

Taking Back Control of Your IT Infrastructure: Consolidating to a Data Center in a Box



Redguides
for Business Leaders



Alex Louwe Kooijmans
Ogechi Nduka
Liv Samuels
John Schlosser
Christian Strauer

- Describes perceived CIO challenges and findings from a CIO Global Study
- Explains the need for server consolidation and the steps to do it
- Examines how IBM zEnterprise System offers a powerful solution



Executive overview

Businesses today are innovating rapidly and the impact on the IT infrastructure is significant. Innovations such as SOA, Web 2.0, analytics, and even cloud computing can increase the IT infrastructure and also increase the complexity of the environment. The IT infrastructure needs to be well designed and planned, and ready to grow with new business initiatives. In the next few years, the required IT infrastructure capacity for any enterprise might double or triple; beyond that, we do not yet know how much more will be needed. It depends on which innovations in our way of doing business await us in the future; some of this is really unpredictable. However, we do know that the IT infrastructure should be designed in a way that it can grow exponentially without resulting in a total gridlock of complexity, and that addresses the limited floor space and the rising energy costs.

If you are already affected by these kinds of issues in your environment, you can be certain they will become worse in the future unless you rethink how to run your data center. One of the most challenging areas in the data center is the large number of servers. Server consolidation is one of the most powerful methods to control your massive IT infrastructure growth and complexity, and also the optimization and simplification of the environment.

Although server consolidation means going from many servers to fewer servers, it is also important to simplify the management of all these servers, even after they are consolidated. Ideally, the server infrastructure will be entirely virtualized and will provide the maximum scalability.

The IBM® zEnterprise System offers a powerful solution for both the “quantity” issue and the “complexity” issue. It provides a lot of power within one single box and is secure, fully virtualized, and supports multiple operating systems. Its massive compute power and scalability can be provided on a small floor space while using less energy. The complexity is reduced by providing a single management interface that is used to provision server capacity. The zEnterprise also has its own internal high-speed and secure network, connecting all the various servers, making it truly a “data center in a box.” If you are convinced that you need to do something in your data center to be prepared for the unpredictable future, IBM can help with various proven studies and methodologies to assess your current situation and estimate the outcome of your transition, measured in real cost and value.

Archived



Business challenges perceived by CIOs today

How are technology leaders helping their organizations adapt to the accelerating change and complexity that mark today's competitive and economic landscape? To find out, IBM spoke with 3,018 CIOs, spanning 71 countries and 18 industries. Public and private sector CIOs described both the challenges and opportunities from increasing complexity. The results of this survey can be found in the 2011 IBM Global CIO Study.¹

Based on this study, Figure 1 on page 4 shows the top priorities of CIOs for 2011 compared to 2009. *Business intelligence and analytics* is the number one priority, followed by *mobility solutions*, and *virtualization*. Compared to 2009, the importance of *cloud computing* has massively grown and overtaken *business process management (BPM)*, *risk management and compliance*, and also *self-service portals* and *collaboration*.

¹ <http://www.ibm.com/ciostudy>




Figure 1 CIO top priorities 2011 compared to 2009, as measured in 2011 IBM CIO study

All of these topics have become important on the CIOs' agenda to respond to the business requirements the CIOs face.

In addition to these priorities, the 2011 CIO study identified four CIO mandates (see Figure 2 on page 5 for more details):

- ▶ Expand
- ▶ Leverage
- ▶ Transform
- ▶ Pioneer

An effective CIO should know and deliver one of these four mandates, based on the goals of the CIO's enterprise.



Figure 2 The four CIO mandates

Note: Especially for those CIOs who have a Leverage mandate, IT standardization and consolidation is a high focus area as mentioned in the 2011 CIO study.

Because IT standardization and consolidation helps enterprises to leverage existing IT components more efficiently, we will have a closer look at this important topic in the next sections. Almost all the top CIO priorities mentioned before have a close relation to IT standardization and consolidation. Therefore, we outline the business challenges associated with these priorities in the context of IT infrastructure standardization and consolidation.

Massive growth of business intelligence workloads creates new demands

Business intelligence (BI) is increasingly becoming more important for enterprises to gain competitive advantage in the market. One of the most effective ways to achieve business intelligence is the intelligent analysis of existing data. Because of that, CIOs worldwide initiate BI projects to support their business in a more efficient way.

These projects create new requirements and challenges in many respects. From an infrastructure point of view, requirements especially regarding the management of the massively increasing amount of data arise. According to the Experton Group, IT organizations maintain scores of replicated databases and spend a significant portion of their resources on archiving, backup and synchronization of the duplicate databases.² This high overhead costs companies a lot of money and prevents them from investing their existing resources in new projects. Furthermore, excessive replication of databases often leads to situations where

² <http://www.experton-group.com/home/archive/news/article/the-zenterprise-a-disruptive-technology.html>

businesses rely on data that is not current at all. Also, data replication creates more hops and therefore increases the IT infrastructure complexity.

Mobile solutions increase IT infrastructure agility

Besides business intelligence, mobile solutions are becoming extremely important to keep pace in today's markets. The mobile channel allows enterprises to reach out to their clients in a new way and helps to gain competitive advantages, for example through new intuitive applications (apps) for smartphones.

To provide excellent service through the new mobile channel, a solid and resilient IT infrastructure is required. It has to be scalable, secure, and highly available. Furthermore, to work smoothly in conjunction with the back-end infrastructure, end-to-end management is required. However, currently many solutions are managed on individual components rather than end-to-end.

Ineffective resource utilization causes server sprawl

Apart from new workloads, such as BI and mobile solutions, virtualization has an important role on the agenda of CIOs.

Because many servers do not use virtualization effectively today, enterprises still have many underutilized server resources. This, coupled with increasing server capacity demands, results in server sprawl. Therefore, many companies have a high energy consumption in their data centers, and floor space is not used effectively. Also, the server sprawl increases the maintenance efforts of companies, increases complexity, and increases overall costs. Furthermore, non-virtualized environments have the disadvantage that software license costs are often higher than in virtualized environments.

Slow IT service provisioning reduces time-to-market

Another important aspect that CIOs today often have to deal with is the response time of IT infrastructure services to new business requirements.

The provisioning of IT resources such as servers, network, and storage often takes a long time and prevents an agile IT delivery to the business. Especially the delivery of desktop environments, application development environments, test environments, web servers, and collaboration environments require a dynamic provisioning, which many companies are not able to fulfill today because of processes that are too complex. This slow provisioning can negatively affect the time-to-market, give the competition an edge, and negatively affect the bottom line. To address these issues, cloud computing and self-service portals have become popular (see Figure 1 on page 4). After all, the promise of cloud computing is to eliminate or at least hide the provisioning complexity.

Often, provisioning spans more than only one IT department, for example, you might need multiple servers with separate operating systems, a network and storage for a web application including a database, an application server, and a web server. This way is especially true in larger organizations, where typically separate IT departments exist for managing the various server platforms, network, storage, and other IT components. Besides the fact that IT service provisioning frequently depends on manual processes and lacks

automation in general, it takes even longer in cross-platform cases because of siloed organizations.

One option to reduce the server provisioning overhead in cross-platform situations could be to reduce the number of platforms involved, for example through consolidation on one. However, the “one-size-fits-all” approach has several major disadvantages because each kind of application has its own characteristics that are driving platform selection. In contrast, a “fit-for-purpose” approach can and should be taken to identify the right destination platform depending on the non-functional requirements (NFRs) and service level agreements (SLAs) of an application.

By simply providing one server platform, enterprises can have difficulty fulfilling all the various application requirements from both technology and cost perspectives. For example, if there is only one application requiring an extremely high SLA, an expensive platform providing this SLA for all applications has to be chosen although most of the applications do need it.

Business process management solutions generate excessive CPU cost and complexity

One important aspect of business process management (BPM) is the ability of a business unit to measure its success. For example, are all health insurance claims handled within an average of two working days? Another aspect of BPM is to speed up processing by means of automated workload management and straight-through processing (STP). These two aspects require sophisticated middleware and integration, ideally based on state-of-the-art technologies.

From an infrastructure point of view, the new arising BPM solutions create new challenges, simply because much more CPU power is needed. Another aspect to consider is integration: as BPM solutions tend to span multiple systems to integrate a variety of components, the management complexity increases. This management complexity goes along with an increased error vulnerability, which can reduce IT serviceability in certain cases.

Managing risk and compliance requirements increase IT security demands

Because of governmental regulations and the potential risk of business exposures, security is an important topic that CIOs should focus on as a part of risk and compliance management.

Providing dynamic and flexible environments based on cloud and virtualization clearly offers many advantages, but at the same time raises more concerns regarding security. Keeping or even improving the existing security level is key to prevent a potential loss of reputation. Virtualized environments that do not provide a high level of isolation can produce a real security issue. This security issue can have a high impact on the business of a company and eliminate the advantages of consolidation within an extremely short time frame.

From a security perspective, isolation of server images is not the only challenge, the network is also a key factor. The complexity in networks can increase drastically in virtualized environments. Therefore, this complexity can lead to security exposures, for example, communication to systems from unauthorized users because of errors in network configurations.

As IT has become a crucial part of many businesses, more governmental requirements arise regarding the ability to have a high availability and disaster recovery concept in place. These are local and often industry-specific requirements; being compliant or prepared to fulfill them is an important aspect for enterprises.

Archived



Solving today's business challenges through server consolidation

The challenges discussed earlier require a superb agile IT infrastructure, with the ability to grow and scale continuously and in significant proportions. Server consolidation is an important step, and in fact one of the first prerequisites for being successful with each challenge.

In the following sections, we explain how server consolidation plays a critical role in each challenge. First, however, we explain what we actually mean by server consolidation and how the behavior of the ideal IT infrastructure should look.

Introduction to server consolidation

Server consolidation is not a new topic and most companies have gone through one or more server consolidation projects. However, the project might not all have been with the maximum possible results. Organizational structures, dependencies and limitations at the application level, security constraints, and other reasons might have prevented executing a “full” server consolidation. Basically, a “full” server consolidation is seen as sharing and virtualizing as many logical servers as possible in as little physical server capacity as possible. The primary goals are obvious:

- ▶ Reduce floor space, energy cost, and software license cost by hosting many server images in one physical server.
- ▶ Improve scalability and elasticity of the environment by sharing and virtualizing resources.
- ▶ Reduce complexity by streamlining server management tools and processes.

Definition: We define server consolidation as an approach to reduce the number of physical servers through server resource virtualization and sharing. Virtualization and resource sharing allow the running of multiple logical instances of servers on one physical server.

The primary goals of server consolidation are reducing data center floor space, energy cost, software license spending and complexity, and also improving scalability and elasticity of the IT infrastructure.

To better understand consolidation advantages, we outline how an ideal server consolidation environment should be.

Figure 3 on page 11 shows an architecture overview of such a consolidation environment that includes the following items:

- ▶ Multiple types of processors in one system to support application requirements based on a “fit-for-purpose” approach
- ▶ Multiple operating systems, providing various qualities of service
- ▶ A storage system that supports these multiple operating systems in one physical box
- ▶ Virtualization at all layers:
 - Servers including central processing unit (CPU) and input/output (I/O) virtualization
 - Network
 - Storage
- ▶ A high level of security because of a high level of isolation between server images
- ▶ Dynamic resource allocation based on application priorities
- ▶ A provisioning engine that supports various kinds of servers, storage, and network
- ▶ A granular accounting system across various operating systems
- ▶ Cross operating system, high availability, and disaster recovery across two sites, supporting various service level agreements (SLAs)
- ▶ A common infrastructure management and user interface (UI) for all servers types, including centrally managed firmware updates

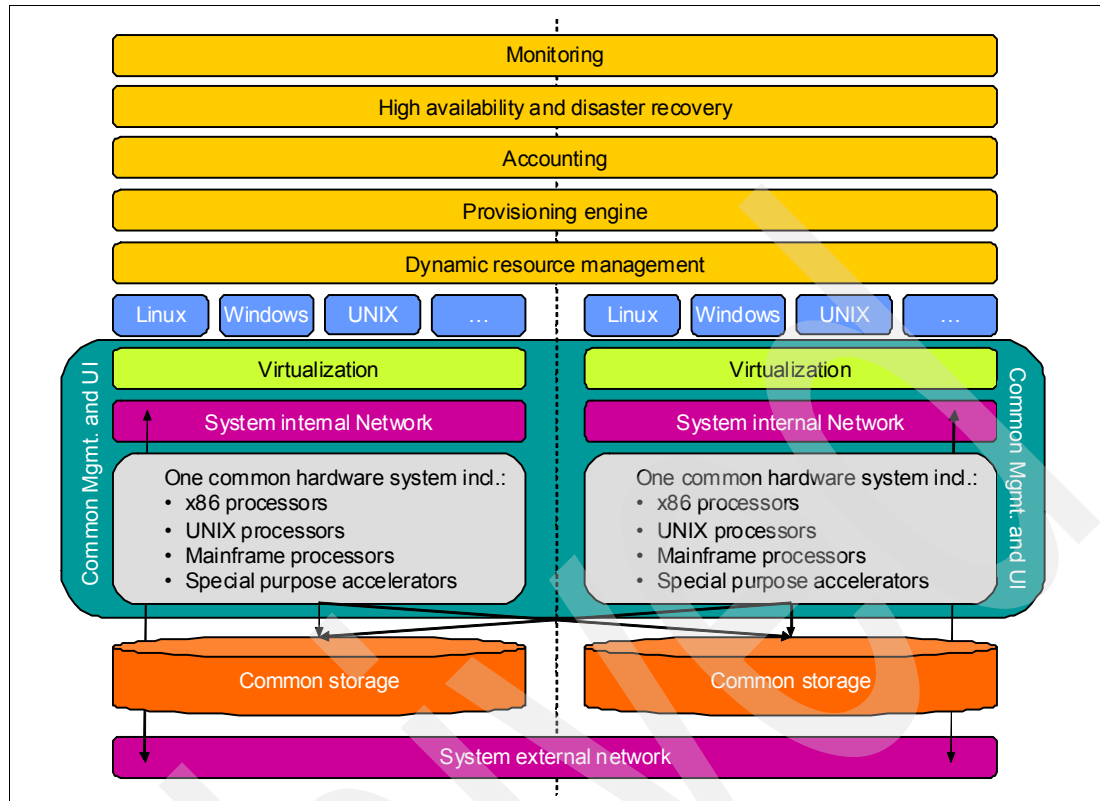


Figure 3 Overview of an ideal consolidation environment

Based on this cross-platform consolidation environment sample, we explain in the following sections how server consolidation helps to address the business challenges mentioned in “Business challenges perceived by CIOs today” on page 3.

Serving business intelligence in a more efficient way

Serving the increasing number of new workloads in the BI space requires both intelligent software and an efficient IT infrastructure. BI solutions typically involve more than one software component, usually at a minimum a database, a data warehouse, an ETL¹ tool, BI software, and a web server. Larger environments have multiples of each of these components running on multiple servers.

Also, many times one or more components are running on a server that is not designed to meet the agreed SLA. A good example of this is a BI tool running on a desktop computer that was once installed for a non-business critical reporting requirement. However, as business evolved, the functions provided by the BI tool slowly became business-critical, but the BI tool was never relocated to a more robust server.

Important: Typical business intelligence environments have multiple components running on multiple servers. Although databases are typically implemented on the “best fit” server, often the other components, most notably the BI software itself, are implemented on servers that do not fit the SLA requirements.

¹ Extract, transform, and load

As explained before, the “one-size fits all” approach has a disadvantage: fulfilling the various non-functional requirements of each workload component in an efficient manner is extremely difficult. Because of that, an ideal consolidation environment should support the “fit-for-purpose” approach, which means the environment has to support various operating systems and hardware implementations under one “umbrella” (see Figure 3). Such a system has the following advantages:

- ▶ Each component of the BI environment is running on the best fitting operating system and hardware depending on its requirements. This allows the support of the SLA of the application in an efficient way.
- ▶ Because of a common management and monitoring infrastructure, the BI application is seen and managed end-to-end instead of having silos with separate views.
- ▶ Accounting can be done end-to-end across all BI application components, thereby providing accurate charge-back to the business.
- ▶ High availability and disaster recovery can be implemented across platforms on an application base rather than on an operating system base.
- ▶ Multi-tier applications can be easily deployed based on a common user interface and fewer organizations.

Siloed organizations often tend to replicate data because each department wants to be the owner of its own data. However, because replication is an asynchronous process, the probability of inaccurate data situations increases. An integrated consolidation system would have the advantage that replication can be reduced to a minimum. This leads to more accurate data that an organization can rely on, and also less network, server, and storage overhead. Furthermore, less replication brings organizations forward on the way to more real-time data access.

Another advantage of consolidating BI environments is data proximity. Servers that have been transferring huge amounts of data over cables, switches, and routers before, can come closer in a consolidated environment. Ideally, they do not even need a physical cable and can communicate cross-memory with each other. This approach improves performance and helps again to improve data currency because of faster ETL processes.

Important: BI environments sometimes result in a total gridlock of massive ETL processes. Co-location of data and processes in the same server environment can significantly reduce ETL processes. As a result, an ETL process can run more frequently, providing more timely information to decision makers.

Server consolidation value in BI applications: Overall, an integrated cross-platform consolidation environment allows the support of business intelligence applications in an effective way. Depending on the required SLAs, each application component can be placed on the best fitting server platform while managing them as whole from an infrastructure perspective. Data replication in BI applications can be reduced drastically in a centrally consolidated environment to improve data currency.

Improving service quality of mobile solutions through better scalability

Similar to BI, mobile solutions are typical multi-tier applications, typically including a database, an application server, a web server and some security components, such as a firewall and proxy server. Therefore, the “fit-for-purpose” support applies to mobile solutions also, based on a consolidation platform to provide an optimum level of service.

Mobile solutions have the potential to generate unpredictable amounts of requests and need a highly scalable server infrastructure to fulfill the required SLAs. Actually, the server environment should be *elastic*, meaning the server capacity is there when needed, but released again when not needed. This elasticity can only be achieved in an environment where resources are shared, servers are virtualized, and intelligent software is in place to manage all of it. In this case, server consolidation definitely makes sense. An ideal consolidated environment is typically more flexible because of virtualization and offers dynamic resource provisioning. Therefore, better scalability can be achieved, which helps to provide customers a consistent responsive service through the mobile channel.

Another aspect that a good consolidation environment should be able to provide in respect to mobile application scalability is *scale-up* (adding resources to a single node) and *scale-out* (adding more nodes to a system) capabilities.

Furthermore, depending on the type of mobile solution, the security infrastructure should be “fit-for-purpose” too, which can be an additional challenge, when combined with a “share everything” approach.

Server consolidation value in mobile solutions: Three items come to mind regarding the mobile channel:

- ▶ Varying and sometimes unpredictable workload
- ▶ Security threats
- ▶ Integration

What is needed is an elastic environment (scaling up and down as needed), with tight integration between the various components, while enforcing strict security. Server consolidation helps to achieve this approach, but the consolidation system has to meet specific requirements to make a good security implementation possible.

Reducing server sprawl through virtualization

Because virtualization is a key factor to reduce energy cost, floor space, and server sprawl, we examine it more closely.

Figure 4 on page 14 shows how virtualization works in general. Typically, virtualization maps physical resources such as memory, disk drives, networks, and CPUs to virtual resources on the target system. Then, the target system sees only the virtual resources, but “thinks” it is using physical resources.

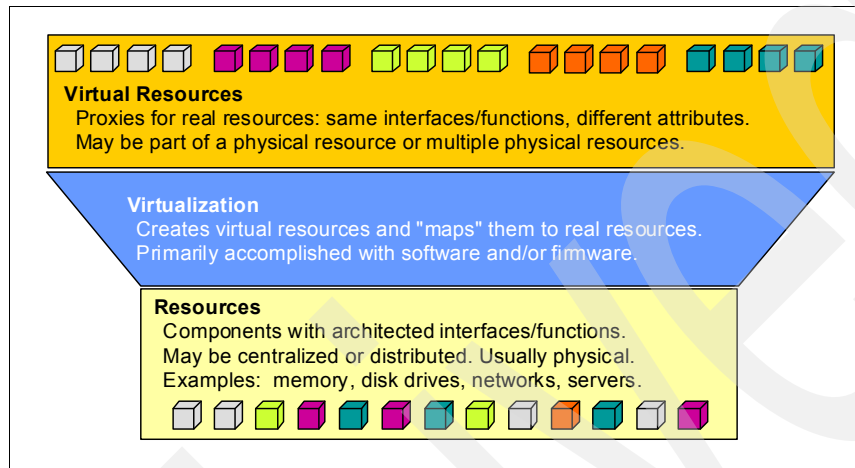


Figure 4 Virtualization in general

Servers that are not using virtualization normally have a varying utilization rate and cause under-utilized resources. However, if you consolidate these servers with various workload profiles on one server using virtualization technology, a more efficient utilization can be achieved, as shown in Figure 5. The ultimate goal is to achieve an average utilization of almost 100%.




Figure 5 Improved server utilization because of virtualization

Therefore, virtualization is a key component to reduce energy cost and floor space through better utilization rates. Consequently, server sprawl, because of increasing server capacity demands, can be reduced or even stopped. Furthermore, asset management becomes easier because of less physical hardware. In addition, less hardware can reduce maintenance and people cost.

Another big advantage of virtualization is that because of less white space, software investments are more effectively used. This way might reduce software license cost and can save companies a lot of money.

An ideal consolidation environment should be able to virtualize all layers:

- ▶ Server (CPU, memory, I/O)
- ▶ Storage
- ▶ Network

As explained before, consolidation environments should also support multiple hardware implementations, such as various processor types. Consequently, virtualization has to be supported for all of them. If virtualization is used on all platforms, a common management interface can be established to build an abstraction layer of the hardware components.

Finally, a central consolidation environment with as few hardware foot prints as possible should be able to dynamically adjust resources based on business priorities. This way includes the ability to add resources such as CPU “on demand” without any interruptions.

Server consolidation value in virtualization: Virtualization enables doing more with less. Hardware and software spending can be reduced drastically and server sprawl can be stopped. To achieve a maximum effect, a central consolidation environment must be virtualized at all layers and across platforms.

Improving time-to-market with the help of automatic resource provisioning

To improve time-to-market, a faster response time of IT infrastructure services to new business requirements is key. Therefore, IT infrastructure provisioning time has to improve. To address this issue, private clouds are a good option because they offer completely automated IT resource provisioning combined with self-service portals.

To provide private cloud services, a virtualized IT infrastructure is required because it provides the interface to IT resource provisioning tools. Therefore, a consolidation environment should be completely virtualized, as explained in “Reducing server sprawl through virtualization” on page 14. To speed up the provisioning process in such environments, completely automated provisioning tools are required.

If a consolidation environment supports multiple operating systems and processor types as explained previously, the provisioning of resources should be performed completely automated across platforms based on the SLA requirements of each application component. With this approach, enterprises can provide IT infrastructure services fast at any time based on business requirements.

Combined with a cross-platform accounting infrastructure, accurate billing can be created for the users.

Server consolidation value in provisioning: Private cloud offerings supporting multiple platforms can improve time-to-market and support business requirements in an optimized way.

Reducing CPU cost and complexity of business process management solutions through tighter system integration

Business process management (BPM) solutions are characterized by strong integration requirements and significant CPU consumption. A modern BPM architecture is based on SOA principles and ties together many services in various applications. Also, a BPM engine typically uses multiple abstraction layers and frameworks, resulting in more CPU consumption (when compared to traditional third generation programming languages).

To be able to support the increasing CPU demand of BPM solutions, a consolidation environment must be able to handle the intensive resource usage in an efficient manner. Various kinds of CPUs or even optimizers in the form of an appliance for various kinds of application characteristics are a good option to reduce the total CPU demand. As certain processors support specific BPM applications in a more efficient manner, the already explained “fit-for-purpose” approach can help to improve BPM service quality and cost.

Because of the strong integration requirements, BPM solutions often result in spanning multiple platforms. This again leads to longer deployments and more complex end-to-end management of the entire BPM solution. A consolidation environment, as shown in Figure 3 on page 11, supports everything under one umbrella and can reduce such impact because of a central management instead of silos.

Server consolidation value in BPM applications: BPM solutions require much integration and CPU power. A consolidation environment with specific processors improves BPM service quality because of a more effective CPU usage. In addition, a cross-platform management infrastructure helps to improve the BPM operating and deployment.

Solving risk and compliance issues because of higher level of security

The loss of a company’s reputation because of IT security exposures is a serious issue that must be addressed. Government regulations also play an important role from a security point of view. Ideally, a server consolidation environment provides the following security components:

- ▶ The server virtualization layer, which must provide a high level of server image isolation to prevent unintended access between systems. A good indicator for this approach is a high *Evaluation Assurance Level* (EAL) certification for the virtualization layer.
- ▶ A simplified, virtualized, and protected internal network that makes attacks at the network level difficult or impossible. Such a simplified network is also easier to maintain and consequently reduces the probability of human errors. Apart from the internal network, communications to the system external network should be reduced to a minimum.
- ▶ A hardware based cryptography acceleration to use cryptography faster and more secure.
- ▶ Fine granular access control for systems management to reduce access to certain critical resources to a minimum.

Because a consolidation server hosts more server images than an unvirtualized single server, the availability of a consolidation environment becomes more critical. For example, in case of a hardware failure, all images are not reachable. Therefore, a consolidation environment should be highly available in general to be compliant regarding potential governmental regulations, which require a certain business continuity. For disaster recovery, these regulations often apply as well also. Therefore, disaster recovery should be done as fast as

possible in a consolidation environment, ideally without any data loss. As explained previously, a good consolidation environment must support multiple operating systems and processor types. Therefore, data integrity must be managed consistently across platforms in an ideal consolidation environment.

Server consolidation value in risk management and compliance: Consolidation environments should provide a server image isolation level as high as possible to prevent business exposures as a result of security problems. Also, a simplified internal network should be used to improve security and reduce complexity.

Finally, a consolidation platform should provide an efficient cross platform, high availability, and disaster recovery to avoid compliance issues.

Archived



Introducing a data center in a box for server consolidation

After having described the characteristics of a consolidation environment and also the value it brings to each of the challenges, we now explain concrete solution approaches.

As explained in the chapter “Solving today’s business challenges through server consolidation” on page 9, a completely virtualized environment supporting various operating systems and processor types under one umbrella is a key requirement to run applications in an optimized way while reducing complexity, resource consumption and floor space. Furthermore, an extremely efficient virtualization, cross platform provisioning and a high level of security have to be supported by such an environment.

Until recently, only selective solutions were available, which fulfill a small part of the overall requirements for a perfect consolidation environment. For example, many virtualization solutions are available, but most of them are quite limited in scalability and only applicable to one specific operating system, processor type, or both. Another example is the virtual local area network (VLAN) technology which already exists for network virtualization. However, VLANs are not integrated into the server management technology because they have a different management interface. These issues force enterprises to implement separate processes and even departments to manage these various ways of virtualization, finally leading to longer server provisioning times.

With the introduction of the IBM zEnterprise System, this approach has all changed. zEnterprise provides the ultimate scalable computer power in one box, choice of processor types, and operating systems, all virtualized and securely connected by an internal network and managed by using one single interface.

Figure 6 on page 20 shows how the zEnterprise System can be used in a consolidation context: storage, network, and server components are completely virtualized; a common management infrastructure controls the infrastructure components. This common management infrastructure, called IBM zEnterprise Unified Resource Manager, provides unified monitoring and dynamic resource adjustments for a hybrid environment of various processor types and virtualization technologies.

Currently, the following environments are supported:

- ▶ IBM z/OS®, running on IBM System z® processors
- ▶ Linux (Red Hat or SUSE) running on System z processors
- ▶ Linux on x86 processors
- ▶ IBM AIX® on IBM POWER7™ processors
- ▶ Windows on x86 processors¹

Based on IBM DS8800 technology, disk storage located in one physical box can be shared with all operating systems running on IBM zEnterprise. Combining high availability and disaster recovery and also common accounting and provisioning based on IBM Tivoli® solutions, the IBM zEnterprise System and IBM DS8800 combined are a real data center in box.

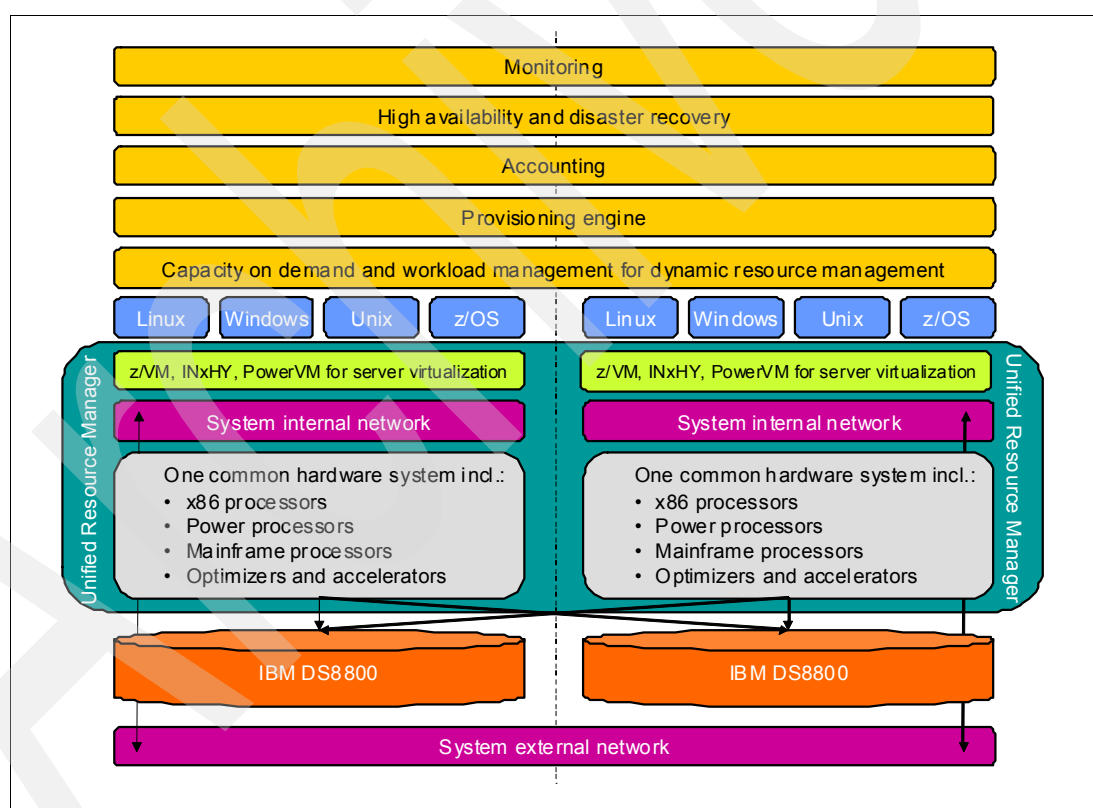


Figure 6 IBM zEnterprise technical overview

Next, we look more closely at how zEnterprise helps to address the specific CIO challenges outlined previously.

¹ Statement of Direction: In the future, IBM intends to offer select IBM System x® blades running Microsoft Windows in the IBM zEnterprise BladeCenter® Extension Model 002.

Running business intelligence in a completely integrated data center in box

For business intelligence workloads, the integrated hybrid concept of zEnterprise is a perfect match. For example, IBM DB2® on z/OS can be used as a central operational database and as a data warehouse. The application servers hosting the BI software could run, depending on their specific non-functional requirements, on another supported platform such as AIX, Linux, or Windows. This integrated concept reduces the communication overhead related to extract, transform, and load (ETL) processes, because the servers running the various operating systems are interconnected through a high-speed (10 GB) and secure internal network, with a minimum of hops through switches and routers. When using z/OS together with Linux on System z, this way is even more optimized because of IBM HiperSockets™ technology. With this technology, communication is between separate logical images based on cross-memory communication technology instead of a physical wire. Also, when using HiperSockets, most TCP/IP overhead is avoided. Furthermore, all BI application components running on separate systems are managed as one workload, making it easier to meet the SLA, which is usually set at the application level, not at the component level.

If you are using DB2 on z/OS for online transaction processing (OLTP) and a central data warehouse, you can reduce data replication time significantly while relying on a more than 45-year proven platform. By having almost no data replication, enterprises can count on real-time data and improve their business performance. Compared to solutions that are commonly used today, this improvement is major, from a data currency point of view.

In addition, with the help of the IBM Smart Analytics Optimizer, which is a specially designed accelerator that can improve the performance of specific data warehouse queries by a factor of up to ten and even more, enterprises can improve their data warehouse performance drastically, which also effects the performance of BI workloads, of course.

Besides less data replication, faster cross-server communication, and improved query performance for BI workloads, another important aspect to gain advantages of consolidation is to have a highly efficient virtualization. See “A virtualized everything environment to achieve optimum resource utilization” on page 23 for further details about how zEnterprise addresses this issue.

Consolidation on zEnterprise value for BI applications: zEnterprise can reduce data replication to a minimum through the combination of a central database supporting millions of users and co-location of all components of the BI solution in one physical server. This enables companies to rely on more current data and improve business performance. Furthermore, data co-location in one single system combined with the IBM Smart Analytics Optimizer improves data warehouse performance massively and makes zEnterprise the perfect match for business intelligence solutions.

End-to-end mobile workload scalability in a hybrid system

zEnterprise is a perfect option as a consolidation platform for mobile solutions because it supports all the specific characteristics that a mobile workload typically has. This support includes a high available database with DB2 on z/OS and also various options for the application servers, web servers, and firewalls. Various criteria might influence the operating system and processor selection for each component:

- ▶ Availability requirements (for example, high for the database; medium for application servers, web servers, and firewalls because they can often manage availability on their own)
- ▶ Skill availability regarding the various operating systems
- ▶ I/O requirements (for example, high for the database; medium to low for application servers, web servers, and firewalls)
- ▶ CPU requirements (for example, medium for the database; high for application servers, web servers, and firewalls)
- ▶ Network overhead because of a high number of database queries. For mobile workloads, this heavily depends on the specific application.
- ▶ Scalability requirements (for example, scale-up for the database, and scale-out for the other components)

Regardless of what operating system and processor type the application components are running on, zEnterprise integrates them in one system managed as a whole. All the infrastructure components of the servers and the corresponding network resources assigned to the mobile solution can be managed end-to-end by using the Unified Resource Manager.

In this hybrid context, the co-location of server images in one system helps to reduce network hops and a virtualized everything architecture ensures an extremely efficient resource usage for up to thousands of virtual servers (see also “A virtualized everything environment to achieve optimum resource utilization” on page 23). Based on this virtualization, the zEnterprise concept offers good scalability across all supported server types because of scale-up and scale-out possibilities. Because mobile solutions typically have strong scalability requirements, this feature is key. Furthermore, an integrated dynamic adjustment of resources for virtualized server images helps to improve mobile application responsiveness.

Finally, zEnterprise offers a good ecosystem of high availability and disaster recovery, which helps to provide a reliable service for mobile solutions.

Consolidation on zEnterprise value for mobile solutions: The zEnterprise system is a well suited consolidation environment for mobile applications because it supports various operating systems and processor types in one system. This environment enables a best-fit approach for the various mobile application components. For all application components, regardless of which zEnterprise part they are running on, high scalability can be offered.

A virtualized everything environment to achieve optimum resource utilization

Based on decades of virtualization experience, zEnterprise offers a solid foundation for consolidation. The zEnterprise is completely virtualized: all server instances are provided based on a virtualization technology. Depending on the destination operating system and processor type, various virtualization technologies are used underneath. However, all of them are managed by using a common user interface and management tool called the Unified Resource Manager. This common user interface reduces silo-based thinking.

The zEnterprise server virtualization is extremely efficient and causes only a small overhead while offering high scalability. For example, with the superior scale-up IBM z/VM® server virtualization technology, thousands of virtual servers can run on one single system. In addition, the zEnterprise system can have up to four books with up to 80 configurable processors, offering significant capacity on which to scale up, and consolidate thousands of virtual servers on a single footprint. Alternatively, efficient scale-out options can be offered based on the POWER7 and x86 blades. These scalability options, combined with a small virtualization overhead portion, makes zEnterprise a perfect platform for consolidation.

Besides server virtualization, the system internal network is also virtualized based on VLANs. This virtualization layer is also managed by the Unified Resource Manager and helps to provide or change network resources faster than in organizations where the network is managed independently. Furthermore, no physical network cables have to be added or repatched because the network is completely virtualized.

With the help of IBM DS8800 technology, physical volumes are also virtualized and represented as logical volumes to the zEnterprise system. Both System z and open storage can be used with one physical footprint.

Consolidation on zEnterprise value for virtualized applications: Overall, zEnterprise provides a solid foundation for consolidation because the whole environment is completely virtualized and offers high scalability. As a result, all components have a more effective utilization rate, which lowers physical equipment, energy, and floor space requirements.

Automated server provisioning in a hybrid environment

To support a private cloud, which provides server provisioning in a hybrid consolidation environment, the virtualized-everything architecture of zEnterprise is a perfect base. Based on this virtualization architecture, server images can be easily created with the Unified Resource Manager for various operating systems and processor types.

With the help of additional IBM Tivoli tools, a self-service portal for automatic service provisioning can be added, and also metering and accounting capabilities, such as for Linux on System z. IBM intends to extend the Tivoli Integrated Service Management for System portfolio to take advantage of the zEnterprise ensemble monitoring and management capabilities that are provided by the Unified Resource Manager APIs².

² Statement of Direction:

<http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&etternum=ENUS111-136>

For further information about these zEnterprise cloud capabilities, see *Cloud Computing and the Value of zEnterprise*, REDP-4763.

Consolidation on zEnterprise value for provisioning: zEnterprise provides a solid base for private cloud services, which allows companies to improve time-to-market through faster IT provisioning.

Integration within a box and integrated appliance for business process management solutions

To address the increasing CPU demands of BPM solutions, zEnterprise offers a special integrated appliance called IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise (XI50z). This appliance can act in various roles such as a security gateway, enterprise service bus (ESB), or simply for parsing XML messages, which is often needed in BPM environments. With the help of the DataPower XI50z, CPU consumption for XML parsing on the servers can be reduced drastically, so these can be used for other tasks. Overall, this way can increase response times and service levels.

When used in an ESB role, the DataPower XI50z can be used for bridging, routing and transforming messages between nodes within the BPM solution. Because the DataPower XI50z is “on board” the zEnterprise server and integrated with the Unified Resource Manager and system internal network, no additional hops are in the network, and security is guaranteed.

Consolidation on zEnterprise value for business process management solutions: zEnterprise can reduce CPU consumption of BPM solutions with the help of the integrated DataPower XI50z appliance. Also, when used in an Enterprise Service Bus role, the XI50z provides sophisticated integration capabilities, hosted and managed within the zEnterprise.

Using a higher level of security in a server consolidation environment

Despite all the advantages, in a “share everything” environment there is usually one serious concern being raised: How much isolation can be guaranteed between virtual server images, networks and storage partitions? The advantages from consolidation can be maximized only if virtual servers, networks, and storage can coexist in one physical environment without security concerns. Weak security offered by a server, network, or storage server can lead again to hosting applications on separate physical resources, causing sprawl.

Therefore, probably the most important attribute of a consolidation system is to be highly secure, meaning guaranteeing full isolation between logical images as designed. zEnterprise offers this high level of server image isolation. For example, the virtualization technology logical partitioning (LPAR) is EAL5 certified and z/VM EAL4+³. This high isolation level ensures minimum security exposure from a consolidation point of view. From a network perspective, the zEnterprise supports the definition of VLANs, that can be easily managed through the Unified Resource Manager.

³ For more information, see: <http://www.ibm.com/systems/z/advantages/security/>

It is crucial that systems management access is handled carefully. Therefore, zEnterprise offers fine granular access to various resources of the Unified Resource Manager, based on roles. This reduces consolidation risks because system administrators have access only to their specific required resources. The Unified Resource Manager has its own highly secured network connections to the various hypervisors in the system.

Another aspect to remember is that a consolidation environment must be highly available in general to reduce impact on a large number of servers. For years, zEnterprise is known for its stability and high availability and disaster recovery functions. This availability is a solid base to count on for a consolidation environment because it helps to reduce business impact and fulfill potential governmental regulations.

Consolidation on zEnterprise value for security: Overall, security has always been a key design point and differentiator of the IBM System z platform. zEnterprise continues these principles and expands these to the POWER7 and x86 blades that can be attached to the system. With zEnterprise, there should not be any security concerns to share many logical servers and networks on the same footprint, even if they have vastly different security levels (trusted versus untrusted).

Archived



Server consolidation approach

A server consolidation project can be a big undertaking for your business. However, a methodical approach to your IT transformation project can simplify the process and ensure a successful outcome with the desired business results. In this chapter, we provide guidance for the steps to be taken for a successful server consolidation project. We outline the steps for the pre-consolidation, consolidation and post-consolidation phases of the project.

The flow chart in Figure 7 on page 28 shows a high level overview of a possible approach for a server consolidation project.

Note: For less complex projects, several steps in the flow chart may be combined in a single step.

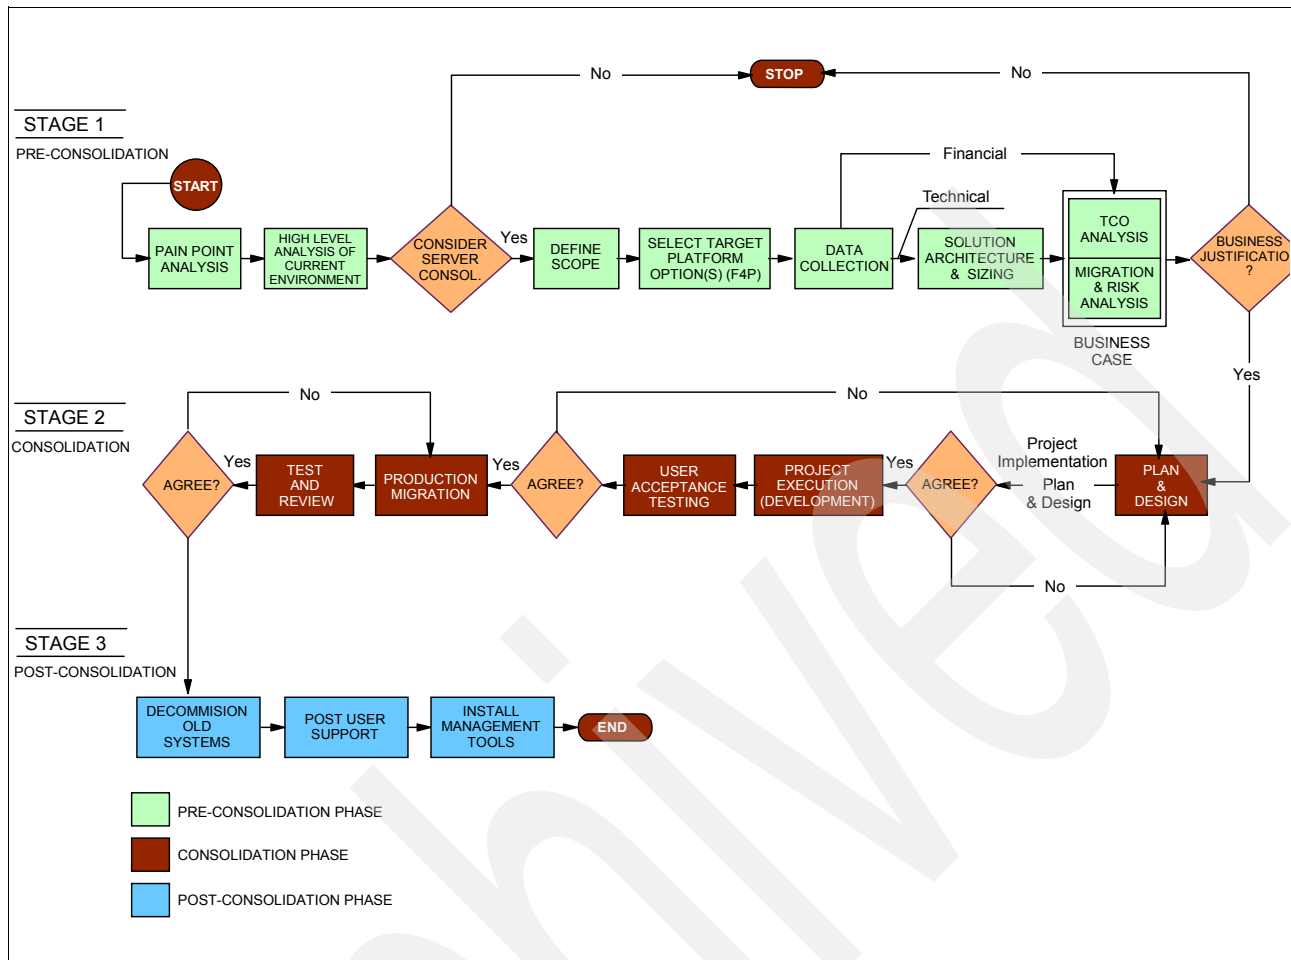


Figure 7 Server consolidation end-to-end process

Phase I: Pre-consolidation

In this phase, you analyze your as-is environment to determine whether a server consolidation is beneficial *and* feasible from the perspective of business optimization, risk, and return on investment (ROI).

At the end of this phase, you should attain the following accomplishments:

- Have clarity about the challenges of your current IT environment and how it relates to your business needs
- Understand the potential impact of a server consolidation IT transformation project to help optimize your business processes and realize your desired to-be state.
- Decide whether there is enough business justification to engage in a server consolidation project.

This phase has seven steps, which are described next.

Step 1: Analyze pain points

As we mentioned earlier in this IBM Redguide™ publication, businesses today want an agile IT environment that is flexible enough to meet their changing business needs while continuing to meet their *service level objectives* (SLOs) and *service level agreements* (SLAs) with their clients.

This initial step involves building an understanding of any IT inhibitors to the company's business now and in the future. It involves understanding your current business optimization objectives or challenges such as simplifying the IT environment, reducing costs and improving service levels. You should also understand how you are performing today with regard to these stated objectives and the positive impact to the business if the desired state is reached.

Note: In this step, you answer the question “Why should I consider consolidation?”

Step 2: Review the current environment

In this step, your objective is to have a clear picture of the existing IT environment and to relate it to the desired objectives of the previous step, with a view of understanding at a high level whether server consolidation will help realize the to-be state. You review the current server inventory, the applications and primary functions, application affinities and dependencies, business constraints, non-functional requirements, SLOs, SLAs, and overall costs.

You should obtain an overview of the server inventory for information about the hardware, operating systems, utilization, and the server primary workload and function such as production, test, and development. For the applications, you want to understand the primary functions and any platform affinities that can provide you an insight to the potential complexity of a consolidation. For example custom-coded applications that are written for a specific hardware might be present a higher level of complexity compared to file and print server applications. Application affinities might require that certain applications are closely located to avoid potential response time issues, and dependencies might require that certain applications must be consolidated together. You should also understand the business constraints in the existing environment that might impede a consolidation opportunity, for example business ownership, location, security, and network.

Step 3: Define scope

If you have determined that server consolidation should be considered in the environment, in this step you determine how big an enterprise IT transformation your business is willing to undergo in a potential server consolidation opportunity. This step both determines the level of effort, complexity, risk, and costs that will be involved, and also determines how big are the enterprise transformation and benefits you can realize from the server consolidation opportunity.

For example, in a simple consolidation project, at a lower level of complexity and transformation, you might only be considering virtual servers on the same hardware platform and operating system. At an average level of complexity, you might be considering consolidating on the same operating system but on separate hardware platforms. At a higher

level of complexity, you might consider a full migration with various operating systems and hardware platforms, and you might also need to recompile code.

Note: In this step, you answer the question “How *complex* should the consolidation project be?”

Step 4: Select target platform alternatives

In this step, you narrow the target hardware platforms to base your business case for server consolidation. In addition to functional requirements, you must consider the client’s non-functional requirements and workload characteristics as you select the target platforms as these have a direct bearing on predicting the potential of a successful server consolidation.

There are several non-functional requirements to consider such as scalability, response time, high availability, continuous availability, disaster recovery, security, skills, workload fit, data center constraints, manageability, standards, strategy, data proximity, data replication, governance, and efficiency. Depending on the relative importance of these factors in your business, certain platforms might be a better fit for the applications in scope.

The workload characteristics for the applications under consideration might have some platform affinities. For example, if the application has a near steady-state usage variability, it might not benefit from a platform that offers superior virtualization capabilities that can fit many virtual servers per core. It might be better suited on a dedicated system.

Note: In this step, you must avoid the “one-size-fits-all” approach.

Figure 8 on page 30 shows a sample bar chart of the possible output of a *fit-for-purpose* assessment. This example shows how the platforms under consideration rank based on the non-functional requirements of a hypothetical client’s environment.




Figure 8 *Fit-for-purpose: Ranking platforms by client’s non-functional requirements*

Step 5: Collect data from your IT environment

In this step, you provide two types of data: technical and financial. This data is necessary for the initial solution design and is especially critical for building the business case.

The technical data can include your server inventory information such as hardware, processor speed, number of cores, utilization, predominant workload, number of servers, and target platform choices.

The financial data can include software maintenance and service, hardware maintenance, floor space costs, energy costs, staff costs, depreciation, software purchase costs, hardware purchase costs, and transition costs.

Step 6: Design solution architecture and size requirements for to-be state

In this step, you design the future state with possible solution alternatives for the desired IT environment based on the technical data that has been provided. Review the data to identify groups of servers that may provide opportunity areas for consolidation by considering the current capacity, workloads, utilization, and operational costs for those group of servers. The sizing process typically involves analyzing the equivalent processor requirements on the target platforms. The output of this process is a solution architecture that recommends servers to be consolidated and the related technical requirements of the to-be state. It may also include an identification of groups of servers that are not cost effective to consolidate.

Note: In this step, several iterations of the sizing might have to be done before a solution architecture can be decided.

Also necessary might be to use estimates assumptions or industry averages (“rules-of-thumb”) in this step.

Step 7: Construct your business case

In this step, the outputs from the solution sizing and the financial data provided are used to construct a business case. The business case articulates the operational savings that can be achieved through a move to the proposed consolidated environment, balanced against the the one time costs of implementing the consolidation. It consists of two key elements:

- ▶ Total cost of ownership (TCO) analysis
- ▶ Migration and risk analysis

TCO analysis

The TCO analysis compares the costs of running on the current environment versus running on the target consolidated environment. It shows the expected return on investment (ROI) of the initial costs of the new consolidated environment, from the operational savings that can be realized.

Figure 9 on page 32 shows a sample business case output for a consolidation opportunity. It shows a cost comparison of running on the current environment versus running on a proposed consolidated environment. See “Business case 2: Linux on non-virtualized x86 servers compared to IBM zEnterprise BladeCenter Extension (zBX)” on page 39.

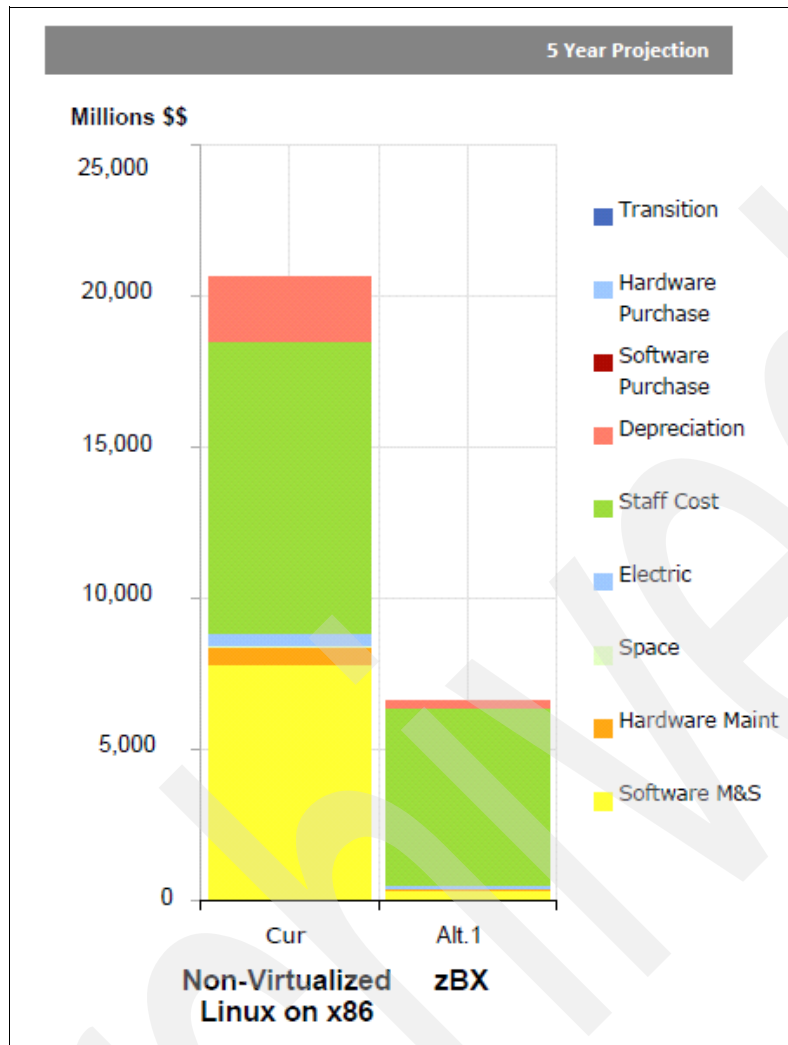


Figure 9 Sample output for a server consolidation business case

Migration and risk analysis

At the end of this step, you analyze the risk of implementing the server consolidation project to deliver the desired IT transformation results and make a judgement about whether the estimated operational cost savings will be worth taking the risk.

Note: The risk is primarily based on non-cost factors. Typically you will look at organizational and technical factors.

Typically, you answer the following questions:

- ▶ What will the transformation entail and is it feasible?
- ▶ What are the potential key issues of the migration and can they be mitigated?
- ▶ What is the migration cost of undertaking this project?
- ▶ How long will the entire project be and can I afford it?
- ▶ How much down time will be required and can I afford it?
- ▶ How many full time equivalent (FTE) hours will I need to invest for planning, reviews, approving and acceptance testing, and can I make this investment?

Based on the results of this assessment and the insight provided, you should be able to decide whether there is a business justification within an acceptable risk for the business. If you decide that there is a business justification, proceed to Phase II, the consolidation project.

Phase II: Consolidation

In this phase, you fine-tune the design of your to-be environment, implement your consolidation project, test and complete the migration of your production environment.

Aim to accomplish the following goals during this phase:

- ▶ Create an actionable design and implementation plan for your consolidation project.
- ▶ Execute your plan and test the new solution.
- ▶ At the end of this phase, roll out the migration of your production environment.

This phase has four steps, which are described next.

Step 1: Plan and design

In this step, you are involved in two processes: implementation planning and design. At the end of this step you produce an actionable IT Transformation Project Plan. You will review the plan with decision makers to get an approval before you can proceed with the project.

In the design phase, you might need to verify the assessment of the as-is environment made in “Step 2: Review the current environment” on page 29, of your pre-consolidation phase. You might have to substitute some of the assumptions and estimates made in the pre-consolidation phase with actual data from the client’s environment. Finally, you will fine-tune the solution design from “Step 6: Design solution architecture and size requirements for to-be state” on page 31, of the pre-consolidation phase to ensure that it will fit in seamlessly with the client’s environment.

Based on your output from the design phase, you work on your project plan in the implementation planning phase. At a minimum, your plan should have the following elements:

- ▶ Schedule of tasks to be completed listed chronologically for a phased implementation of the project
- ▶ For each task, identify the owners with roles and responsibilities
- ▶ Time required for each activity
- ▶ Regular checkpoints to review progress, get consensus or modify design and project plan as is necessary
- ▶ Time buffer to allow for unplanned events

Step 2: Development and user acceptance testing

In this step, you are involved in the basics of the consolidation project. This step might involve activities regarding migration of servers and operating systems, updating IT addresses, changing application codes, installing and integration of new systems into the existing environment and network, data and application migration, and so on.

You might typically start by deploying the solution on a small scale and then have users test and approve the solution before you start a large scale deployment or migration. Depending

on the complexity of the project, this step might involve activities such as pilots, proof of concepts (POCs), and bench mark studies.

Before proceeding to the next step: You might have to return to the “drawing table” to further fine-tune your design and plan based on the results of this step, before you can proceed to the next step.

Depending on the complexity of the consolidation project, you might have to repeat this step for several phases of the project.

Step 3: Production migration

If there is a successful user acceptance test result in the preceding step, you receive approval to continue with a full deployment of production. The key difference in this step and the previous step is that the deployment is on a much larger scale.

Step 4: Test and review

This final step of the consolidation phase is included in the preceding step. As each piece of production is rolled out, test, review, and measure it against the expected results before moving to the next phase of the migration. This step can ensure that bottlenecks are identified earlier in the process and managed before they can stall the project.

When the consolidated environment is running and all tests and reviews have been done and approved, you have completed the consolidation phase of the project and you have commissioned your new consolidated environment.

Phase III: Post-consolidation

In this phase, you are engaged in activities that lay a strong foundation for your new environment, running as planned and delivering the desired business results.

Decommission old systems

Old systems that are no longer in use must be decommissioned, disconnected from the network, and disposed of, as is necessary.

Post user support

A range of services are available to provide support for the on-going operation on the new consolidated environment for example, technical support of new users and IT staff, warranty and maintenance, managed services, and so on.

Install management tools

With the new consolidated environment running, you might have to install certain management tools for performance monitoring and reporting.



Case studies

In this chapter, we discuss two real-world server consolidation case studies with targets of zEnterprise z196 and zBX. These business cases were done by the IBM Scorpion team.

Business case 1: ISV database servers on non-virtualized x86 compared to Linux servers on IBM zEnterprise z196

This case study concerns a server consolidation from non-virtualized distributed x86 servers to a virtualized environment on zEnterprise using Linux for System z and Integrated Facility for Linux (IFL) processors.

Business challenge

The client expressed interest in a fact-based infrastructure assessment study to gauge the level of efficiency of the client's current state and to explore potential opportunities for improvement.

Description of the current environment

Of several server optimization business cases developed for this client, the current state island we describe, consisted of approximately 375 database instances deployed on 57 physical non-virtualized x86 servers (predominantly 2-socket) running the Linux operating system. The total core count/socket count was 240/109. The business case consists of the following server inventory:

- ▶ 25 Linux on Intel x86 2.93GHz servers
- ▶ 32 Linux on Intel x86 3.0GHz servers

Approach

All technical and financial data was real and furnished by the client through IBM Scorpion team templates. The templates were specifically designed to capture all relevant data points essential to the generation of business cases that accurately and with integrity demonstrated the realm of possibility if the optimization recommendations were embarked on.

Tried and tested tooling in conjunction with the IBM Scorpion team consultative methods were used to capture all-inclusively the current state and to develop the target state alternatives.

Overview of solution architecture

Two variations of zEnterprise (with Linux) environments were recommended as follows:

- ▶ 10 IFLs on 1 Book zEnterprise 196 (M15) Linux Server ("AltCase1")
- ▶ 13 IFLs on 2 Book zEnterprise 196 (M32) Linux Server ("AltCase2")

The key difference for these two alternative solutions is additional scalability and high availability provisioning provided by having an additional book on the second alternative zEnterprise solution.

Results of the study and conclusions

The financial merits of the business case is clearly stated in Figure 10 on page 37 and Figure 11 on page 38.

x86 Linux ISV Database Server				
Sizing	5 Year Projection	Current	AltCase2	AltCase1
	Server Type		z196-IFL	z196-IFL
	Total Cores	240.00	13	10
	Used Cores	240.00	13	10
	Total Sockets	109.00	13	10
	#Logical Servers	57.00	57	57
	#Physical Servers	57.00	13.00	10.00
	Ave.Log.Srv RIP	4,486.0	1,548.3	1,191.0
	Total RIP Capacity	255,701.1	88,252.6	67,886.6
	Total RIP Workload	56,424.2	56,424.2	56,424.2
	Ave %Utilization	22.07%	63.93%	83.12%
			2 book (M32)	1 book (M15)
Annual Operating Costs (AOC)				
	Staff Cost Code			
	Software Cost Code			
	Software Cost /CPU	15,154	27,412	27,412
	Software Cost /Lsrv	0	0	0
	Software Cost /Psrv	2,544	0	0
	Software M&S	1,796,804	356,356	274,120
	Hardware Maint	62,700	0	0
	Space	269,529	15,371	15,286
	Electric	72,996	29,665	28,607
	Staff Cost	228,523	205,671	205,671
	Depreciation	12,000	0	0
	Total AOC	2,442,552	607,063	523,684
	est.potential saving /yr		1,835,489	1,918,868
One Time Costs (OTC)				
	ELS Solution Purchase		1,315,501	1,027,312
	Hardware Purchase			
	Total OTC	0	1,315,501	1,027,312
	Investment incl.Write-Off		0	0
			1,315,501	1,027,312
5 Year Projection				
	OTC + 5x AOC	12,212,760	4,350,816	3,645,732
	5 yr saving		7,861,944	8,567,028
	% age AOC Reduction		75%	79%

Figure 10 Detailed cost comparison of current state and alternate solutions for business case 1




Figure 11 Five-year cost comparison for business case 1

The resulting benefits of moving from the distributed non-virtualized x86 servers to Linux servers on a z196 according to solution alternative 1 ("AltCase1") are outlined as follows:

- ▶ Potential savings of \$8.5M over five years (79% reduction in cost)
- ▶ Annual operating costs reduced from \$2.4M to \$523K per year (79% per year)
- ▶ Total physical servers reduced from 57 x86 servers to one z196 server (98% reduction)
- ▶ Annual power and floor space costs reduced from \$341K to \$43K (87% reduction)

Savings: Moving to Linux servers on zEnterprise could save this client \$8.5M over five years.

Other frequently stated Fit-for-Purpose benefits for ISV database servers in a Linux on System z environment are as follows:

- ▶ Leverage classic strengths of zEnterprise
 - High availability
 - High I/O bandwidth capabilities
 - Flexibility to run disparate workloads concurrently
 - Excellent disaster recovery capabilities
 - Security
- ▶ Shortening end-to-end path length for applications
 - Collocation of applications
 - Consolidation of applications from distributed servers

- Reduction in network traffic
- Simplification of support model
- ▶ Consolidation effect
 - Reduced power consumption
 - Reduced software costs
 - Reduced people costs
 - Reduced floor space costs
 - Scalability through full virtualization

Business case 2: Linux on non-virtualized x86 servers compared to IBM zEnterprise BladeCenter Extension (zBX)

This case study concerns a server consolidation from non-virtualized distributed x86 servers to a virtualized environment on zEnterprise using blades in the zEnterprise BladeCenter Extension (zBX).

Business challenge

This client had a large, complex IT infrastructure with plans in place to simplify using Linux, as many clients do. Budget and energy costs were a major concern with virtualization projects well underway on all platforms to provide relief.

Description of current environment

Linux was already in production and running many workloads on all platforms but there was no clear differentiation between the platforms. Management required an accurate understanding of current IT costs and strengths so they could better access new ideas and technologies.

Approach

Fit-for-purpose methods coupled with cost-of-operation comparisons were needed to give direction to further IT simplification efforts and to define an optimized future state with lower costs and better flexibility.

This client had a mature environment and had done a number of platform comparisons in-house. The approach used for this study had to be at least as comprehensive to provide credible results for the client. Therefore, the team chose the IBM STG Lab Services IT Optimization (Scorpion) Study approach because of the depth of its analysis.

Six key characteristics of the study led to a successful outcome:

- ▶ First, the team evaluated “platforms” within the environment, including all servers of all types to simplify the process. One of the most difficult things to get right in an analysis of this type is an exact match between the technology and the associated costs. The easiest way to do this is *not to limit* the technology scope to a few machines or a single application but to expand it to match all technology in identifiable categories of the IT budget. This way also prevents the zEnterprise solution from being compared to the “best-of-breed” subset of the environment that is not reflective of the whole environment.

- ▶ The second characteristic was to focus on a representative application and include all the pieces for solutioning. By definition, a critical application will usually require specialized or dedicated infrastructure. Each of these components and their associated costs were included in the comparison.
- ▶ Third, the team considered target utilizations that were realistic for the environment. Our experience is that the most important aspect of performance analysis within a cost-of-ownership is not which vendor claim or bench mark is used as a base but rather the following aspects:
 - What system utilizations are “normal” in your current environment?
 - What is a reasonable expectation for the future?
- ▶ Fourth, we considered the labor impact across all of IT, not only direct infrastructure support. High Full Time Equivalent (FTE) ratios in the Information Technology Infrastructure Library (ITIL) structured “Run” group have been an IT management target for years with improvements based on *either* of two strategies:
 - Real efficiency improvements
 - Pushing work onto other parts of the organization

The extent to which the second strategy is used differs by platform for a variety of reasons but the result is always the same. Any cost of ownership analysis that limits labor calculations to IT infrastructure support headcount will likely miss major portions of the real supports costs.

- ▶ Fifth, we used fit-for-purpose to identify key non-functional requirements and assigned costs to those architectures that were lacking. The mainframe had superior quality of service in a variety of areas and those supported the favorable cost comparison. The Unified Resource Manager extended those qualities of service into the distributed infrastructure in a way that other solutions could not.
- ▶ Finally, we emphasized incremental costs and growth, and de-emphasized costs and technology refresh for all platforms. This allowed us to compare the impact of System z special pricing and the superior price/performance of the newest technology of z196, and the zBX.

Overview of solution architecture

The solution is presented in multiple business cases. The case included here includes the zBX portion of the solution only. The zBX is a single frame with a partially populated BladeCenter and IBM PS701 Blades configured in Automate mode.

The zEnterprise solution was defined in multiple cases to allow different workload characteristics and fit-for-purpose categories to be considered. The solution was not compared to alternative non-IBM solutions, or to alternative resident solutions.

Results of the study and conclusions

The resulting business case and benefits presented to the client are summarized as follows:

- ▶ Potential savings of \$14M over the next five years (total costs are 65% lower per year)
- ▶ Reduced annual operating costs by \$775K per year (40% decrease) for total savings of \$3.8M
- ▶ Total software savings of \$7.5M because of 95% annual decrease in maintenance and support

- ▶ Reduced physical server count from 200 or more to a single zBX frame; reduced cores by a factor of 10
- ▶ Facilities savings of over \$300K as a result of consolidation to more energy efficient zEnterprise servers

Important: Consolidating onto the IBM zEnterprise zBX could save the client \$14M over the next five years.

Figure 12 on page 41 shows a graphical representation of the cost comparison of the current state and the proposed future state running on zBX.




Figure 12 5-year cost comparison for business case 2

Overall, the client was pleased with the comprehensive analysis but did challenge the results. IBM was able to detail the source of all the data, and also the method used to project the estimated savings and overcome the client's bias. Additionally, IBM was able to demonstrate the practical value of fit-for-purpose concepts and their role in cost analysis.

Lessons learned

Making the investment of time in doing a more comprehensive analysis of the client's environment and to-be state is necessary to provide the insight necessary to make a business case for consolidation and ultimately the decision to consolidate.

IBM services and offerings

IBM has a rich portfolio of offerings to assist you in your consolidation projects. IBM also has many experts with deep technical skills and years of server consolidation experience for all levels of complexity. These experts can guide you to avoid pitfalls and manage bottlenecks that can stall your consolidation project.

Table 1 summarizes IBM offerings, and the services that you can engage for support in the pre-consolidation phase of your project.

Table 1 Pre-consolidation phase: IBM server consolidation services and offerings

Services and offerings	Description	Time frame for a typical engagement										
			Pain point analysis	Current state analysis	Define scope	Select targets	Data collection	Solution architecture	Sizing	TCO analysis	Migration and risk analysis	
Alinean	System consolidation evaluation tool to provide a preliminary high level analysis of the estimated potential savings by consolidating your distributed UNIX and x86 applications to a current Linux on IBM System z solution. ^a Little actual client data is required.	< 1 day	X	X		X				X		
Fit-for-purpose workshop	Workshops to determine best fit platform based on client's non-functional requirements and workload characteristics. Client participation is required.	1 day		X	X	X	X					
Linux on System z Server Consolidation (SCON) analysis	Sizing to project the processor requirement on the target System z environment in a consolidation.	1 week					X	X	X			

Services and offerings	Description	Time frame for a typical engagement	Pain point analysis	Current state analysis	Define scope	Select targets	Data collection	Solution architecture	Sizing	TCO analysis	Migration and risk analysis
RACEv study	Server virtualization analysis and cost tool offering to analyze a set of subject servers to generate a cost and value analysis for targets. Estimates may be used in addition to client data. Might not involve on-site data collection.	1 - 10 days		X	X	X	X		X	X	
Eagle study	TCO study comparing costs of operating platforms for business applications over a five-year period prior to and after consolidation. Typically involves on-site data collection and consultation with the client. However, some estimates may be used as inputs.	4 weeks	X	X	X	X	X	X	X	X	X
Scorpion study	Server agnostic cost and value analysis utilizing mostly technical and financial inputs from your environment to build a comprehensive business case for IT optimization (virtualization, consolidation, energy consumption reduction). Typically involves on-site data collection and consultation with client.	2 - 4 weeks 6 - 8 weeks	X	X	X	X	X	X	X	X	X
IBM Services	World-leading portfolio of assessment, infrastructure, business consulting, systems integration, and outsourcing services to provide support for every phase of your end-to-end IT transformation consolidation project ^b .	Varies	X	X	X	X	X	X	X	X	X

a. <https://roianalyst.alinean.com/stgz/>

b. <http://www.ibm.com/services/>

Table 2 on page 45 lists the IBM service that you can engage for support in the consolidation phase of your project.

Table 2 Consolidation phase: IBM service

Service	Descriptions	Time frame for a typical engagement	Plan and design	Project execution (dev)	User acceptance testing	Production migration	Test and review
IBM Services	World-leading portfolio of assessment, infrastructure, business consulting, systems integration, and outsourcing services to provide support for every phase of your end-to-end IT transformation consolidation project.	Varies	X	X	X	X	X

Table 3 lists the IBM service that you can engage for support in the post-consolidation phase of your project.

Table 3 Post-consolidation phase: IBM service

Service	Descriptions	Time frame for a typical engagement	Decommission old system	Post user support	Install management tools
IBM Services	World-leading portfolio of assessment, infrastructure, business consulting, systems integration, and outsourcing services to provide support for every phase of your end-to-end IT transformation consolidation project. Services to optimize utilization of your data center solutions and for delivery of new technologies ^a .	Varies	X	X	X

a. <http://www.ibm.com/systems/services/labservices/>

Note: The time frame for a typical engagement listed can vary significantly depending on the complexity of the project and the number of iterations required.

Archived



Summary

Server consolidation has definitely been proven as a way to reduce complexity and lower cost in the IT infrastructure. In the world of server consolidation, “bigger is better” is a true statement and the more servers can be consolidated into a larger footprint, the better. At least, this is true under the condition that the server is able to virtualize at all levels, share resources effectively, and provide adequate isolation between resources.

IBM is in the lead by introducing zEnterprise, a server that is truly a “data center in a box.” The system provides seven operating systems¹ and three processor types², managed by a single management tool and all interconnected by a high-speed and secure internal network. You can consolidate heterogeneous workloads currently running on various types of stand-alone servers onto one single zEnterprise server. By consolidating to zEnterprise the path length between application components can be shortened and networking infrastructure can be simplified by eliminating routers and switches.

If you are interested to learn what a consolidation to zEnterprise could mean in your specific situation, IBM can help you with a range of offerings.

Other resources for more information

See the following resources:

- ▶ 2011 IBM Global CIO Study:
<http://www.ibm.com/ciostudy>
- ▶ Experton Group; The zEnterprise – A Disruptive Technology:
<http://www.experton-group.com/home/archive/news/article/the-zenterprise-a-disruptive-technology.html>
- ▶ *Cloud Computing and the Value of zEnterprise*, REDP-4763

¹ z/OS, Linux on both System z and Intel processors, IBM z/TPF, z/VM, IBM z/VSE™, AIX; and Statement of Direction: Microsoft Windows

² System z, POWER7, and Intel

The team who wrote this guide

This guide was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO).

Alex Louwe Kooijmans is a project leader with ITSO in Poughkeepsie, NY, and specializes in service-oriented architecture (SOA) technology and solutions using System z. He also specializes in application modernization and transformation on z/OS, both from an architectural and tooling perspective. Previously, he worked as a Client IT Architect in the financial services sector with IBM in the Netherlands, advising financial services companies about IT issues, such as software, hardware, and on-demand strategies. Alex has also worked at the Technical Marketing Competence Center for IBM zSeries® and Linux in Boeblingen, Germany, providing support to customers starting up with Java and WebSphere on System z. From 1997 to 2000, Alex completed a previous assignment with the ITSO, managing various IBM Redbooks® projects and delivering workshops around the world in the area of WebSphere, Java, and e-business technology using System z. Prior to 1997, Alex held a variety of positions in application design and development, product support, and project management, mostly in relation to the IBM mainframe. Alex has 23 years of IT experience and has been the lead author of many IBM Redbooks and IBM Redpaper™ publications.

Ogechi Nduka is a Client Technical Specialist in the Global Techline organization at IBM in the US. She holds a Master's Degree in Management Science and Engineering from Stanford University and a Bachelor's Degree in Electrical/Electronic Engineering from the University of Port Harcourt. Her areas of expertise include System z Performance Analysis and Capacity Planning, Linux on System z Server Consolidation Sizing and Total Cost of Ownership Assessment, and IBM System z Energy Efficiency Analysis. She has written extensively about System z Energy Use and Efficiency Estimation and has a patent filed for her work on System z Energy Analysis.

Liv Samuels has 32 progressive years of mainframe technical experience. Tenure includes time with outsourcing organizations in the UK and USA whereby engaged in numerous projects covering a broad variety of workloads spanning various business sectors. Positions have included Operations Supervisor, Operations Production Support Supervisor, and mainframe Systems Programming Project Leader/Specialist. Primary focus areas included all the technical phases and disciplines in risk analysis and contingency planning related to mainframe operating system support, workload migrations, consolidations, and rationalizations. Currently Liv provides leadership as a member of the IBM STG Lab Services and Training Scorpion team engaging in a variety of consultative-focused brand-independent client IT infrastructure assessments and optimization solution studies (range of engagements complexity are workshops to rapid-assessments to in-depth).

John Schlosser is currently a Senior Managing Consultant for the Scorpion (IT Rationalization and Optimization) practice within IBM Systems and Technology Group - Lab Services and Training. He is a founding member of the group, which was started in 1999. He has developed and modified many of the methodologies that the team uses for IT infrastructure cost analysis. John works with other team members to execute customer CIO-sponsored IT infrastructure and cost of ownership studies. These studies focus on IT systems rationalization, optimization, and virtualization along with cross-platform cost of ownership analysis for all server and storage platforms. He has personally lead or assisted in over 91 studies and has authored or contributed to several white papers on the subjects of cost of ownership, platform selection, and workload consolidation. John regularly presents to customer IT executives, at customer conferences, and to IBM executives. He has presented these methods at Guide and Share and is published on these topics in industry periodicals.

Christian Strauer is an IT Architect from IBM Germany. He works as a Client Technical Architect in the Financial Services and Financial Institutions sector. His areas of expertise are solution design, infrastructure platform selection, IBM System z, and disk storage. He has worked for IBM since 2002 when he started his studies at the University of Cooperative Education in Mannheim. Before his current job role, Christian worked as an IT Specialist in the IBM System z Field Technical Sales Support team Germany. He holds a diploma of Applied Computer Science and has authored various IBM Redbooks publications.

Thanks to the following people for their contributions to this project:

Kurt Acker
New Workload Field Technical Sales Specialist for System z virtualization

Ellen Carbarnes
System z Product Marketing Manager, IBM Systems & Technology Group

Frank DeGilio
Distinguished Engineer and CTO Cloud computing, IBM Systems & Technology Group

Robert Green
Cloud Solutions Architect and System z Solutions FTSS

Darrell Hawkins
Manager, IT Optimization Consulting, Systems Lab Services & Training

Sebastien Llaurency
Technical Solution Architect, IT Optimization, IT Service Management, Provisioning, Monitoring

David Petersen
Distinguished Engineer, IBM GDPS® Chief Architect, System z Products

Ella Buslovich, Elsie Ramos, and Diane Sherman
International Technical Support Organization, Poughkeepsie Center

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks


IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US

registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>



The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
BladeCenter®
DataPower®
DB2®
GDPS®
HiperSockets™
IBM®

POWER7™
Redbooks®
Redguide™
Redpaper™
Redbooks (logo) ®
System x®
System z®

Tivoli®
WebSphere®
z/OS®
z/VM®
z/VSE™
zSeries®



The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.