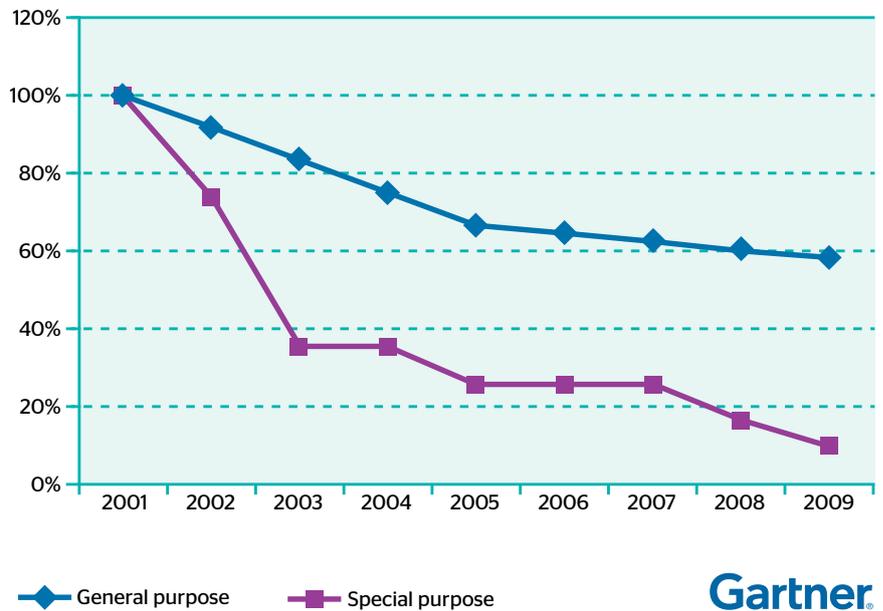


Source: ITCandor, January 2011

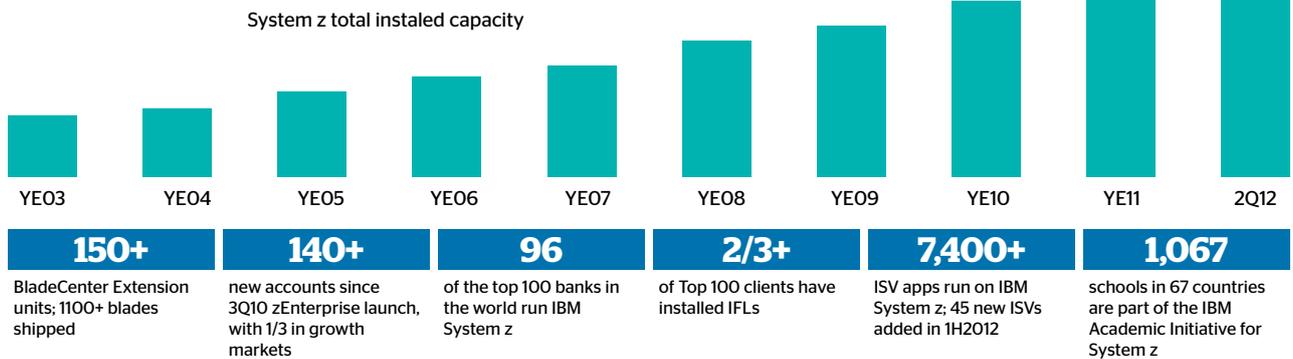
3. Start Planning for the Next Generation of IBM Mainframes, Gartner, GOO226761, January 2012

In summary, there are lots of numbers which are still going in the right direction for the mainframe, even after all these decades.

IBM cost per MIPS



The growing IBM zEnterprise System ecosystem



(source: IBM2012)

3.2. Availability, performance and utilization

Mainframes provide very high levels of availability and fault tolerance, and are designed not to have single points of failure. More than any other sizeable platform, these systems have redundant internal components and engineering: IBM System z servers execute every instruction twice, in parallel, and compare the results ('lock-stepping'). Workloads can be shifted in-flight to spare processors, without impacting the running systems.

In days when many suppliers talk in terms of 'five 9s' availability⁴, it is salutary to recognise that most mainframes already currently deliver 100% availability (although committed service hours may be less than that figure). They do so with high levels of performance (there is typically no lower 'Quality of Service (QoS)' option with mainframe services) and with a Mean Time Between Failure (MTBF) of over 20 years. Indeed, there seems to be some debate in the industry as to the MTBF for the mainframe platform: it is variously described as over 20 to over 30, or even 50 years.

An Eagle Rock Alliance survey from as long ago as 2001 identified monetary downtime costs for a range of users: if was greater than \$50k/hour for more than half and over \$1m/hour for 8%. Systems have probably become even more critical since then, with continuing adoption of real-time access and commerce via the Internet. That survey emphasises why it may be worth paying more for a high level of availability.

The mainframe platform is also inherently secure because there is no threat of viruses and limited threat of hacking: viruses on mainframes are made impossible by its architectural structure. It is not so publicly understood or available a platform as is Windows, so teenage hackers cannot hone their skills in the comfort of their bedrooms. Not only would such a threat be harder to produce, but it would be harder to implement: service providers don't tend to set up mainframes unprotected in open environments like the cloud, as they do with some other servers. And even if it could be hacked, it would have very limited impact because of the architectural design and the internally-secure nature of the platform: everything runs in its own address space, kept apart from each other by mechanisms within the hardware and microcode.

The mainframe can also deliver high levels of service due to its server architecture, consisting of separate Central Processors (CPs) and System Assist Processors (SAPs) like dedicated Input/Output Processors (IOPs), processors for cryptographic services (CryptoExpress) and with spare processors ready to take over in flight if any of them fails. The processors are deployed using dispatching algorithms which ensure that, unlike some other platforms, processors do not have to wait for I/O or other offloaded processing like encryption/decryption to complete.

The mainframe is also remarkable for its scale, in terms of the sheer volumes of data processed (processing capacity of up to 78.400 MIPS in one physical box), data stored and transferred (up to 288 GB per second), up to 3TB of real memory and, since the introduction of zEnterprise, also leading in clock speed (5.5 GHz processor in a zEC12).

In Triggers for refreshing servers, June 2005, and thus pre-mass-virtualization, Forrester stated that "The current average capacity utilisation for Wintel-based server environments is typically 8% to 15%, and 28% to 45% for Unix/RISC-based systems. For mainframes like IBM System z is 65% to 75%". Mainframes can, and do, happily run at 100% utilization for long periods.

So the mainframe is still unchallenged, in that it can be fully loaded with an enormous volume of work which it will continue to process indefinitely and uninterruptedly, with no fuss or bother.

3.3. Software and standards

In the first 20 or so years of the mainframe's life, IBM invented its own de facto standards; indeed, it is arguable that it had to do so. But since then, with the advent of distributed computing and the internet, IBM has been busy supporting industry and de jure standards.

In line with other aspects of the mainframe, though, you can continue to use the old standards alongside the new: the mainframe supports TCP/IP networking (as well as SNA and VTAM), LDAP directory services (as well as its own directory and security catalogues), Open Source, DB2 connectors, file sharing, etc.

A mainframe environment is inherently versatile: it can run a range of IT environments:

- ▶ Old software: programs written decades ago in Cobol, PL/1, Fortran or even Assembler still work on the most modern system
- ▶ Multiple application environments can co-exist: batch, online, interactive and web access
- ▶ The most modern platforms are also supported: Java, Service Oriented Architecture (SOA, using Websphere), C++, XML, etc.

⁴ Readers should always be wary of easy claims of high levels of availability. Costs increased exponentially for every '9' added, and 99.9999% availability actually equates to only 5.26 minutes downtime per year, or 6 seconds per week: see https://en.wikipedia.org/wiki/High_availability.

A mainframe can run those 25-year-old Cobol and CICS applications, alongside those which have been recently written using SOA and Java. This, although the hardware, operating system and middleware have all undergone major changes in that time; backwards compatibility has been ensured to protect investments in historic software. That, indeed, is why there is so much legacy software to worry application managers: it wasn't broken, so it didn't need fixing, sometimes for decades.

There are over 1,300 independent software vendors (ISVs) supporting the mainframe platform, a number which is actually increasing according to Gartner.

3.4. Speciality engines

Most multi-processing environments consist of processors which all access the same resources and perform equivalent functions: Symmetrical Multi-Processing or SMP. Sometimes, though, a processor can be configured to take on particular tasks, such as the maths or graphics processors in a PC.

IBM has announced a number of special-purpose processing engines which can be installed into a mainframe. They are, in order of appearance:

IFL

- ▶ Integrated Facility for Linux
- ▶ Introduced in 2000
- ▶ From 127 MIPS per processor in 2000 to approximately 1550 MIPS in 2012.

zAAP

- ▶ System z Application Assist Processor, for Java and XML
- ▶ Introduced in 2004
- ▶ Starting with z10, merged with zIIP.

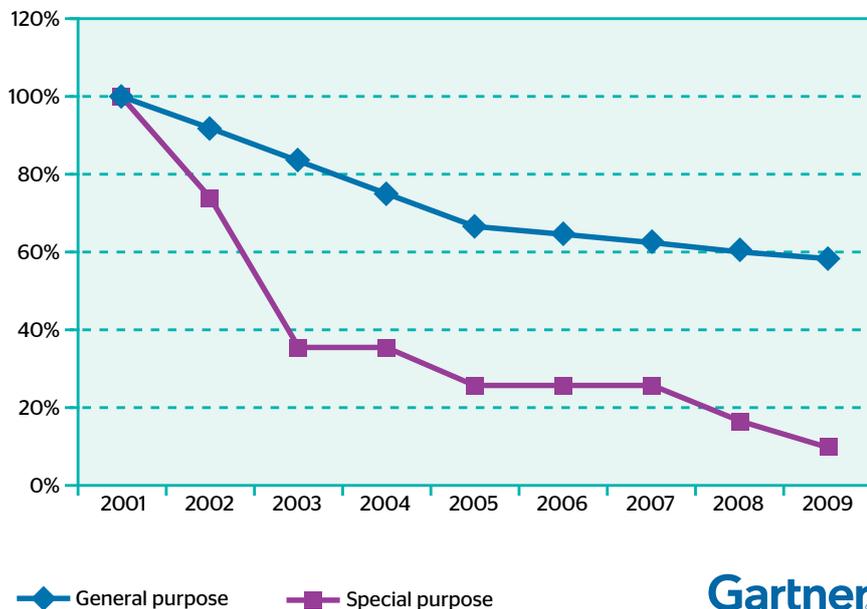
zIIP

- ▶ System z Integrated Information Processor
- ▶ Introduced in 2006
- ▶ Initially for DB2 only, later also for other database-related workloads.

The advantages of these speciality engines come partly from the fact that they 'offload' some of the processing from the general-purpose engines within the system.

The main reason however, is that they come with an increasing price advantage, both for the hardware (see graph) and also as regards charging for the software which runs within the system: this is normally charged on processor capacity basis, yet the capacities of these special engines are not counted for the normal license charges.

IBM cost per MIPS



Gartner

3.5. Cost of ownership or usage

In spite of their reputation for being expensive, there are mainframe advocates (e.g. the Arcati research group) who claim that mainframes are very cost-effective, when compared against Unix, Linux and Windows running on distributed systems, on a per-user basis.

They claim costs as follows:

Mainframe	\$4,500
Unix	\$5,400
Windows	\$8,000

But it is difficult to make an apples-with-apples comparison in this context, and differing analyses do produce different results. This favorable analysis is seen as being especially true where there are a large number of end-users wanting to use the same functionality, because the other platforms do not scale so well, and over a five-year period, because we are considering stable workloads.

Also running mainframe services at an external service provider (ESP / outsourcing) will benefit the business case in favour of mainframe by exploiting the economies of scale and scalability capabilities in a shared environment. The following table, from Gartner, illustrates that by showing the average cost per installed MIPS for a range of scales:

The cost per Installed MIPS has reduced considerably over recent years, from measured data, as tracked by Gartner.

IBM themselves have undertaken studies, endorsed by Illuminata, showing mainframes having 5-60% cost advantages over Unix, Linux and Windows alternatives. This is predicated not on a one-for-one comparison, but on the cost of running 10-50 applications on one mainframe, versus 10-20 blades or a grid of 50 systems.

Some specific aspects of this study are worthy of further highlighting and consideration:

- ▶ The mainframe requires less electricity and aircon than many 1U-racked servers running the same load. This is very significant in a modern data centre environment, where these components are much more of a constraint than space
- ▶ The balance of cost components have changed radically in the last decade:
 - people costs increased from 14 to 29%
 - hardware costs decreased from 65 to 14%.

The current distribution of mainframe cost components is shown in the following diagram, from Gartner group.

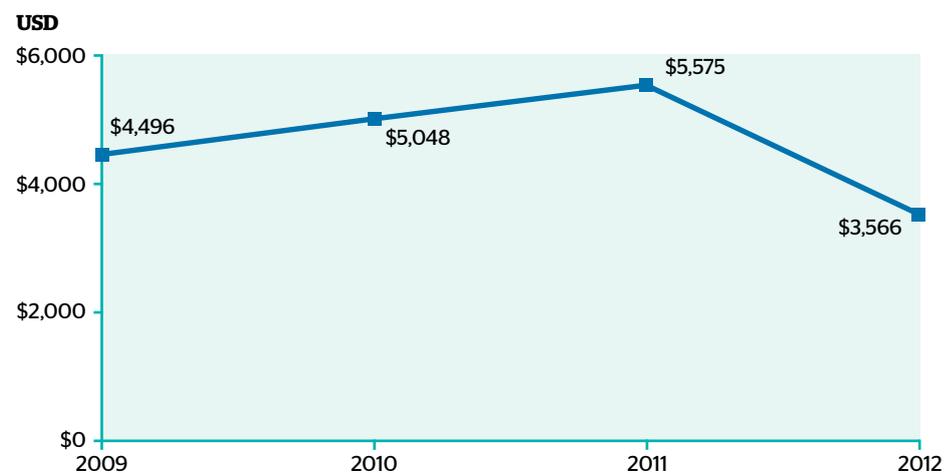
There are claims of benchmarks showing that an SAP environment with more than 1,500-2,000 users is least expensive to run on a mainframe. Whilst IBM is strengthening their alliance with SAP, initially in the DB2 area as a competitive product to Oracle, SAP is reciprocally pushing the IBM DB2 platform, including SAP on mainframe as the high-end resilient solution.

The one weakness in the cost analysis for mainframe environments is the software component. Indeed, according to Gartner this can amount to 44% of the cost of ownership. This is partly because some ISVs treat the platform as a 'cash cow'. Developments such as the provision of a Linux environment may help to address this aspect. Another possible remedy is the advent of organizations which specialize in migrating environments from known 'difficult' ISVs packages to less troublesome software.

The cost of any conversion also needs to be brought into the balance: Gartner Group has stated that, "While the overall number of IBM mainframe customers is declining (as the number of smaller customers migrating off the mainframe each year exceeds the number of first-time customers), most large enterprises have accepted that their legacy application suites (typically written in Assembler or COBOL) will need to remain on the mainframe environment. Even if the same performance could be achieved on different platforms at a lower cost, the porting or rewriting needed to achieve this is often prohibitively expensive."

Also, there are assertions that any mainframe environment of more than 1,500 MIPS (although the figure varies) is too difficult and complex to move to another platform⁵.

Annual mainframe cost per installed MIPS, 2009 - 2012



Source: Gartner IT key metrics data (December 2012)

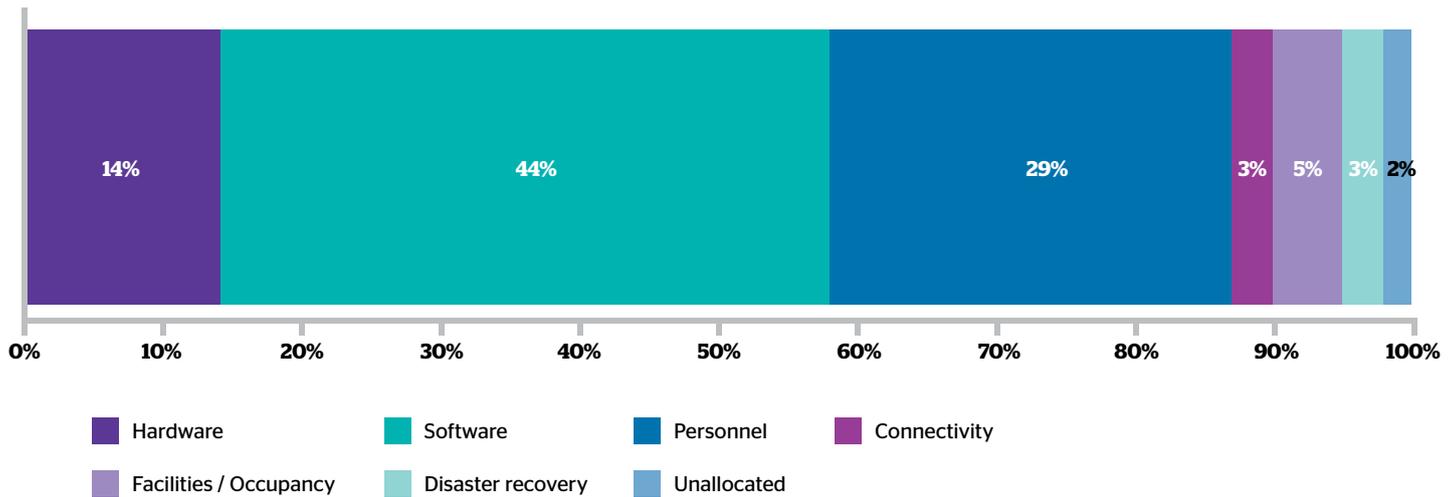
Mainframe Utilization Rate: by Environment Scale

Environment size installed MIPS	Small <2,000 Installed MIPS	Medium 2,000 - 6,000 Installed MIPS	Large >6,000 Installed MIPS
Average	61.2%	65.1%	72.5%

Source: Gartner IT key metrics data (December 2012)

5. Proof Points For Mainframe Application Migration, Forrester Research, September 2011

Distribution of mainframe costs



Source: Gartner IT key metrics data (December 2012)

3.6. Environmental considerations

In terms of power and cooling, mainframes were previously (pre 1995) a concern because of their need for water-cooling. But a highly-utilized mainframe is much more efficient than a room full of under-utilized Windows servers (as per Forrester's estimates).

Their low environmental impact is now seen as one of the strengths of the mainframe platform. With the introduction of the zEnterprise zEC12 system, IBM has re-introduced the option for water cooling for enterprise, focusing on 'Green-IT'.

3.7. Mainframes and Cloud Computing

Cloud computing is flavor of the moment in IT, but is this all quite as new as we are led to believe? Or is the mainframe a classic implementation of a cloud computing environment? 25 years ago, many companies were 'doing utility based computing' on mainframes:

- ▶ They ran a standard (albeit proprietary IBM) architecture: S/370
- ▶ Multiple business units were processing in the same environment
- ▶ The systems were often owned by suppliers, whether internal or external
- ▶ They undertook capacity planning, performance and availability management
- ▶ They used sophisticated tooling to automate, charge, secure, archive.

Then along came 'Distributed Computing', consisting of:

- ▶ Differing proprietary infrastructure architectures: every box was different
- ▶ Dedicated box(es) per application, leading to 'server sprawl'
- ▶ Dedicated storage per processor box, causing a nightmare for capacity planners and storage administrators, etc.

Later developments in distributed computing led again to re-centralizing and 'Utility Computing' initiatives.

New technologies now provide the needed aspects to allow these other environments to deliver using some of the same service configurations as the mainframe always has done: e.g. SANs that work across multiple platforms, architectural compatibility between systems, sophisticated tooling to manage the workload and provision systems to run it.

These facilities are now being extended by the widespread use of virtualization: separating the logical application environment from the physical system(s) on which it runs, and allowing multiple applications to run on the same box. Virtualization is now well-established in the Unix world and is making great headway for Windows systems. This is an approach, indeed a technology, which has been present in the mainframe world since the 1970s.

So, in many ways, current moves towards cloud computing are reinventing the mainframe.

Especially after the latest announcement of zEnterprise with hybrid capabilities, mainframe is very well positioned to play a large role in cloud computing, specifically to those shops that have already mainframe expertise in-house.

3.8. Greying of the workforce

From a completely different perspective, there are fears of a shortage of available mainframe skills, as the IT support population ages and retires, but many of these and other developments may also help address that concern:

- ▶ Mainframes are much easier to support these days: previously an army of highly-trained specialists was required⁶, but the tools to manage a mainframe platform have improved enormously, so routine functions can be done by lesser-skilled resources with much less risk and exposure:
 - z/OS has a build-in 'Health-checker'. It provides a foundation to help simplify and automate the identification of potential configuration problems before they impact system availability
 - System dumps no longer need to be analyzed in-house: rather the system will indicate to which of your software suppliers it should be sent for resolution
- ▶ The development of a Mainframe Community as part of the Mainframe Charter (see later) includes making mainframes available to colleges across the world, so that students can gain hands-on experience as part of their learning of IT skills. IBM now has more than 150 universities around the world (50% in the Americas, 30% in Europe and 20% in Asia) enrolled in its Academic Initiative program
- ▶ Linux already has an enormous amount of adherents, is very well adopted within academic circles and has its own community support. Combining this Open Source approach onto the mainframe platform could also yield enormous benefits and make skills much more available
- ▶ The advent of architectures such as SOA means that those using the platform as a component of their applications environment need much less knowledge of the underlying platform. Indeed, they may not even know on what platform a specific function is provided
- ▶ IT service supply models such as outsourcing and off-shoring mean that users can contract any remaining problem to someone else, who can benefit both from much-greater economies of scale, and from the use of a well-trained and -disciplined workforce in other geographies.

6. The key measure for staffing in this environment is installed MIPS per mainframe FTE, which is enormously improved but still varies depending on the scale of installation (source: Gartner).

4. New developments within the mainframe world

4.1. The IBM Mainframe Charter

IBM does seem to recognise that the mainframe platform has a perhaps-unfair reputation as being 'a bit of a dinosaur'. This may be unfair if the technology is judged dispassionately, but partly arises because of some genuine concerns, such as whether customers will still be able to retain the services of suitable technicians once the 'first generation' of support staff has all retired?

There are good commercial reasons why IBM should try to maintain and enhance this platform as long as it can: it is very profitable and there is very little direct competition.

IBM responded to this situation in 2003 with the creation of the 'Mainframe Charter': this was and is a set of nine statements (see box) in which it promised to continue innovation, offer value and foster a support community.

This seems a worthy intent, albeit one which is in its own interest. Mainframes do seem to engender loyalty, and there genuinely is, and has for a long time been, a sense of community within the mainframe world, somewhat redolent of that within the Open Source community. Indeed, with the use of Linux on z, those two communities do seem to fit well together.

4.2. Hybrid computing

With the 2010 introduction of the next generation mainframe (code named 'the Gryphon' - half lion / half eagle), IBM System z made its first steps into the world of Hybrid computing.

In fact the introduction was threefold:

- ▶ A new generation of mainframes - zEnterprise - with the typical characteristics of a new machine: better, bigger, faster, more of everything
- ▶ zBX: the System z Blade center eXtension, capable of running IBM Power7 (AIX and Linux) and x86 blades (specific appliances and accelerators) in tightly-coupled racks. In 2011 this was extended with Windows
- ▶ Unified Resource Manager: allows the supported zBX platforms to be virtualized into a single system for management. It also allows for the prioritization of certain workloads in the system. URM can monitor the various platforms for signs of bottlenecks or faults, and modify the overall system to recover, maintaining a specified quality of service level at a workload level.

Together with the introduction in 2012 of another new zEnterprise system (zEC12), the IBM DB2 Analytics Accelerator (IDAA) was also brought to light. It runs complex queries up to 2000x faster, while retaining single record lookup speed, and eliminates costly query tuning while offloading query processing. It is comparable to Oracle Exadata or SAP HANA technologies, but then completely within a zEnterprise complex.

4.3. Workload consolidation

Linux seems to be a very attractive option for the mainframe environment, both for systems to be converted from a previous z/OS (MVS) environment and also as a target for consolidation. Many Linux systems with variable load can be consolidated into one mainframe environment, with the added advantage of in-machine transfer speeds between them.

Linux runs 33% of the net new MIPS being shipped, according to Gartner Group, and 60% of mainframe users are using some Linux on the platform. 'Almost 225 ISVs currently deliver over 600 applications running on Linux on the IBM mainframe', according to Bloor Research.

Linux skills are readily available and can easily be ported to this platform.

The IBM Mainframe Charter

It is our intention to continue to:

Provide Innovation:

- ▶ Provide leadership in innovation to enhance the use of IBM eserver zSeries to support increasingly integrated and flexible business processes for the on demand business
- ▶ Maintain zSeries position as a benchmark for flexible, efficient, and responsive platforms for highly complex, integrated environments running a wide range of mission-critical workloads
- ▶ Improve the automatic and self-managing capabilities of the zSeries while working to simplify user processes and system administration tasks.

Offer value:

- ▶ Enhance the value proposition and lower the cost of computing of zSeries solutions in a way that is compelling, clear, and consistent
- ▶ Extend the on demand characteristics of zSeries servers highlighting its strengths as an environment for usage-based computing
- ▶ Increase the ability to account for allocation and use of zSeries resources in an on demand environment.

Foster a community:

- ▶ Support programs designed to foster vitality in zSeries community, helping to promote a strong application portfolio and world-class service
- ▶ Provide the skills and expertise to assist customers in designing, developing, and deploying on demand solutions built on a foundation whose cornerstone is zSeries
- ▶ Leverage key open standards and common structures to enhance the use of zSeries in large, heterogenous environments.

4.4. Workload modernization

There are a number of known and proven techniques that can be used when faced with old applications running on the mainframe, which can range from 'freshening them up' in situ to complete conversions to another platform, with different levels of changes in their end-user appearance.

Many of these can take place as 'soft migrations', where the existing and new environments co-exist seamlessly, and the users transfer in groups or the applications transfer module-by-module.

The general approaches include:

- ▶ Upgrade: the deployment of new software technologies on the existing platform
- ▶ Migrate: transition of existing applications to a new optimal platform
- ▶ Regenerate: generation of new code based on the existing application model and/or application source code (e.g. Cobol to Java)
- ▶ Replace: development of complete new code, and/or purchase of commercial-off-the-shelf (COTS) software.

It will be recognized that these are not exclusive approaches, given the modularity and complexity of most applications environments, and that a blend of approaches may be appropriate, depending on the circumstances.

Some examples of how modernization can be applied are:

4.4.1. GUIfying: putting a new Graphical User Interface (GUI) in front

To replace 3270 ('dumb') terminals with a modern web-based front end.

4.4.2. SOA (Service Oriented Architecture) front end

Leaving most of the existing applications code in place, if it is suitable, and 'opening up' the environment by putting an SOA interface in front of it. The application can then take its place in a more complex multi-application environment, connected in various ways (e.g. via an Enterprise Service Bus (ESB)).

4.4.3. Migrate to Java

Typically, where applications are running on a DBMS (e.g. IMS or CICS), written in COBOL.

IBM's Rational™ EGL (Enterprise Generation Language⁷) may be used as a means of conversion: it is similar in many ways to languages such as COBOL and Java, but with enhanced modelling capabilities. The resulting software can use special-purpose engines (e.g. zAAP), yielding considerable cost savings.

The code is converted (typically, 98% automatically) from languages such as COBOL to EGL, which can in turn generate Java code. Some specific subroutines (e.g. in Assembler) may need to be converted manually.

The database may be able to continue in use as-is, or be replaced. If it is replaced, there may be a need for some data conversion, e.g. from EBCDIC to ASCII.

There are multiple companies delivering similar solutions, some as a one-time service, others as a software product, e.g. MicroFocus, Bluephoenix, PKS.

4.4.4. Re-hosting

Migrating mainframe workloads to what are perceived as being lower-cost platforms⁸. Typically, these are to other dedicated platforms, one of the options for which is Linux on z; some organizations envisage doing so to Cloud environments, but these are typically Private rather than Public Clouds. Various package solutions exist to enable this approach, either by re-hosting or code transformation.

4.5 New Workloads

To counter the application migrations off the mainframe and to attract new workloads, in late 2009 IBM introduced the Solution Edition (SE) program for the zEnterprise. This delivers a zEnterprise mainframe as a bundle of hardware, software, middleware, and maintenance at a steep discount, but is workload-specific. The program addresses some specific applications and workloads, including Application Development, Cloud Computing, Enterprise Linux, and SAP.

The SE for SAP program claims to make it more affordable for companies to benefit from the strengths of System z for their SAP environment: if they are already running SAP in a distributed environment, they can bring it to the mainframe through SE for SAP. It can combine with DB2 on z/OS at a special price and (optionally) Linux on System z by taking advantage of the System z SE for Enterprise Linux.

Hybrid computing environments also can be accommodated under the SE program. For example, when implementing the SAP database server with z/OS and DB2 on zEnterprise, companies may choose to implement the SAP application server on zEnterprise with Linux on z, or on the zBX with POWER7 blades and AIX. This comes at a higher cost saving as both the zBX and Unified Resource Manager are considered SE optional products.

7. See [https://en.wikipedia.org/wiki/EGL_\(programming_language\)](https://en.wikipedia.org/wiki/EGL_(programming_language))

8. Rehosting Mainframe Workloads in the Cloud, Gartner G00218073, November 2012

9. See The mainframe and the cloud, The Bathwick Group, November 2010

4.6. Mainframes and Cloud

For Cloud environments, there are pushes to position the mainframe as a suitable vehicle for the provision of Private Clouds: those hosted within and for the sole use of an enterprise⁹. It is based on a 'scale-up', rather than the usual 'scale-out' model, but has inherent advantages of elasticity and multi-tenant security, with an increasing ability to integrate with other platforms. But mainframes have also found their way into Public Clouds, such as Amazon.com

A specific configuration, the zEC12, has been created, which can run thousands of distributed Linux systems, and can be coupled with the zBX to support Windows and AIX alongside.

4.7. Service Oriented Architecture (SOA)

Many businesses continue to rely on their mainframe-based application systems to provide critical transactional business support. But they need to find a way to 'open up' that platform, stop it being such an 'island', and allow it to take its part in the wider set of applications and IT systems within the enterprise.

Service Oriented Architecture (SOA) is an application development and processing model which is very modular, and allows application functionality to be broken into components and exposed to other applications.

SOA:

- ▶ Standards-based components reused and combined as required
- ▶ Self-describing: standard format of request and response
- ▶ Self-contained: not tightly coupled
- ▶ Location independent: in terms of network or geography.

Web Services refers to the way in which these SOA service components can be made accessible via a web (intranet or internet) interface, making them accessible from other parts of the organization or even other organizations.

SOA depends upon having building blocks of functionality: a system or systems which host particular information or perform particular tasks when invoked in a specific way. This succinctly describes many legacy mainframe environments, which is why they are particularly amenable to being 'wrapped' in an SOA environment.

Technologies to SOA-enable legacy (e.g. CICS) applications use Websphere as a front-end. Websphere is the predominant, Java-based, platform for these developments in an IBM environment.

If legacy systems needed to provide business process support that can be front-ended by an SOA environment, they can then be 'knitted together' with other systems, making up a complete environment from these modular components.

5. The mainframe business case

Rather than being treated as being 'legacy', the mainframe platform can just be seen as another platform within a range of current options, with its own performance and cost characteristics. And as one which fulfils that role while it is still running the vast number of business-critical applications it brings as its heritage.

That said, the mainframe has unique characteristics. In most cases, it effectively delivers 100% availability: there is no 'low Quality of Service' option for mainframe services, which have a MTBF of over 20 years.

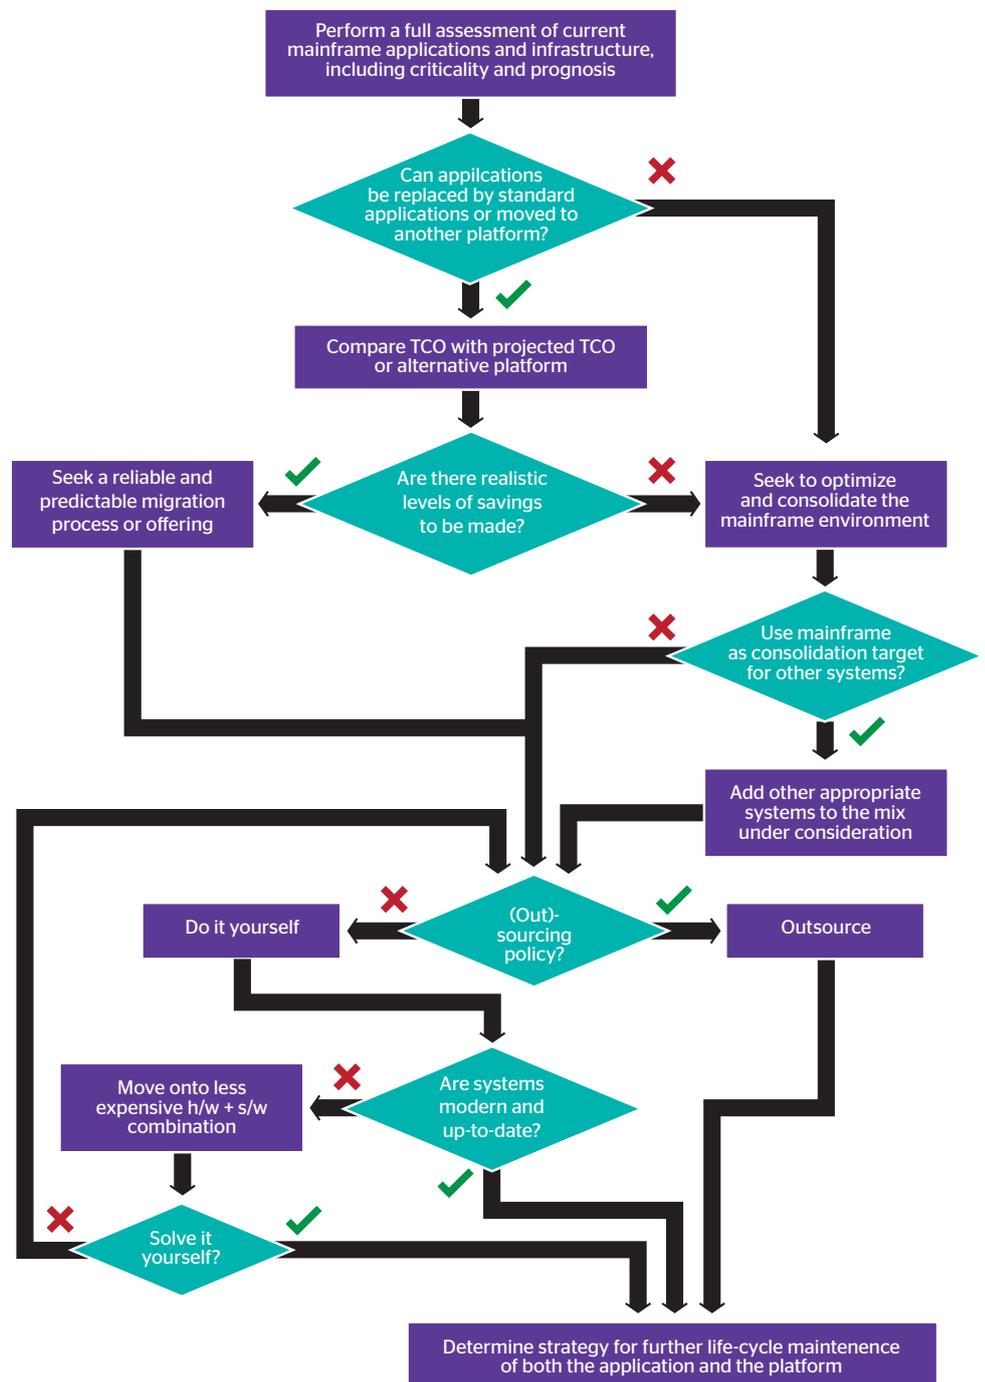
To run mainframes efficiently, however, you do need a certain critical mass. It is debatable what that level is, but the breakpoint seems to be in the scale of 7,000 MIPS. At that point you can make efficient use of the Parallel Sysplex environment and the scaling licence fees. So, many users are now opting to have their mainframes run by a service provider, who has the critical mass of both systems and skills in order to maintain and develop the environment.

Conclusions:

- ▶ If you already have a mainframe environment, you may very well find that the best approach is to keep it, and even consolidate other systems onto it
- ▶ Possibly you could look to get it managed for you, even if you retain other platforms yourself: you don't have to own a mainframe to be able to use one
- ▶ Simplification is key to the management of these environments: create a single, coherent environment
- ▶ These systems could be the target for large-scale consolidation: you can run hundreds of Unix or Linux images on one mainframe.

Users of mainframe environment can also best position for the future by reducing the barriers between mainframe and non-mainframe platforms. This includes using technologies such as Linux and SOA, described elsewhere, and standard operating processes, such as ITIL. The less a mainframe sits in a stand-alone 'silo', the more it can be integrated to take its part in future IT services, and the less risk there is of having difficulties supporting it in the future.

As has now been explained at some length, the 'slumbering dinosaur' is awake, and is ready and able to maintain its rightful place in the portfolio of modern IT solutions.



6. Colophon

6.1. References

- ▶ Architecture of the IBM System/360, Amdahl et al, IBM Journal of R&D, 1964
- ▶ IT Key Metrics Data 2013: Key Infrastructure Measures: Mainframe Analysis: Multiyear, Gartner Benchmark Analytics, G00246747, December 2012
- ▶ Progress Report for IBM's zBX, Gartner, G00231027, March 2012
- ▶ Rehosting Mainframe Workloads in the Cloud, Gartner, G00218073, November 2012
- ▶ Proof Points For Mainframe Application Migration, Forrester Research, September 2011
- ▶ Server Sprawl And Cost Reduction Bring New Customers To Mainframe/Linux Platform, Forrester Research, July 2011
- ▶ Comparing Virtualization Methods for Business, Solitaire Interglobal, 2012
- ▶ Triggers for refreshing servers, Forrester Research, June 2005
- ▶ The Arcati Mainframe Yearbook 2013, Arcati Research, 2013
- ▶ Shaping the Cloud, Atos, November 2011
- ▶ Blog: Dancing Dinosaur (www.dancingdinosaur.wordpress.com).

6.2. Acknowledgements

The editorial team for this white paper was led by Mick Symonds, with contributions from Herbert Voskuil, Paul Bles and other Atos staff.

For more information, please contact your local Atos representative, email to mainframe.services@atos.net or visit www.atos.net/mainframe

Appendix A. A short history of the mainframe

A short history of the mainframe	
7 April 1964	IBM announces the System/360: a family of five increasingly powerful computers that run the same operating system and can use the same 44 peripheral devices.
1968	CICS (Customer Information Control System) is introduced. It allows workplace personnel to enter, update and retrieve data online.
1968	S/360 model 85 introduces high speed cache memory, making high-priority data available 12 times faster than ever before - and laying the foundation for the same cache memory found across much of today's computing technology.
1970	The Plug Compatible Manufacturer (PCM) market arises: by 1970, over 80 companies have over 200 products compatible with the S/360 architecture and its peripheral devices.
1970	The S/370 as successor of the S/360 is introduced. The S/370 Model 145 is the first computer with a fully integrated monolithic memory (all resistors, capacitors and diodes on a single slice of silicon). Before this semiconductor technology, magnetic core technology was used.
1971	Amdahl comes with a clone of the S/370. Gene Amdahl was chief engineer for the IBM S/360 program, he starts the Amdahl cooperation. The Amdahl machines are air-cooled whereas the IBM machines are water-cooled.
1972	IBM announces VM virtualization, with the VM/370 operating system.
1972	SAP develops a revolutionary ERP-system for S/370. For the first time, companies can place orders and track inventory in real time.
1976	SAS software creates a new competitive edge: business intelligence. With SAS you can get intelligence out of raw data.
1981	The successor of the S/370 is the 3081, which uses less energy and space than its predecessors. With it comes eXtended Architecture (XA), providing dynamic channels and 31-bit (2 GB) addressing.
1983	MVS/XA shipped: an operating system to exploit the new architecture.
1984	Amdahl introduces MDF (Multiple Domain Feature) the first hardware partitioning tool on the Amdahl 470 mainframes.
1988	IBM introduces PR/SM, the answer from IBM to Amdahl's MDF. Now hardware partitioning with Logical Partitions (LPAR's) are possible on IBM mainframes.
1988	DB2, IBM's relational database management system is introduced.
1988	MVS/ESA and VM/XA operating systems were introduced, relieving the memory constraints which limited the size of applications by using new architectural constructs.
1994	The Parallel Sysplex is announced, which makes it possible to share 32 mainframes in a single cluster.
1994	Unix on the mainframe was announced. OpenEdition, later renamed to Unix System Services is an integral part of the operating system OS/390 (In 2000 the OS is z/OS)
1995	CMOS based processors are introduced into the mainframe environment, setting the new roadmap for modern mainframe technology. CMOS machines required 1/20th of the energy and 1/10th of the floorspace compared to the older machines.
1999	Capacity Upgrade on Demand (a utility-like commercial construct) provides extra capacity which can be turned on as dictated by business needs.
1999	Linux appears on the mainframe, combining the flexibility of open source computing with the legendary scalability and reliability of the mainframe.

2000	IBM introduces the eServer family of servers where the various sub-brands were rebranded all at the same time. IBM S/390 became IBM eServer zSeries (z for Zero downtime). At the same time introduced the 64-bit version of the operating system z/OS.
2003	The Mainframe Charter is introduced by IBM. It articulates IBM's commitment to deliver innovative solutions, expand zSeries value and foster a zSeries community in order to meet customers' on demand business requirements.
2000	The IFL (specialized Linux processor) is introduced on the zSeries.
2004	The zAAP (Application Assist Processor) is introduced on the zSeries dedicated to offload specific Java and XML workloads.
2005	IBM zSeries became the IBM Systemz, again as an IBM rebranding.
2006	The zIIP (System z Integrated Information Processor) is announced for the System z servers. Initially introduced to relieve the general processors (CPs) of specific DB2 processing loads, but currently used to offload other z/OS workloads as well.
2010	Introduction of zEnterprise, zBX and Unified Resource Manager (URM): the start of hybrid computing: z/OS, Linux, IBM Power (AIX) and appliances.
2011	Extended hybrid computing with Windows capabilities.
2012	IBM Mainframe = System z = zEnterprise family of servers is introduced delivering granularity in the size of servers. With the introduction of the zEC12 in 2012, ranging from 26 to 78,000 MIPS on one machine and, at 5.5 GHz, industry-leading on processor speed for high-end servers.

Appendix B. What are MIPS, MSUs and Service Units?

Mainframes performance has traditionally been described in terms of millions of instructions per second (MIPS), as a simple indication of speed. This is not as simple as it sounds, because the mainframe uses a complex instruction set, so that the contents and effect of one instruction can vary enormously.

The difficulty is compounded by factors such as virtualization and multi-processing, and the need to balance I/O processing with execution of instructions.

The term MIPS has often been parodied as 'Meaningless Indicator of Processing Speed', but no simple alternative has emerged. Indeed, the emerging market for general-purpose utility and cloud computing suffers from the lack of a simple unit of measure equivalent to the MIP.

Within the system, Service Units are used for workload despatch and balancing. There are approximately 50 SU's per MIPS.

Separately, the concept of Millions of Service Units (MSUs) arose for aspects such as software licensing. Confusingly, these are not millions of the Service Units measured above: there are between five and eight and a half MIPS per MSU; the figure varies because IBM uses them to encourage migration to latest-generation mainframes.

IBM have tried to facilitate workload estimation and capacity planning by publishing Large System Performance Reference (LSPR) tables, but these are far too complex to allow simple comparisons.

There is an independent analyst, Cheryl Watson, who has made a niche business since 1965 out of publishing charts¹⁰ of the various systems performance for different workloads and operating systems levels, along with Rules of Thumb as to how they should be deployed.

10. See System z CPU Chart, www.watsonwalker.com

About Atos

Atos is an international information technology services company with annual 2011 pro forma revenue of EUR 8.5 billion and 74,000 employees in 48 countries. Serving a global client base, it delivers hi-tech transactional services, consulting and technology services, systems integration and managed services. With its deep technology expertise and industry knowledge, it works with clients across the following market sectors: Manufacturing, Retail, Services; Public, Health & Transports; Financial Services; Telecoms, Media & Technology; Energy & Utilities.

Atos is focused on business technology that powers progress and helps organizations to create their firm of the future. It is the Worldwide Information Technology Partner for the Olympic and Paralympic Games and is quoted on the Paris Eurolist Market. Atos operates under the brands Atos, Atos Consulting & Technology Services, Atos Worldline and Atos Worldgrid. For more information, visit: atos.net

For more information, contact: mainframe.services@atos.net

atos.net

Atos, the Atos logo, Atos Consulting, Atos Worldline, Atos Sphere, Atos Cloud, Atos Healthcare (in the UK) and Atos Worldgrid are registered trademarks of Atos SA. March 2013 © 2013 Atos.