IBM System z total cost of ownership

What it means to you and your business

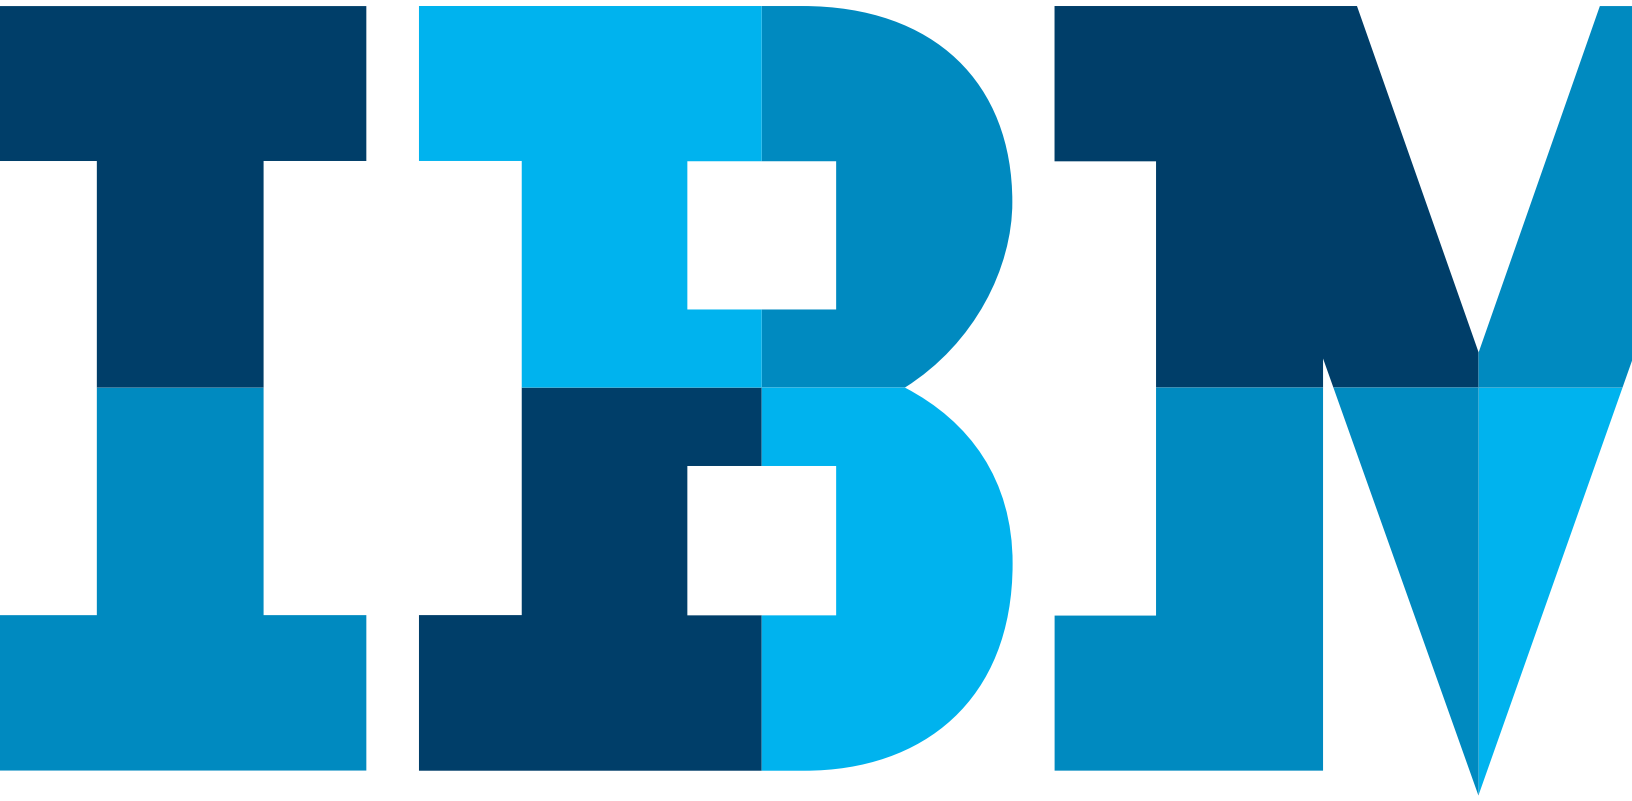

IBM

Contents

Overview

If you are an IT professional or leader who wants to explain how IBM® System z® helps reduce total cost of ownership (TCO), this paper is for you. Organizations are constantly making decisions about how to most effectively adjust the IT infrastructure to support changing business requirements and new workloads. These discussions are difficult if decision makers don't understand the unique TCO capabilities of DB2® for z/OS®, IMS™ and System z. This paper can help you as you make these decisions, with clear explanations and examples about System z and some key TCO-related questions to consider.

This paper can help you effectively explain many unique TCO benefits of System z, including:

- Industry-leading hardware, with today's fast microprocessors, the most I/O channels, and hardware-assisted clustering to support scale-out technology
- Industry-leading software, including the operating system and middleware that are trusted by some of the most demanding workloads and businesses on the planet
- An industry-leading platform featuring high levels of efficiency, extensive automation, a minimal administrative footprint, strong analysis and support capabilities, and minimal environmental impact—including less electricity, networking and floor space per unit of work

Much of the information in this paper is based on the experiences of the IBM Worldwide Eagle TCO Engagement Services team (also known as the Eagle TCO team). This team has worked with many clients around the world to help them understand total cost of ownership issues in their IT environments.

What does "total cost of ownership" mean, and why do I care?

"Total cost of ownership (TCO) is a financial estimate whose purpose is to help consumers and enterprise managers determine direct and indirect costs of a product or system."

— from Wikipedia

In an IT context, TCO includes the direct cost of acquisition of hardware and software, as well as indirect factors including the cost of maintenance, management, staffing, facilities, availability and security. If you look only at acquisition cost and not at the other factors, you are less likely to make good decisions. These suboptimal decisions can end up increasing the costs to your enterprise.

Another way to think about TCO is to think about a car purchase. When you buy a car, the purchase or lease price of the car is a big factor in your decision. However, you probably also think about gas mileage, safety, reliability, availability of replacement parts, and costs and ease of maintenance and repair. Therefore, when you decide what kind of car to buy or lease, you probably do not base your decision solely on price. The same thought process can also be applied when deciding what environment to use for the IT workloads in your enterprise.

The constant pressure on IT cost is relentless. A 2012 survey of 100 mainframe enterprises, sponsored by Arcati Limited[1], indicated that 82 percent of respondents "thought mainframes are too expensive (or appear to be)" and that these costs have a negative effect on mainframe acceptance. As a decision maker or stakeholder in an IT organization that manages large databases and high transaction volumes, you feel pressure to reduce the costs of doing business. For many clients, the mainframe budget shows up as one of the biggest items. The intense pressure to reduce costs can readily lead to the false assumption that moving workload off the mainframe will save money.

In the early days of IT, the only occupant of the data center was the mainframe. Today, many mainframe budgets include items that have no actual connection to the mainframe, including distributed infrastructure, workforce, enterprise network, electricity and cooling costs. Corporate jets have even found their way into the mainframe budget. Because of this situation, an important first step when looking at cost is to understand exactly what costs are covered by the mainframe budget and to apportion data center costs and other costs appropriately.

What makes System z so special?

If you work in a System z environment, you probably know that today's mainframe is more than just a successful legacy. Throughout recent decades, technical innovation is a key reason why System z continues to be recognized and depended on to deliver the highest qualities of service for availability, reliability, security and scalability. Although some of these innovations have been partially emulated in other solutions, the IBM System z technical evolution continues today, and will continue in the future. System z technology remains a key cornerstone for demanding enterprises around the world.

The following unique characteristics of System z account for real cost savings in many enterprises.

Centralized computing model

System z is architected and designed to provide the highest levels of availability, reliability, security and scalability for applications and data. This end-to-end design spans the operating system, middleware and hardware, including storage. System z is built to support a centralized computing model, which, from an economy-of-scale standpoint, is generally more efficient than distributed, modular models (similar to traditional UNIX infrastructures). Modular models require relevant data to be copied and moved, usually multiple times, so that it can be available for each specific business application. Capacity growth is contained through the proliferation of additional cores, which in turn drives increases in software and hardware costs — as well as labor costs associated with managing, securing and maintaining the distributed environment. The proliferation of cores also adds data and conversion legacy, as well as network overhead and costs to connect the modules.

System z and its system-level centralization offer simplicity because all components reside in one place. Therefore, a System z solution can significantly reduce the proliferation of cores, reduce mistakes, and enable a centralized management, maintenance, security and regulatory infrastructure. System z solutions can also provide performance and cost advantages by removing data latency. Removing data latency can lead to a reduction in conversion costs and networking overhead, while driving high levels of system and business efficiency.

Industry-leading availability

Since the introduction of Parallel Sysplex® and data sharing technologies, System z has been a recognized leader in availability by providing an end-to-end, almost continuously available solution. It is built on a proven stack of robust, fault-tolerant components that are designed to eliminate single points of failure. Hardware and middleware components in a System z solution are architected together to maximize application and system-level availability.

Why does industry-leading availability matter? Businesses today run 24x7, with round-the-clock online shopping and mobile banking on demand. If your system is down, your business may also be brought down, leading to lost orders and revenue, late payments and lost customers. The probability of costly outages is a key component of any TCO discussion because when applications are down or data is not available, these outages are likely to have a negative impact on potential business resiliency and customer satisfaction.

According to a 2013 EMC report[2], scheduled outages (such as maintenance, migrations, backups, batch jobs and upgrades) can account for up to 85 percent of all outages. The Parallel Sysplex technology in System z allows you to manage planned outages, such as rolling upgrades and online migrations, without causing application outages. System z continues to provide some of the highest levels of continuous availability and resiliency, with some customers reporting an average mean time to failure measured in years or decades.

Industry-leading multi-site disaster recovery capabilities

Availability during planned outages is crucial, but availability during unplanned outages is also important. Because enterprises need to be prepared to recover from disasters, disaster recovery is another key factor in any TCO assessment.

With Parallel Sysplex technology, unplanned outages in a System z environment can be minimized. System z is designed to ease disaster recovery with cost-effective solutions that are built into the hardware, middleware and software. These solutions include robust disaster recovery support in the System z database management systems, IMS and DB2 for z/OS. In contrast, disaster recovery in a distributed environment is decentralized. Each application that needs disaster recovery support must be coded with the required complexity to properly support the failover. In addition, initial and periodic testing of the disaster preparedness of all distributed applications across an entire data center is a difficult and large task, in both time and effort for the IT staff.

Investment protection

System z continues to evolve to support and smoothly integrate with the latest technologies. New versions of the hardware, middleware and software include multiple, cost-saving enhancements. For example, the latest versions of DB2 for z/OS and IMS contain enhancements that increase performance and simplify or automate administration or programming tasks. These improvements can reduce labor costs and increase throughput.

For enterprises with core business applications that run on System z, this continuous evolution means that IT teams can take advantage of new and emerging capabilities in the industry. In other words, benefiting from advances in technology does not mean that you need to discard what you have already invested in and start over. Some of the latest technology that is already integrated with DB2, IMS and System z includes Java, XML, JSON and big data. IT organizations with investments in System z technology can move to new technology areas to meet business needs, while still getting the most out of their current investments.

Dynamic system-wide workload management

The system-wide workload management technology of System z is unique and among the most powerful in the industry. This technology is designed to:

- Run both high-priority and low-priority workloads together, within the same operating system and database instance, or across database instances (even across DB2 for z/OS and IMS instances)
- Help meet response-time service-level requirements
- Reduce costs through very high utilization rates

With System z workload management capabilities, an enterprise can prioritize its system-resource usage based on its own unique and dynamic business needs and policies, across various workloads at a process level. This prioritization is extended across the entire system, including storage, with a single, system-wide view and management system. The workload management software also handles priority queuing, which can prevent high-priority requests from completely dominating lower-priority requests. As a result, you can establish work priorities and policies at the outset. The workload management software then selects what work is most important at a given time. Static configuration is not required to trigger processes.

This workload management capability is interesting because it enables you to run your highest-priority workloads—meeting or exceeding service-level obligations—and to assign less important work to any unused CPU. As a result, significantly higher system utilization can be achieved. According to a 2013 article in Enterprise Executive[3], "zEnterprise processors are architected for maximizing throughput and system utilization when consolidating multiple workloads on a server complex. Mainframes can consistently handle utilization levels of 80 to 100 percent without freezing or failing." These kinds of utilization rates can translate to a better ROI for IT managers. You pay only for what you use, not for idle capacity.

Shared-everything storage architecture

Mainframe systems offer the benefit of supporting a shared-everything storage architecture. Processors, cache, memory, I/O and storage can all be shared by many different workloads that run in mainframe systems. System z data management and storage management practices and processes help organizations efficiently and effectively manage a centralized, shared environment that supports databases and applications that run the enterprise. As a result, multiple mainframe applications can share a single instance of a database, which is not necessarily the case in distributed environments.

Ultimate scale-up and scale-out capabilities

System z technology provides industry-leading scalability and flexibility for your environment, supporting both scale-up and scale-out models. With some systems that are CPU rich but I/O poor, getting workloads to scale uniformly is difficult. In contrast, System z provides large memory and extensive CPU and I/O bandwidth. The machine, therefore, is efficient and effective across a wide variety of CPU-intensive and I/O-intensive workloads. A special hardware-assist clustering capability, delivered in the coupling facility, is exploited by z/OS and IBM middleware to facilitate very efficient scale-out capabilities. These integrated scalability solutions provide near-linear scalability to handle the most demanding workloads.

What goes into a TCO analysis?

Building on the earlier car analogy, you need to look at more than just direct cost when investing in IT technology. Here are some of the key factors to consider:

- Components of your environment, including hardware, software, people, network, storage and facilities
- Time-related factors, such as resources for migrating to new hardware or software, the amount and timing of organic business growth, planned business changes that affect capacity needs, and periodic changes, such as hardware refreshes
- Indirect factors that dramatically affect various qualities of service for your enterprise, including availability, security and scalability

A valid TCO analysis also includes costs for any nonproduction systems that are associated with each production system. Development, test, quality assurance and disaster recovery systems that support each production system all include hardware, software, labor, storage and other components listed above. Therefore, costs for each nonproduction system need to be identified, weighted for the needs of the business, and added to the TCO assessment, as shown in the following figure.

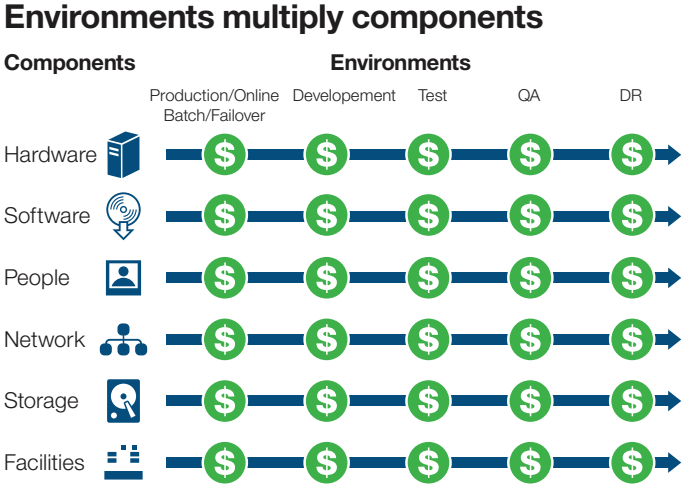

Figure 1: The IT environment can significantly increase the number of required components

System z can often facilitate multiple environments on a single logical partition (LPAR), whereas the distributed solution requires replication of servers. In many cases, a mixed workload favors System z due to many factors, including the z/OS Workload Manager.

How does System z workload consolidation drive down TCO?

System z is a large system that is designed to provide a pronounced economy of scale. Information from multiple clients indicates that a single System z core can run many times the amount of work that a single x86 core can run. For example, in 2012, the IBM Eagle TCO team worked with a client on a consolidation project, which demonstrated that one System z core (using the latest z/OS technology, IBM zEnterprise® EC12) was equivalent to 180 x86 cores. Replicating the work that a System z

core can handle on distributed Intel cores would mean that the software bill — one of the biggest expenses — might need to increase dramatically to cover those additional cores. Most enterprises use System z as a central hub, which provides key centralized services and data to multiple applications. If you visualize a bicycle wheel, think of the applications generally being on the outside "spoke" of the wheel, with System z as the hub; they are not typically directly linked together.

You may hear some people say that System z is expensive. Consider, however, the frame of reference. If they are comparing a System z environment to a single Intel box, then they are probably right — System z is more expensive. However, they also need to recognize that a single Intel box is designed for something completely different than System z. System z is designed and engineered to deliver a decreasing cost of work as workloads scale. For a very small system, distributed environments are typically less expensive. System z is not optimized when it runs a small, standalone application. Rather, System z is made to run a large number of databases of different sizes, and to manage them as a single environment for greater efficiency and cost effectiveness. As more workload is added to a System z environment, the average cost of a unit of work goes down, as shown in the following figure.

Example of cost difference between System z and distributed environments

Why do distributed database cost per unit of work **rise** as workloads grow?

- Core proliferation
  - Software priced per core
- Database software costs
  - Largest component of cost
- Hardware cost escalates faster as number of cores grows

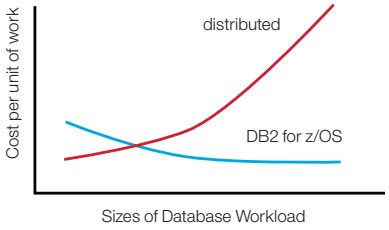

Figure 2: An example of cost differences between System z and distributed environments

Many System z clients find themselves with a large number of data warehouses and data marts that are built from their core operational data on System z. When you look at each individual data warehouse or data mart on its own, it may look efficient. However, when you look at the entire collection of data warehouses or data marts for a single enterprise, you begin to see the cost and inefficiencies that this model can bring.

The following figure illustrates an example of a European bank, for which 28 percent of its distributed systems simply moved data around the enterprise. These same systems used 16 percent of the System z MIPS to pull the data out and move it somewhere else.

Enterprises with complex data movement solutions and unforeseen costs

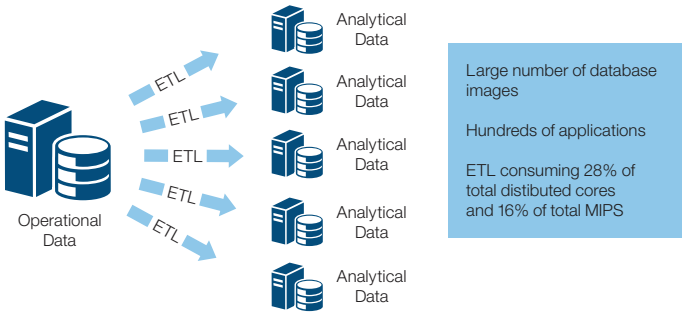

Figure 3: Enterprises with complex data movement solutions and unforeseen costs

This figure does not reflect the additional costs for administering the systems and managing security and compliance for those systems. Nor does the figure take into account the business impact of making decisions based on multiple versions of the truth, which may not correspond with actual operational data.

Additionally, leaving the data on the mainframe can deliver additional important benefits, including:

- The simplicity of running the business from one place, without wasting precious resources to move data around the business
- Real-time analytics for real-time business decisions and value
- A single, secure domain, which the organization can manage and control in one place with minimal resources

How does System z optimize workload consolidation?

Today's applications are more complex than in the past. They have a broad mixture of high-speed transactions and complex queries that must run simultaneously with very low latency of data. Numerous technologies have been added to each platform to accommodate the needs of a growing business. Clients who embrace a fit-for-purpose deployment strategy typically move workloads around. They then add workloads to where they can best optimize the use of technology, and provide the right economics at the right qualities of service. Therefore, a fit-for-purpose deployment strategy often results in a hybrid environment, with some workloads running on distributed systems and others being consolidated on a centralized, mainframe system.

System z provides an enterprise-class solution that brings high-volume business transactions, batch reporting and complex analytic queries together, running concurrently, in a mixed workload environment. This mixed workload environment is available on a single, integrated platform that uses a network-attached accelerator for easy and cost-effective management, which is transparent to applications.

With the introduction of hybrid computing on System z with the zBX, the IBM DB2 Analytics Accelerator, and what IBM has offered for some time with Linux on System z, clients now have a robust choice to host workloads where they best fit. Many successful patterns for fit-for-purpose deployment exist. One such deployment in a multi-tiered environment could include SAP, where data, application-serving layers and presentation layers can be deployed across the System z integrated resources. Data can reside on z/OS, and application-serving and presentation layers can run on the zBX.

Another strong use case exists for real-time analytics. For example, clients in the insurance industry want to prevent the payment of fraudulent claims before they happen, rather than trying to recover payment after funds have been disbursed. With the IBM DB2 Analytics Accelerator and SPSS® software, System z can provide a single-system approach to delivering these capabilities with superior economics. There are also significant, ongoing trends with both new and existing clients to utilize Linux on System z as an economical, adaptable and powerful consolidation engine. In many situations, the modern mainframe offers a broad array of capabilities to host new applications and help clients deliver on significant industry trends in analytics, cloud and mobile, as seen in the following figure.

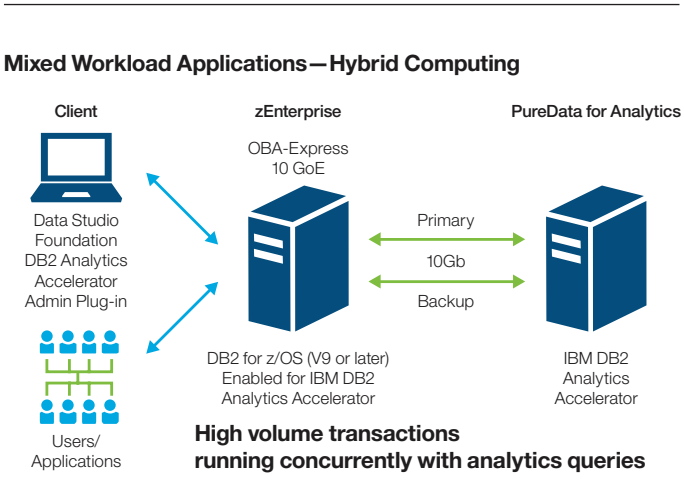

Figure 4: Mixed Workload Applications—Hybrid Computing

What impact does software have on TCO?

Because software is typically one of the largest expense items for IT organizations, it is a significant variable in the TCO equation. To effectively compare software costs on System z to software costs in a distributed environment, a product-by-product analysis must be performed, with pricing on a per-core basis. This comparison can be challenging because of the need to:

- Consider all of the core systems, including what needs to be provisioned for high availability, backup and restore
- Identify and estimate costs for all third-party software products
- Add increased labor costs for functions that cannot be replicated in a distributed environment

Transactional costing is a financially appropriate measure to use in this sort of comparison. For System z software, cost per transaction decreases as volume increases. IBM offers flexible, on-demand hardware and software pricing, which can significantly lower transaction costs for medium to very large workloads. Other factors that make the price go down include:

- Consolidation to fewer, bigger boxes
- Lower incremental costs for large enterprises
- Smaller cost increases as new work is added
- Ability to add all System z workloads together to get a lower price per transaction
- Performance improvements in each software release, leading to fewer resources that can often support the same workload

Because of these factors, the mainframe can deliver a substantial economy of scale as it grows to meet changing capacity requirements. For example:

- In the experience of the IBM Eagle TCO team, doubling capacity on a mainframe creates a cost increase of as little as 30 percent for IBM mainframe software.
- In 2013, an IDC report summarized the results of in-depth interviews of six organizations, on four continents, that had consolidated on zEnterprise and System z Business Class models. The report indicated that software license costs for these organizations dropped by an average of 71 percent.[4]

The distributed story, in contrast, is very different, with a more linear cost increase as systems grow. Each new server may well require an additional software license. In addition, the shared-nothing architecture cannot take advantage of shared database resources like System z enterprises can. System z enterprises can combine development, test and production environments in the same workload with the z/OS Workload Manager, thereby providing the most critical transactions with needed resources while still allowing lower-priority work to complete. With the distributed approach, each separate environment typically needs to be stocked with a full capacity of servers that can handle peak workload.

When looking at software costs as part of a TCO analysis, be sure to factor in all application and middleware software license and maintenance fees. Availability and backup and recovery systems should also be included for a complete TCO comparison.

How do people factor into TCO?

To be complete, a TCO analysis considers not only hardware and software acquisition costs but also costs associated with the people who support the IT environment. This analysis includes database administrators, system administrators, programmers, operators and managers. Even though the labor cost is often the most expensive segment of the TCO equation, it is often forgotten or overlooked.

All core systems, regardless of platform, have administration and maintenance duties to be performed. Every physical device must be racked, stacked, powered on, managed and maintained. Software upgrades, security and maintenance patches, and hardware maintenance and upgrades need to be rolled out and managed on an ongoing basis.

Designed to be a single point of maintenance and control, System z is built to minimize the human resources required to manage the environment. Because you can scale so many applications and so much work on a single box, you can significantly minimize the work required to manage the environment. Security processes can be defined and managed for the entire system. Software and hardware maintenance and patches can be rolled out across the entire environment with more efficiency and less labor. According to a 2013 report[5] that summarized an IDC study of organizations that consolidated servers, these organizations found that IT staff productivity increased for day-to-day server support activities and maintenance. These productivity increases resulted in an average reduction in operational costs of 57 percent. When new workloads are implemented on a mainframe system, they can generally be managed by the existing staff.

In the distributed model, growth is supported by adding more physical systems. The more physical core systems you have, the more time it takes to administer those systems. Administrative tasks add up and can consume valuable resources, thereby keeping IT staff members from working on other innovative tasks that could achieve key business goals. When a distributed environment grows to thousands of servers, the labor costs can be significant. In many cases, the labor costs of managing servers and rolling out patches are harder to calculate because they exist in various business departments.

If System z is so great, why is it such a large IT expense item?

With the help of DB2 for z/OS and IMS database management systems, System z has been a mainstay of many data centers for the last 40 years. To help enterprises manage and track cost, a variety of accounting and billing capabilities are available to be deployed to the cloud. Over time, fixed costs that are more difficult to allocate to a specific project are sometimes added to the mainframe budget so that the project can be billed back and spread across multiple organizations. As long as each area of the enterprise continues to use the mainframe, costs can be appropriately allocated to each area.

However, recent incidents have occurred in which enterprises have included the recovery cost for things like corporate jets in the mainframe budget. When these internal bills get passed along to the various internal business units, the mainframe expense item is viewed as unnecessarily high. Suddenly, people start talking about moving workload off the mainframe to reduce costs.

Months after moving the workload off the mainframe, many enterprises experience an even bigger problem: chargeback costs (for the corporate jet or similar non-IT expenses) are readjusted. They then realize that they should have been looking at the real cost of the mainframe, not the chargeback costs. Unfortunately, this realization usually occurs only after incurring quite a bit of risk and expense trying to replicate and test the environment on a different platform.

How do upgrades and migration projects affect TCO?

Replacing outdated technology, including hardware, has become a regular occurrence in IT departments. Most people take into account the cost of replacing the hardware itself. However, they often forget to include the time and effort involved in refreshing the hardware and software, and in successfully moving to the new technology. System z is designed to help IT departments streamline these upgrades. In fact, some clients are able to install new System z hardware on Friday night, work on some cabling and testing over the weekend, and have all of the applications running again on Monday morning. In contrast, a hardware upgrade scenario is very different in a distributed environment that includes many cores and lacks the centralized computing model of System z. Just how long might such an upgrade take? How many people would need to be involved? All of these expenses need to be included when adding up the total cost of an upgrade.

For projects that involve migration of IT workloads from a mainframe system to a distributed environment, the challenge is to accurately estimate how long the migration will take. A common best practice for do-it-yourself home improvement projects is to automatically add 20 percent to your budget. An IT migration or re-hosting project, from a time-based perspective, is also likely to grow beyond scope and time estimates.

How does provisioning affect TCO?

Availability and failover requirements and costs differ by platform, client needs and the importance of the system in question. For example, what you need for a less important nonproduction environment (such as test, development, disaster recovery and QA systems) can be very different than what you need for a core 24x7 production system. These different requirements lead to very different provisioning needs. In general, nonproduction environments require fewer resources on the mainframe than in a distributed environment. Based on the feedback and experiences of various clients that IBM has worked with in TCO studies, some general guidelines have emerged for provisioning to meet high availability needs for a production environment:

- Clients who use a dedicated failover solution on a **distributed environment** can expect to provision 2.5 times their production environment. Clients who choose an N+1 clustered approach can expect to provision two times their production environment.
- **System z** clients, in contrast, need significantly less resources, typically targeting the same amount of resources as they have in production. In many cases, clients can provision even fewer resources and still meet their availability requirements.

Provisioning ratios for nonproduction environments have slightly different ratios, but they are still significantly smaller on System z. The average size of the machine (production, development and test) is only about 20 percent more than it would be if the machine supported only production workloads.

Distributed solutions have a wide range of provisioning requirements. The size of the production environment plus an additional 200 percent is often set aside, as shown in the following figure. The total provisioned capacity for a distributed system that supports production, development and test can therefore be three times what would be required for production.

Provisioning can clearly result in a significant cost savings for System z solutions.

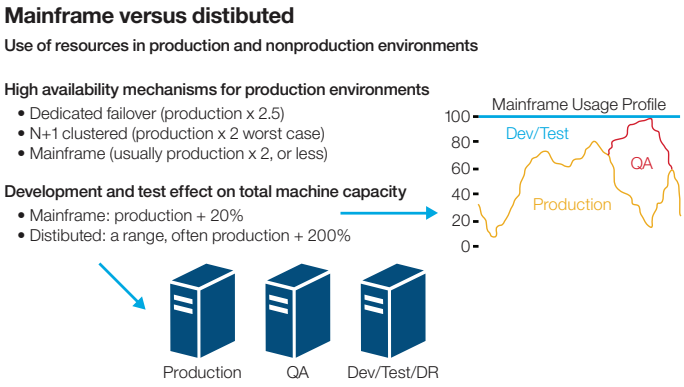

Figure 5: Use of resources in production and nonproduction environments

How does System z minimize my facilities costs?

An often overlooked component of the TCO equation is the cost of facilities, such as floor space and power costs. Although facilities costs typically do not top the list of IT expenses, they can be an important component of the overall operating costs for a production workload infrastructure. A 2011 US Department of Energy paper[6] titled "Best Practices Guide for Energy-Efficient Data Center Design" indicates that, for a typical data center with a highly efficient cooling system, IT equipment loads can account for over half of the entire facility's energy use. Further, in challenging economic times, organizations look for ways to control costs however they can. According to a 2012 issue of CIO Insight[7], 52 percent of organizations have asked their IT groups to reduce energy costs.

Facilities costs, including energy costs, can spiral out of control in a distributed environment because of the failure to maximize the economies of scale of the capacity that may already be available. The same 2011 US Department of Energy paper noted above states that rack servers account for much of the wasted energy and represent a significant portion of the IT energy costs of a typical data center. The paper notes that servers, in particular, take up most of the space. Most of these servers operate at or below a 20 percent utilization rate most of the time, even though they draw full power throughout their operation.

Although the x86 deployment model comes in at a lower percentage of the total cost, it results in a higher bottom-line cost over a five-year period. The percentage of total costs for power and floor space continues to hold steady across various deployment sizes, failing to deliver the economies of scale that are normally expected with growth.

Looking closely at facilities costs, you can see where the total cost savings associated with System z start to surface. The ability of System z to handle mixed workloads and to scale translates to the need for fewer CPUs and servers, which results in a reduced need for floor space and power to run these workloads. The IBM zEnterprise 196 (z196), for example, offers increased capacity, an increased number of available processor cores per server, and reduced energy usage and floor space. With its attractive energy efficiency, the z196 can save up to 75 percent on energy costs, compared to virtualized x86 alternatives.[8]

In 2013, IDC[9] studied organizations that had consolidated on zEnterprise and System z Business Class models. These organizations reported that their facilities costs dropped an average of 66 percent. With System z, organizations can allocate a smaller portion of the budget to facilities costs as they grow. IT leaders can then deliver cost savings to the business and redirect financial resources to other important initiatives.

What is the cost of outages to my enterprise?

The importance of availability in the event of planned and unplanned outages has been well established. Although most enterprises place a high value on availability, they are often unaware of the actual costs associated with outages. Every business has different levels of tolerance for outages. In any case, the high cost of downtime can be minimized by the high availability of System z.

If you do not have your own internal measurements, you can use averages for the financial impact of an hour of downtime. A 2011 paper[10] by Somers Associates, Inc. and W.H. Highleyman, titled "The Availability Digest," provides the following summary table on the estimated direct hourly costs associated with outages, by industry.

Average hourly costs (in US dollars) for downtime, by industry

Brokerage: $6.5 million	Financial: $1.5 million	Chemicals: $704,000
Energy: $2.8 million	Manufacturing: $1.6 million	Health care: $636,000
Credit card: $2.6 million	Retail: $1.1 million	Media: $304,000
Telecom: $2 million	Pharmaceutical: $1 million	Airlines: $90,000

Sources: Network Computing, the Media Group and Contingency Planning Research

Direct costs in these estimated averages include decreased user productivity, regulatory and litigation costs, and opportunity costs such as lost revenue.

According to a report by the Ponemon Institute[11], Dun & Bradstreet reported that 59 percent of Fortune 500 companies experience at least 1.6 hours of downtime per week (about 83 hours per year). The actual length of outage and direct cost for those outages in any given enterprise vary. However, this data illustrates the potential for significant costs. Most enterprises would prefer to avoid these potential costs, and should therefore consider comparing the TCO of different environments.

How does security enter into a TCO discussion?

Security is more than just access to a database. It includes access to the entire system, from the application, to the database, to specific data in a database table, to the operating system, and even to printers. Actual costs and impact naturally vary, but research by the Ponemon Institute gives you an idea of the importance of factoring in the costs and impact of security breaches when you discuss TCO for your environment:

- A 2011 survey of 583 IT security professionals[12] claimed that 90 percent of their employers had suffered at least one data security breach.
- A 2013 report about the cost of security breaches[13] indicates that the average cost for a US company that reported a security breach in 2012 was USD5.4 million.

Beyond the financial cost, a firm's reputation may be damaged. Publicity about a security breach can cause a number of problems for the business.

The threat of security breaches is a reality that warrants planning and preventive actions. When looking at the cost of a security breach, account for the cost of detection, notification, further response or action, compliance and audit actions and repercussions, and, of course, lost business and damage to company reputation.

System z is designed to provide industry-leading levels of security. Distributed solutions often provide security that must be managed independently in the database, hardware and operating system. In contrast, System z provides a proven, integrated intrusion-detection environment, with system-level RACF® protection (from logon to data encryption), and consistent, policy-based security management. With these integrated capabilities, you can provide a cost-effective, manageable enterprise security vault that can save money and pain by preventing security breaches in the first place.

Is there any real-world proof to back this up?

The short answer is yes. Here are some examples of actual client situations.

Example 1

Should I move less-critical workload to a distributed environment?

In one example from the Eagle TCO team, a client had an under-utilized mainframe system that ran only non-critical applications. The client decided to move these applications to a distributed environment, and the results (from 2008) are shown in the following figure.

Core proliferation for a small workload shows relative efficiency (based on an actual client situation)

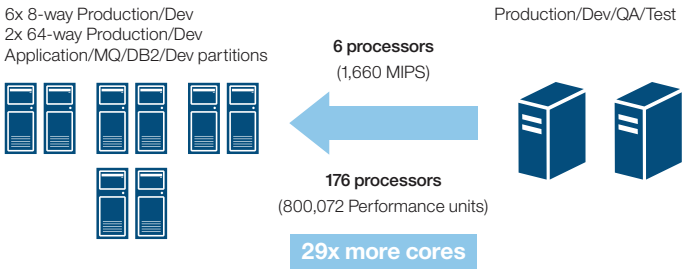

Figure 6: Core proliferation for a small workload shows relative efficiency

The deployment model involved moving each application, covering each function role (such as test, development and disaster recovery), and then breaking those into modular components. After that work was completed, additional capacity provisioning was needed so that workload spikes could be properly handled.

You might think that virtualization technology could handle potential workload spikes. Unfortunately, that was not the case. For these applications, especially for those that ran in production environments, the client needed an environment that would help avoid running out of capacity during workload peaks. By reserving the capacity, they provisioned for their peak workload. Virtualization helped get rid of multiple frames, but the client still ended up with just as much reserved (and generally idle) capacity as if the boxes had been provisioned.

Example 2

Decision to migrate from mainframe to distributed, with costly results.

An ISV launched a new development effort.[14] This organization:

- Opted to build a .NET application to replace an existing mainframe application, an effort that:
  - Spanned nine months
  - Cost USD38 million
  - Resulted in only 20 percent of the required functionality that the mainframe application provided
- Abandoned the effort because of the poor results
- Ultimately developed a web-based, front-end application to modernize the existing mainframe application, an effort that:
  - Was completed in 29 days, with no defects
  - Cost only USD2 million
  - Satisfied 100 percent of the functional requirements for the application

By choosing to develop a new distributed application rather than reusing and modernizing an existing mainframe application, this ISV spent 94 percent more than it needed to and wasted quite a bit of valuable time.

Example 3

Better TCO for new workload on System z. A large not-for-profit US healthcare system uses both Oracle and DB2 for z/OS for its data warehousing applications. In planning for a new enterprise data warehouse, decision makers leaned toward using Oracle. However, after extensive testing, they became convinced that DB2 for z/OS was a better overall solution.

- **Background:** This healthcare system embarked on a high-profile research initiative that required a clinical, dimensional warehouse to manage a massive amount of data about patients, caregivers, medical procedures, diagnoses, calendars, facilities and much more. Having background and experience in both Oracle and DB2 for z/OS, decision makers needed to decide which environment would be best for this new enterprise data warehouse. They knew that a correct decision would be a critical success factor for their strategic research initiative.
- **How this healthcare system evaluated options:** To set up a fair, valid comparison between Oracle and DB2 for z/OS, decision makers planned the competition so that the test environment would simulate their planned warehouse environment, including the number of tables, indexes, rows, security requirements and more.

The client's own comparison of Oracle to DB2 for z/OS uncovered some interesting findings, which pointed to the following benefits of using a DB2 for z/OS and System z environment:

- The System z environment is 100 percent compatible with the proposed enterprise data warehouse model.
- Existing SQL statements could be used, whereas using Oracle would have required new and complex SQL coding.
- No server or disk upgrade was required.
- No incremental software costs were incurred because DB2 for z/OS is already installed.
- The new workload qualified for new-workload pricing deals.
- More robust and less expensive tools were available.
- Centralization was far more energy efficient and green, compared to a distributed server farm solution.
- Quick ability to scale made the enterprise poised for growth.

Recent technological advances at IBM (such as big data, new capabilities around security, IBM DB2 Analytics Accelerator, and Business Analytics for Healthcare) will be extremely useful as the organization drives future extensions to its enterprise data warehouse.

In 2012, the client decided, after extensive testing, to implement its new enterprise data warehouse on System z for a wide variety of reasons, summarized below.

The final decision

Agility
Ability, in just one month, to:

- Load millions of rows of data
- Tune performance
- Test backup and recovery
- Set up automation and monitors for performance
- Run performance benchmarks for SQL and LOAD utility jobs

Security

- The most secure server on the planet.
- DB2 certified for CIA security standard.
- System z has never fallen victim to a virus

Availability

- Client experienced high availability with DB2 for z/OS over 10 years.
- Last unplanned System z outage was 9.5 years ago.
- Non-disruptive daily updates.
- Non-disruptive database utility jobs.

Support
The client shop already has the following in place:

- On-call database support and world-class DB2 expertise
- Automated utilities, metrics, notification, and quality measures

Compatibility

- 100 percent compatible with client data model.
- Compatible with existing SQL.
- Unicode data supports every written language on the planet

Toolsets

- Eighty existing toolsets already run with DB2 for z/OS.
- Operating system and database monitoring tools are in place.
- World-class database utilities.
- Fully integrated BI suite that is available today.
- IBM BI tools also support cross-platform federated data.

Costs

- No server upgrade or disk upgrade is needed.
- No incremental software costs.
- Text mining extender included in DB2 for z/OS license.

How do I get started doing a TCO analysis?

Ask this question, and you have already taken the first step. Decisions about the best fit for each IT workload must be based on more than just the hardware and software acquisition costs. As discussed, the total cost of ownership is influenced by many other factors. Although the task of estimating TCO is not easy, IBM offers resources to help you along the way.

Information Management TCO developerWorks page

Visit the IBM System z Total Cost of Ownership page in developerWorks® **ibm.com**/developerworks/community/groups/community/systemztco for links to a variety of different resources that provide additional details about TCO issues and IBM products that can help you minimize your TCO. This page will be updated periodically with new resources. To be notified of updates to this page, join the developerWorks community, and click **Follow this Community**, as shown below:

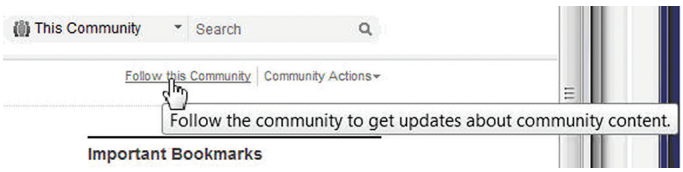

The IBM Worldwide Eagle TCO Engagement Services team

Also known as the Eagle TCO team, this group of subject-matter experts has been working with different enterprises to understand their costs and challenges for several years. These enterprises include:

- Those who are debating whether to move part or all of their workload off of System z
- Those who already moved (or tried to move) off of System z with the hope of saving money, but are not sure they made the right choice
- Those who are trying to decide what platform they should deploy their new workload on for the best business value

When working with a typical enterprise, the Eagle TCO team meets with representatives from the enterprise in person to kick off the TCO study and tackle Steps 1 and 2 below. Then, the Eagle TCO team works on its own to complete the study (Steps 3 through 5), which usually takes about 30 days. The steps for an Eagle TCO study are:

1. Establish the scope of the study (applications, platforms, and so on)
2. Gather information about the enterprise, including planned growth
3. Build the cost model
4. Review findings with the client, and iterate as needed
5. Produce a final report

Interested in having the IBM Eagle TCO team do a TCO study for your enterprise?

For no charge to you or your organization, this team of IBM subject-matter experts can gather details about your enterprise and help you decide on a best-fit solution for your IT workload. They will help guide you to making the best-fit decision for your particular situation. Contact eagletco@us.ibm.com to learn more.

References

1 Arcati Mainframe Yearbook 2013, published by Arcati Limited, 2013 http://www.arcati.com/newyearbook13

2 "*Transforming Disaster Recovery with Continuous Availability*," published by EMC Forum 2013 Lead Your Transformation, 2013 http://www.emc.com/campaign/global/forum2013/pdf/trasforming-disaster-recov.pdf

3 "*IBM System z: The Lowest-Cost Database Server Solution*," written by Cal Braunstein in Enterprise Executive, published by Enterprise Systems Media, April 2013 http://enterprisesystemsmedia.com/article/ibm-system-z-the-lowest-cost-database-server-solution

4 "*The Business Value of IBM zEnterprise System Deployments*," sponsored by IBM, published by IDC, July 2013 http://www.techrepublic.com/resource-library/whitepapers/idc-the-business-value-of-ibm-zenterprise-system-deployments/

5 "*The Business Value of IBM zEnterprise System Deployments*," sponsored by IBM, published by IDC, July 2013 http://www.techrepublic.com/resource-library/whitepapers/idc-the-business-value-of-ibm-zenterprise-system-deployments/

6 "*Best Practices Guide for Energy-Efficient Data Center Design*," published by the US Department of Energy, revised March 2011 http://www1.eere.energy.gov/femp/pdfs/eedatacenterbestpractices.pdf

7 "*IT Management Slideshow: Saving Money By Cutting Data Center Energy Costs*," by Dennis McCafferty, published by CIO Insight, April 2012 http://www.cioinsight.com/c/a/IT-Management/Saving-Money-By-Cutting-Data-Center-Energy-Costs-673784/

8 "*Green IT – Energy Efficiency*," published by IBM, April 2013 http://www-03.ibm.com/systems/z/advantages/energy

9 "*The Business Value of IBM zEnterprise System Deployments*," sponsored by IBM, published by IDC, July 2013 http://www.techrepublic.com/resource-library/whitepapers/idc-the-business-value-of-ibm-zenterprise-system-deployments/

10 "*The Availability Digest*," published by Sombers Associates, Inc., and W. H. Highleyman, June 2011 http://www.availabilitydigest.com/public_articles/0606/value_of_availability.pdf

11 "*Disaster Recovery & Cost of Downtime*," published by 10gea Trusted Telecom Solutions, 2013 http://www.10gea.org/articles/disaster-recovery-cost-of-downtime/

12 "*Perceptions About Network Security: Survey of IT & IT security practitioners in the U.S.*," published by Ponemon Institute, June 2011 http://www.juniper.net/us/en/local/pdf/additional-resources/ponemon-perceptions-network-security.pdf

13 "*2013 Cost of Data Breach Study: Global Analysis*," published by the Ponemon Institute, May 2013 https://www4.symantec.com/mktginfo/whitepaper/053013_GL_NA_WP_Ponemon-2013-Cost-of-a-Data-Breach-Report_daiNA_cta72382.pdf

14 "*Enterprise Modernization: A Rational Approach*," published by IBM, July 2011 http://public.dhe.ibm.com/common/ssi/ecm/en/raw14274usen/RAW14274USEN.PDF

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
November 2013

IMW14745-USEN-000