Redbooks
ibm.com/redbooks

Scale up for Linux on IBM Z

Edi Lopes Alves
Eric Everson Mendes Marins

Linux

IBM

Redpaper

International Technical Support Organization

Scale up for Linux on IBM Z

December 2017

REDP-5461-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (December 2017)

This edition applies to IBM LinuxONE and Linux on IBM Z.

This document was created or updated on December 12, 2017.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Bluemix®	IBM LinuxONE Rockhopper™	Redbooks (logo)  ®
DB2®	IBM z®	Storwize®
DirMaint™	IBM Z®	System z®
FlashCopy®	IBM z Systems®	z Systems®
HiperSockets™	IBM z13®	z/Architecture®
IBM®	PR/SM™	z/OS®
IBM Blue™	RACF®	z/VM®
IBM LinuxONE™	Redbooks®	z/VSE®
IBM LinuxONE Emperor™	Redpaper™	z13®

The following terms are trademarks of other companies:

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

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.

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

Preface

This book was written by IBM® IT specialists who have experience implementing IBM Z® solutions, especially Linux on IBM LinuxONE™ (LinuxONE) or IBM Z servers. Therefore, the contents of this book follows the guidelines from Linux and IBM z/VM® regarding LinuxONE and IBM Z installations. The preferred practices described in this book are gathered from the experiences of those specialists in hundreds of projects at IBM and customer environments.

This IBM Redbooks publication provides you with all of the information that you need to decide the best scaling architecture when implementing Linux on IBM Z or LinuxONE. This book has the following goals:

- To inform you about x86 sprawl problems
- To inform you that x86 Vertical Scale out architectures are problematic going forward
- To provide solutions to x86 server sprawl problems
- To inform you about the LinuxONE and IBM Z differentiation for each x86 server sprawl problem
- To provide virtulization and security options for LinuxOne and IBM Z

The scaling up and scaling out architectures enable you to scale the capacity of an existing system to attend a sporadic application demand or an application workload. This gives you some freedom to operate in the environment. However, if this activity is performed without correct planning and the correct architecture choice, it leads to a server sprawl situation where your environment houses more servers than it should based on its current and predicted requirements.

Although scaling out on x86 systems is a common form of scaling because of their popularity, the x86 systems were originally designed as cheap computers. Unfortunately, the scale out on x86 can easily become a problem in terms of total cost of ownership (TCO) when the environment starts to increase in terms of number of physical servers.

The LinuxONE and IBM Z servers solve the sprawl problem caused by the scaling out of x86 servers, and are an excellent choice for cloud, mobile, big data, blockchain, analytics, and other workloads that require a robust and flexible environment.

This publication describes the advantages and disadvantages of the scaling options.

The audience for this publication consists of the following groups:

- Customers, IBM Business Partners, and IBM consultants planning and installing Linux on IBM Z, IBM Z family or x86 platform
- System administrators administering the Linux Systems

For more information about understanding the advantages of server consolidation on LinuxONE over distributed servers, see the following links:

- IBM LinuxONE solutions
- IBM LinuxONE Redbooks

If you are a customer considering LinuxONE and IBM Z family as a platform for your applications (analytics, blockchain, cloud, or other) or a pre-sales person, read those publications.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Edi Lopes Alves is a Senior IT Specialist in Brazil working with IBM z® Systems for the GTS team. She has more than 25 years of experience working as a z/VM and Linux on IBM Z specialist. Edi is IBM L2 IT Specialist certified and has a master's degree in e-Business from ESPM Sao Paulo. She supports z/VM and Linux on Z for an American Express account. She has supported z/VM environment and cloud initiatives for Banco do Brasil and IBM Global Accounts (IGA) for several years as part of IBM Green and IBM Blue™ Harmony projects, and supported z/VM Field Test at the Endicott Lab. Edi has co-authored the IBM Redbooks® publications: *Securing your Cloud: IBM z/VM Security for LinuxONE and IBM Z*, SG24-8353; *IBM Wave for z/VM Installation, Implementation, and Exploitation*, SG24-8192; *Using z/VM v6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039; *Linux on IBM System z: Performance Measurement and Tuning*, SG24-6926, and *Introduction to the New Mainframe: z/VM Basics*, SG24-7316.

Eric Everson Mendes Marins is a Senior IT Specialist for IBM Global Technology and Services in Brazil. He has more than 15 years of experience in configuring Linux on IBM Z. He is currently working as a Linux on IBM Z Architect for IBM Global Account (IGA) in Brazil and implementing several cloud solutions on IBM Z family. He is IBM L3 IT Specialist certified (Thought Leader) and an IBM L2 IT Architect certified (Expert). Also, he holds a degree in Computer Science from Faculdade Ruy Barbosa and a post-graduate degree in Computer Information Systems (Database Management). His areas of expertise include Linux, z/VM, high availability, IP networking, and server management. Eric has co-authored several IBM Redbooks publications, including *Advanced Networking Concepts Applied Using Linux on IBM System z*, SG24-7995, and *Security for Linux on System z*, SG24-7728.

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

Lydia Parziale, Robert Haimowitz
International Technical Support Organization, Poughkeepsie Center

Sandy Bulson
IBM US

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:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form:

  ibm.com/redbooks

- Send your comments in an email:

  redbooks@us.ibm.com

- Mail your comments:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

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

Introduction

This chapter describes the goals of this publication, the differences between the scale-out (horizontal) and scale-up (vertical) architectures, and how the IBM LinuxONE (LinuxONE) and IBM Z platforms can help solve the x86 server sprawl problems.

This publication is a guide and is targeted at non-technical and technical staff who want to perform data center modernization, with server consolidation and virtualization onto LinuxONE and IBM Z.

This chapter contains the following sections:

- The x86 sprawl problems
- Scale-up and scale-out definitions
- LinuxONE and IBM Z versus distributed x86
- Considerations about the LinuxONE and IBM Z platforms

1.1 The x86 sprawl problems

Server sprawl is a growing concern in data centers today. When choosing your next platform or evaluating your current environment, consider the hard effects of server sprawl, and their impact on your IT environment.

Server sprawl typically refers to the number of underutilized computer systems running in a data center. To control this environment, the system administrator has to deal with licensing, support, security, and compliance issues for each system. As the number of physical systems increases, the costs to manage it becomes more expensive and requires more efforts from support staff to keep the environment running and compliant with security standards.

Usually, server sprawl is associated with typical x86 scale-out environments because they had an uncontrolled proliferation in past years, and also because x86 technology was the first choice for many organizations. Figure 1-1 illustrates a typical x86 scale-out environment.

Alternatively, to solve the x86 server sprawl problem, the LinuxOne and IBM Z platforms have become popular to deliver more applications and services, and as an important component to support compute-intensive Business Analytics and big data workloads. Also, these platforms help to consolidate hundreds of underutilized systems onto a single machine, and help your application take advantage of all of the platform benefits.

Figure 1-1 Typical x86 scale-out environment

The main problems of x86 scale-out sprawl include the following hidden costs:

- Low resource usage
- Increased security problems
- Wasted energy consumption
- Increased server management

1.2 Scale-up and scale-out definitions

When a system or an application reaches its maximum capacity, the system administrators have some options to make it scale to improve performance or add capacity for a new application feature. Knowing the strengths and weaknesses of each option enables them to make the best decision to improve an existing system that is expensive to operate, or prone to errors, or difficult to support.

As cloud computing platforms and mobile applications become common across organizations around the world, the importance of having a versatile infrastructure to be more agile and respond faster to changing demand is a key element for success today. Scaling an application to support a new functionality or build a new system is never an easy task, and requires huge effort to ensure that everything is correctly implemented at the opportune time.

In the same way, the Application Owners (AO) require less time to create server instances and add new system resources to help them deploy new applications and workloads faster and more frequently, with no need to wait for weeks. The crucial point is that infrastructure needs to be flexible and enable the AOs to manage scalability where new workloads are needed, and when they are no longer required. Therefore, one of the biggest features of the infrastructure is the ability to scale. There are two ways to accomplish scaling, one is called *vertical scaling* and the other is *horizontal scaling*.

Scale up (or vertical scaling) is enlarging the capabilities of a single system. It is typically the ability to increase the capacity of existing hardware or software by adding resources to the system. Adding processors, memory, or any other resource is an example of scale up. Its important to note that this type of scalability is limited by the size of the server.

Scale out (or horizontal scaling) is adding more hosts to a tier. It is typically the ability to scale larger to deal with a new workload. Usually, the system resources are spread across multiple servers (called *clusters*). Spreading the resources reduces the load on a single system, makes the environment more available, and avoids the “single point of failure” problem. Reliability is improved.

Both methods require planning, and making sure that resources are available. It is also important to ensure that the architecture where applications run can handle the scalability.

New application demands, budgets, smartphones, and other digital devices are driving the definition of scale-up and scale-out architectures today. Some “horizontal” applications, such as data analytics, can take advantage of the scale-out approach to increase performance and provide availability. The scale out is basically a large number of nodes that work together to provide an aggregated performance that cannot be achieved by making a single large node work on its own.

When a system administrator needs to grow the infrastructure from two systems to five or more to support a mission-critical workload, he uses the scale-out approach, which increases availability even if one of the servers crashes. In case of a crash, you usually only lose a minor portion of the virtual infrastructure, as opposed to losing the entire cluster.

The main idea is to aggregate enough computing power to support large operations and complex applications using more than one node. The LinuxONE and IBM Z platforms support both scale up and scale out architectures.

These options are available to be used to scale a system. However, it is important to know the benefits and negatives of scaling out and scaling up architectures, to ensure that you make the best decision where to deploy a new workload or a complex application.

On LinuxONE and IBM Z platforms, these applications can be scaled out or scaled up with superior performance and security using virtual Linux servers.

Table 1-1 provides an overview of the benefits and negatives of scale out architecture.

Table 1-1 Overview of common pros and cons associated with horizontal scaling

Benefits	Negatives
Typically simple (buy and add new node to the cluster)	Requires more power and space for the data center
Flexible (add or delete nodes when needed)	Increases complexity (more nodes to manage)
Easy to perform upgrades on nodes (no application downtime)	Latency issues might happen between nodes
Easier to run fault-tolerance	Additional software licensing
	Requires more cooling
	Generally requires software for clustering to implement high availability
	More intrusion points
	Several points of management

Table 1-2 provides an overview of benefits and negatives of scale up architecture.

Table 1-2 Overview of common pros and cons associated with vertical scaling

Benefits	Negatives
Flexible to resize a virtual server	More expensive than scaling out (specialized servers)
Lower energy consumption (compared with scale out)	Sometimes hard to upgrade a node (application outage is needed)
Minimizes the need for physical nodes	Need reliable hardware
Requires less cooling	
Generally less complex to implement (compared with scale out)	
Fewer intrusion points (tighter security)	
Central point of management	
Increases resource utilization	

1.2.1 Scaling up (vertical)

The vertical scaling is typically used to address the growing demands of an application, where a system administrator can scale an application capacity up or down according to application business needs. The scale-up approaches have been around longer and are more familiar to administrators, generally offering a good feature set and functionality to meet their purpose.

It is commonly used when a server with high load needs more resources or when the application owner needs to run an extra critical workload in a determined period of time.

However, the resources cannot be freely added or it can affect the overall performance of the box. Another crucial point before scaling up an application is related to the platform that the application is running on. With the LinuxONE or IBM Z platform, the CPU utilization can run at over 90% for sustained periods of time without fear of failover. Alternatively, x86 typically runs at 40-60% utilization, which in turn requires more server cores to meet scaling requirements.

Unlike other platforms, LinuxONE and IBM Z were designed to run at very high CPU and I/O utilization because of the balance between CPU performance, I/O capabilities, and access to memory. Of course, running at 100 percent for extended periods of time has never been recommended, because it doesn't leave room for temporary workload spikes, or for any sort of recovery processing that might need to be done in case of a subsystem failure.

Figure 1-2 illustrates the scale-up architecture.

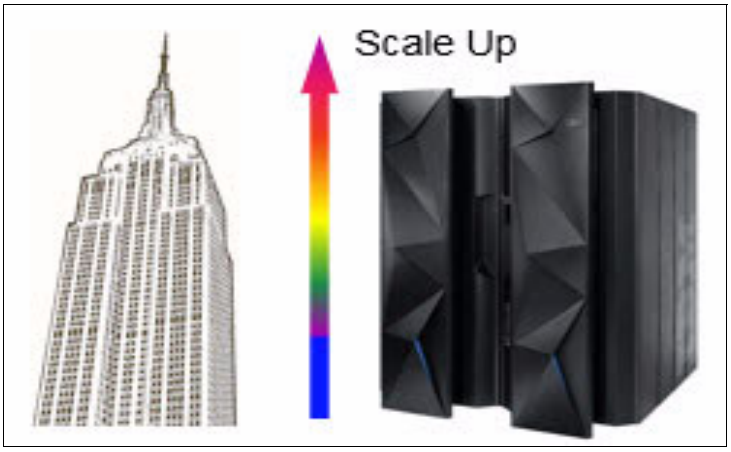

Figure 1-2 Scale-up architecture

1.2.2 Scaling out (horizontal)

This is commonly used in cloud because this architecture addresses some application requirements, such as high availability, flexibility, and scalability. The most important advantage is to provide an easy way to increase capacity dynamically. It means if more resources are required to improve performance and provide high availability (HA), a system administrator can scale out the system by adding more servers into the pool.

It has both advantages and disadvantages. Per instance, adding inexpensive commodity computers to a cluster might seem to be a cost-effective solution at first glance, but it is important for the system administrator who controls the environment to know the licensing costs for those additional servers, the additional operational costs (powering and cooling), and the amount of space they will use in the data center.

Figure 1-3 illustrates the scale-out architecture.

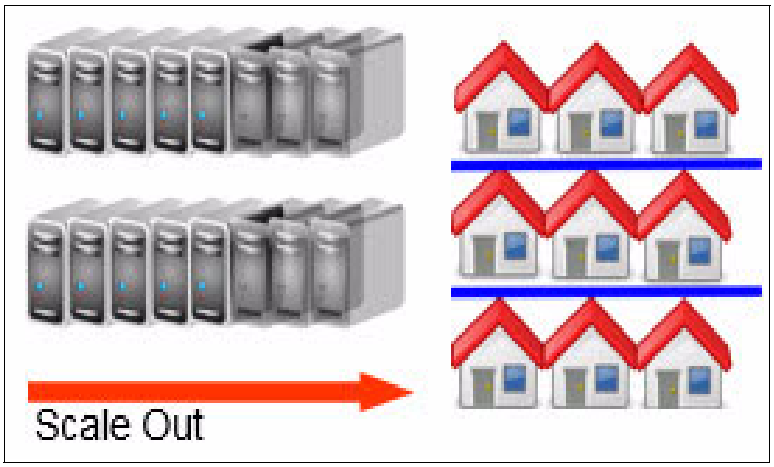

Figure 1-3 Scale-out architecture

1.3 LinuxONE and IBM Z versus distributed x86

Through technology, organizations can reduce costs, increase efficiencies, and increase profits if the capability of systems can be dynamically increased to improve service. With the new business model (mobile, social, big data, and cloud), a system architecture can dictate the performance, costs, and response time for critical applications.

As a consequence of the rapid growth of mobile internet and Internet of Things (IoT) for collecting data to enable people to communicate and collaborate, the big data analytics applications need special resources and scale capability to process massive amounts of data. The system users can take advantage of LinuxONE and IBM Z to enjoy good response time and performance.

The LinuxONE and IBM Z platforms have compelling total cost of ownership (TCO) and total cost of acquisition (TCA) metrics that show considerable cost saving over x86. It is a viable option for the scale-out applications and complex workloads. It is important to note that the ability to scale out on the x86 platform can easily become a problem in terms of TCO when a business grows in a way where hundreds or thousands of nodes are in use to support a popular application. In other terms, a system administrator can replace hundreds of x86 distributed servers with just a few LinuxONE or Z boxes running Linux. This simple choice can substantially decrease the TCO and make the infrastructure operates with less costs.

The key points that differentiate Linux workloads on LinuxOne and IBM Z against any x86 platform are:

- Technology maturity. Server virtualization on IBM z/Architecture® has been in the mainstream since roughly the 1960s, when z/VM started this virtualization concept.
- High utilization levels. Levels upwards of 90% are possible on a sustained basis.
- Performance.
- Supports Virtualization (z/VM and KVM).
- Floor space. Supports up to 8000 virtual guests.
- Encryption accelerated.

Research produced by Robert Francis Group (RFG) in 2015 provides evidence about how IBM LinuxONE and IBM Z platform are a better choice than the x86 platform. The information discussed in this research demonstrated Z and LinuxONE as the most reliable and scalable computing platform available, and ideal for consolidating hundreds of distributed servers. The reason is because these platforms deliver on the promise of a flexible, secure, and intelligent IT architecture that can be managed seamlessly to meet the business requirements.

Additionally, the IBM Z and LinuxONE return on investment (ROI) is quicker, and the TCA is less than that of x86 Linux solutions for organizations that have more than 20 Linux servers. The research can be found in the article entitled "10 Reasons LinuxONE is the Best Choice for Linux Workloads."

Another tool that helps to calculate the financial advantage of IBM LinuxONE for an organization is the IBM LinuxONE calculator. You can input data, such as hardware, workload type, and software, to get an assumption of savings when comparing LinuxONE to x86. The tool is available on the IBM site at IBM LinuxONE TCO calculator. Anyone can take advantage of this tool to calculate the savings by comparing LinuxONE and x86 servers.

Figure 1-4 illustrates the IBM LinuxONE TCO calculator tool.

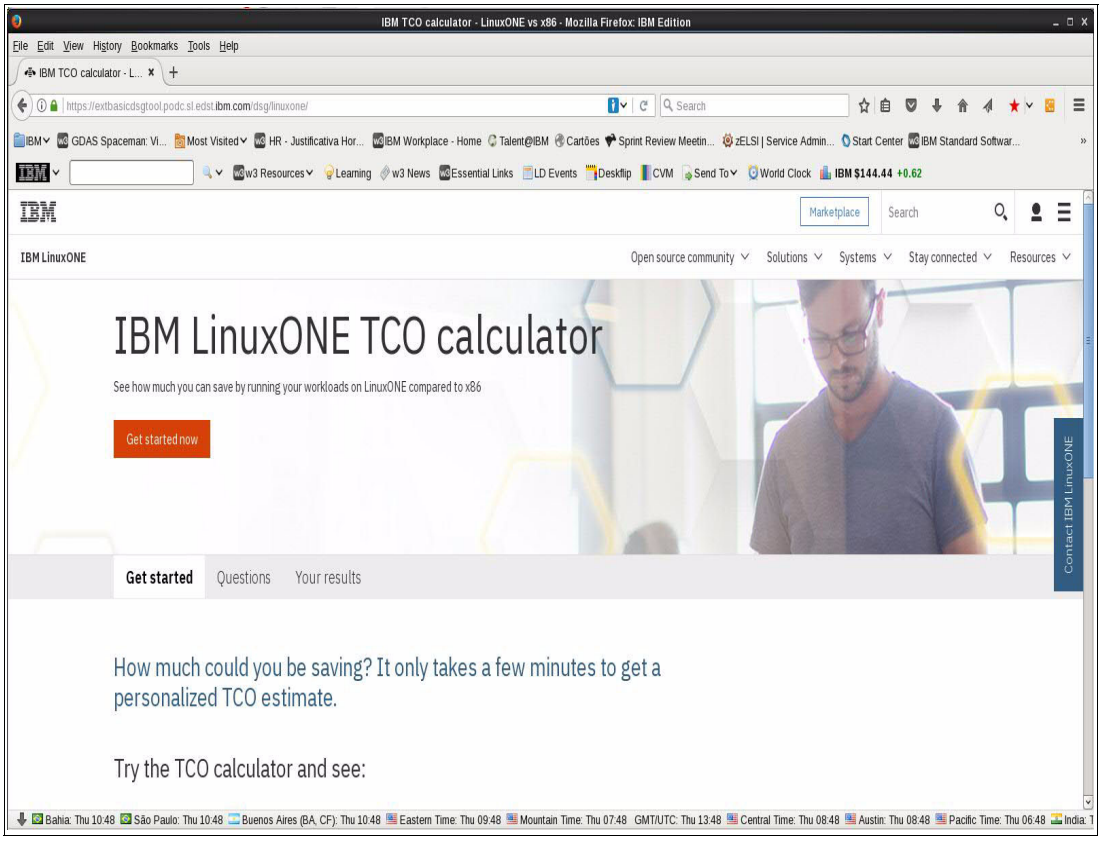

Figure 1-4 IBM LinuxONE TCO calculator

Linux has been a major factor in IBM for years now, and this new venture for LinuxONE and IBM Z is based on IBM z13® and z14 mainframe computers, which were designed to process high-volume mobile transactions and hybrid cloud workloads. The Linux applications can use the strengths and reliability features of the IBM Z hardware, and take advantage of the mainframe's speed, security, scalability, and availability. The server consolidation onto the new machines allows the complex applications to use the efficiencies of LinuxONE and IBM Z, as well as provide the flexibility of a bottom-up provisioning capability.

Additionally, the platform delivers various tools for systems administrators to help to keep the Linux-based services running at the most optimal level, and to use resources efficiently to react quickly to challenges and opportunities. The fact that Linux on IBM Z and LinuxONE run in a virtual environment does not affect the actual task of porting an application. IBM chose the LinuxONE and IBM Z platforms to centralize and consolidate IT workloads, significantly reduce the costs, improve availability, and achieve better performance and utilization.

It is common to see applications and Linux workloads deployed in x86 systems because of the low price that they have, but the hidden costs keep growing. The costs can be considerably higher if you take into consideration other factors, such as system availability, support, licensing, reliability, scalability, and performance. If you are interested to find and determine the best return on investment (ROI) and technical solutions for your workload, you can visit the IBM IT Economics - TCO Studies.

Some IT specialists have the assumption that a x86 server is cheaper than other platforms (like IBM Z or LinuxONE) because they do not include the many downstream costs that they might not be aware of. But it is important to note that hardware needs to run an operating system and applications, so the costs of hardware *and* software must be considered. In addition, the x86 architecture disadvantages can become increasingly onerous when the number of physical servers in a data center reaches a point where the system administrator can no longer manage them effectively.

The reason is that each server has licensing, support, security, and compliance issues. The system administrator needs to know how many servers are running on a network, where their storage is located, what software or operating system (OS) licenses are associated with them and who deployed them. Other not obvious costs include increased infrastructure costs (hidden costs).

The number of new servers added to the data center directly increases the electrical and cooling costs, in addition to increased requirements for floor space and networking. This phenomenon is commonly known as server sprawl, and is one of the top worries of many businesses and organizations today.

It is common to deploy each application on its own server, but it might result in an inefficient system without financial benefits. The organizations often face issues to keep the operation from running at top performance because of sprawl challenges. The x86 server sprawl has caused high maintenance costs, high power and cooling costs, and more space in the data center.

It is important to note that this dimension must be considered because server sprawl is affecting a high number of organizations and data centers around the world today. It is time to begin considering a range of increasingly vertically and horizontally converged solutions that vendors are offering for mission-critical applications.

Large data centers face a number of major issues. The most critical is their complexity, which requires highly skilled and experienced personnel to manage and control the environment. Other considerations include their power and cooling requirements, and the challenges of managing the lifecycle of so many assets. For example, when old servers are retired, are they turned off and removed from the data center? Alternatively, are they left continually powered-up with nothing running on them because no one is sure if any workloads are running on them?

Being intelligent and planning more for the criteria required to deploy an application can help eliminate many of these problems.

In general, the number of x86 servers in today's data centers has grown, which in turn leads to enormous challenges in designing an efficient and cost-effective data center infrastructure.

Although the x86 servers became more efficient and powerful, they face some architectural design challenges to scale compared with other platforms, such as LinuxONE or IBM Z. The x86 vertical architectures are problematic going forward, and can lead to performance inefficiencies when applied to modern applications, such as analytics, cloud, and big data, because of the following reasons:

- Data center space: New applications means new servers (More data center space).
- Cooling: Ensure adequate distribution of cool air, removal of hot air.
- High power demand: More servers means more power.
- Lack of flexibility: Hard to upgrade physical servers when production applications demand 24 x 7 availability. It means application downtimes.
- Hardware and software end of life (EOL): A technical refresh is needed when EOL is near and a replacement is required.
- Support: More system administrators.
- Security: Updating a high number of servers to fix a security exposure can be a hard and complex task.
- Software licensing: High costs.
- Availability: High cost to maintain two servers with the same resources.

1.3.1 Data center space

This is a critical factor that should be considered when planning the infrastructure for an application. To respond to a new application or to data growth, most decision-makers choose to add another system into the data center by default. As a result, more space is needed to accommodate the new racks, cabling, switches, and servers. These components need to be considered when deploying or resizing an application.

The decision-makers need to consider all data center functions and space requirements before drawing up final architecture plans and determining a budget for the data center building space.

Consolidation is the key to resolve data center management challenges and act efficiently to manage the new demands. The IBM Z and LinuxONE add and deliver services faster and more efficiently for shrinking those server farms back to a manageable size.

With greater reason, organizations need to develop a plan to support the cloud, mobile, and big data applications while freeing up space in the data center and reducing power and cooling requirements. This is the main challenge today.

The migration of the x86 servers onto a LinuxONE platform and IBM Z entails using the classic strengths of the IBM Z solutions, such as high availability (zero-downtime), high I/O bandwidth, the capacity to run diverse concurrent workloads, an extremely high level of security, and a superior disaster recovery environment. So, considering the Linux on IBM Z or LinuxONE as the default choice makes it the ideal platform for Linux Consolidation.

The LinuxONE and IBM Z platforms are the better way to reduce the need for electricity and data center space to support the new application demands (analytics, cloud, big data, and blockchain), while modernizing the Linux environments in data centers. IBM Z is actually designed to reduce energy usage and save floor space when consolidating distributed servers.

The need to maintain and keep the applications up-to-date costs a lot of money and energy, which can become a problem to a small business. However, the LinuxONE and IBM Z can be used to optimize the overall IT environment and improve service delivery by consolidating hardware, data, and software.

The consolidation onto a LinuxONE or IBM Z platform appears to be the long-sought remedy for x86 server sprawl and space problem. So, moving the distributed x86 servers to virtual machines (VMs) decreases the number of physical servers required in a data center, which results in significant data center space savings. It is also important to mention that the number of network connections, network switches, and other connections needed for the x86 servers are also reduced because an IBM Z machine and LinuxONE can accommodate them using just a few connections and specialized components.

Figure 1-5 illustrates the benefits of consolidating distributed servers onto a single LinuxONE or IBM Z. The gain in space and energy are significantly high compared with the x86 platform, which makes the IBM Z and LinuxONE platforms the best choice for Linux servers and critical workloads.

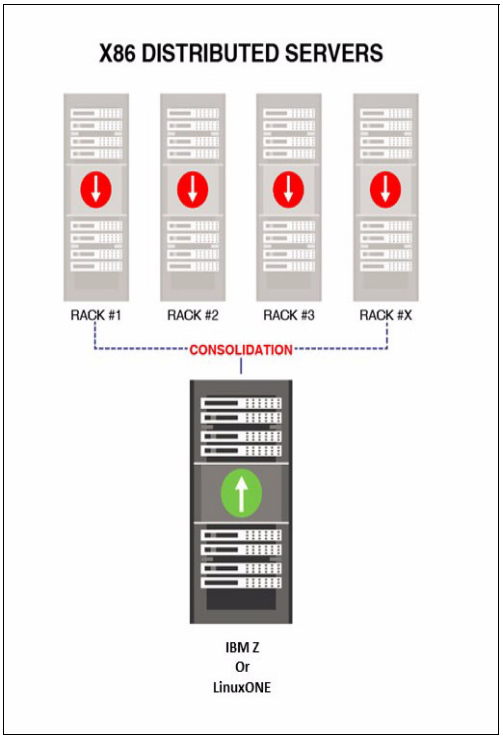

Figure 1-5 x86 scale-out infrastructure consolidated to a single IBM Z machine.

The innovation of the business is forcing organizations to establish a flexible infrastructure where a server can be quickly built in just a few seconds and quickly deleted when it is no longer needed. Following this further, new applications are demanding more resources and availability, and if you choose an IBM Z or LinuxONE platform, it can provide you the flexibility to respond quickly to real demands.

Traditionally, applications are deployed on x86 distributed systems, having hundreds of servers occupying a large area in the data center. Therefore, managing an environment like this can become laborious.

Sometimes this is an unmanaged activity, which is a problem for an organization that needs to respond quickly and adapt to new application requirements.

One of the major advantages of LinuxONE and IBM Z platforms is the ability to consolidate hundreds of scale-out distributed servers onto one machine. You can scale up to more than 1,000 virtual machines (VMs) within a single system (optimally, 200 - 8,000 VMs) and run up to 100% system utilization.

In summary, scaling out your environment using distributed servers seems to occupy more space in your data center and also increase the sprawl problem. However, consolidating the physical servers and scaling up onto an IBM Z or LinuxONE platform can allow more applications to occupy the same square footage area in your data center resulting in significant cost savings.

1.3.2 Cooling

Cooling is another critical part of a data center's infrastructure that should be taken in consideration. Although many still believe the IBM Z and LinuxONE platforms are more expensive than x86 distributed systems, this is not always the case. A number of studies have shown that LinuxONE and IBM Z servers can reduce power consumption and cooling costs by up to 40%.

It is essential that data centers measure just how much energy they use for non-computing functions, such as cooling. Several factors account for effective data center cost, but the number of servers per square meter (density) and power consumption are responsible for a good part of it. For more effective data center management, the hidden costs must be considered.

Remember that each server takes up space. Although this might seem harmless, in a distributed scale-out environment those systems are added so rapidly that their use of space, power, and cooling becomes a major challenge. Although the servers are well-organized in racks, the number of x86 systems is higher, and you probably need more cooling and so more space for cooling them.

It is a bit complicated to compare prices between LinuxONE, IBM Z, and x86 architectures. The x86 seems cheaper at first glance. However, when you consider that x86 systems generally need more space and cooling and are considerably more complex, The LinuxONE and IBM Z servers emerge as more cost-effective in the long term.

LinuxONE and IBM Z were developed to be the coolest processors, with less power consumption and cooling costs than the distributed scale-out systems.

Figure 1-5 on page 10 illustrates the space and cooling benefits when you consolidate x86 scale-out servers onto a single IBM Z or LinuxONE system.

1.3.3 High power demand

Power used by equipment in the data center is one of the more critical issues facing organizations today. Although the cost of the power is important, many clients must manage to power limits in their data centers, both limits in the physical power distribution capacity within the data center and possible limits imposed by power suppliers.

The most energy-efficient equipment is equipment that is no longer in use, whether it is a server, a router, or a storage device. Therefore, virtualization enables a client to combine several stand-alone systems onto fewer, more efficient pieces of equipment, reducing equipment footprint and power use.

These factors lead to a realization of savings and efficiencies that have been difficult to achieve through the design of even the greenest systems or buildings.

To achieve power efficiency in your data center, you need to solve (or minimize) the x86 scale-out sprawl problem. More importantly, ensure that your environment can respond quickly to changes, while controlling support and security.

Although x86 scale-out servers can bring value to your application, they do not have the power and performance efficiency of the LinuxONE or IBM Z platform. Having a high number of x86 servers can increase your data center costs and workload to manage them. The LinuxONE and IBM Z servers can significantly reduce the strain on energy consumption and administrative workloads when compared with x86 systems. So, combining your workloads on one LinuxONE or IBM Z server can also improve performance, increase security, and simplify management.

The LinuxONE and IBM Z platforms are extremely energy-efficient and save up to 75% on energy costs compared to x86 alternatives. These platforms are a good choice when power efficiency is needed without losing performance and are the strongest platforms to take better control of energy usage in your data center.

1.3.4 Lack of flexibility

All operating systems and applications need to be updated to fix a software problem or to get a new application functionality. This work is usually performed by the system administrator or someone else responsible for that specific product.

The upgrade of an operating system or an application seems a simple and quick task. However, upgrading 100 (or 500) servers can make it complex and hard to complete.

One of the disadvantages of x86 scale-out architecture is the lack of flexibility to perform required updates and changes. It is difficult to add or change pieces of code individually, without affecting the whole system.

The reason is that the x86 servers are generally not members of any cluster solution. Therefore, when a server reboot is required, the system administrator must negotiate with the customer and application users the most appropriate time to complete the work, and when to reboot the server to load the new code.

Additionally, there are some required updates that must be applied in a short period of time to close security software exposures and keep the system safe from cyber attacks. Therefore, the number of x86 scale-out servers in your infrastructure can directly increase your security problems and make things even more complex.

Although not mandatory, it is preferable to consolidate your x86 scale-out servers onto LinuxONE or an IBM Z machine. It enables you to grow with virtually limitless scale to handle the most demanding workloads and achieve the flexibility, scalability, and performance for your critical Linux applications.

1.3.5 Hardware and software end of life

As part of this ongoing evolution, the hardware and software is periodically discontinued. This provides straightforward product lifecycle policies that help you better design, deploy, and manage your environment.

Whatever your application deployment strategy might be, it is important to decide when your servers and software are obsolete and have a good plan to refresh them. Upgrading your assets too early, or waiting too long for upgrades might be risky to your business and return on investment (ROI).

Today, the administration of hardware and software assets has become much more complex to track, manage, and support. Without good inventory control, you can face challenges to manage your infrastructure, which can lead you to disorganized workflows, wasted time, and unnecessary delays. In other words, these situations can impact the entire organization.

Although the x86 scale-out architecture is common in the distributed world, their high numbers of small servers and difficult inventory control management can create problems for you. The reason is that you need more time and effort to track your hardware details and software assets, such as version, installation date, latest software updates, purchase date, usage, warranty expiration date, and location for every server. As the number of servers grow, the negative effects grow more rapidly than do the positive effects.

The LinuxONE and IBM Z scale-up architecture can help to simplify your infrastructure, making your asset inventory control and organization much easier and more accurate. Additionally, the server consolidation onto the LinuxONE and IBM Z platforms helps you eliminate the x86 sprawl problems and effectively use your resources to meet growing business needs, without increasing the data center's footprint.

With adoption of a LinuxONE and IBM Z solution and its virtualization technology, you can expand your services without adding any hardware. Having said that, scaling up can seem expensive. However, it does make it easier to control your systems. Also, scaling up provides better cost savings when compared to the x86 platform, including costs related to cooling, data center space, inventory control, and others.

1.3.6 Support

The administrative labor required to support hundreds of x86 scale-out systems is normally significantly high, and usually requires special administrative skills from the technical teams to keep the environment running at top performance. The nature and number of workloads has changed the way that system administrators support their environment. It is because they need to learn about new technologies quickly, and then try to reduce operational costs by their implementation.

One of the most obvious differences between scale-out and scale-up architectures is the cost perspective. Both depend on the technical support teams to make sure that the system infrastructure works as expected. However, the distributed scale-out approach creates a more complex infrastructure where every server needs to be purchased, configured, and operationally managed. This approach requires more system administrators, which results in higher operational costs and less efficient use of resources.

In fact, creating and maintaining a high-level quality of service from a large collection of distributed components demands significant knowledge and effort. It implies acquiring and installing extra equipment and software to ensure availability and security, monitoring, and managing.

Additional manpower and skills are required to configure, administer, troubleshoot, and tune such a complex distributed environment. Server consolidation can help to reduce the number of physical servers, and consequently the number of staff to maintain those qualities of service or serviceability.

The consolidation and virtualization of x86 scale-out servers onto the LinuxONE or IBM Z platform can eliminate the effects of x86 server sprawl with compelling performance advantages. This reduces labor costs, is more reliable, and can be much less expensive to install and support your critical applications.

1.3.7 Security

Security is the number one business priority for most organizations today. If you do not take measures to secure your environment, you will probably spend more money to recover your business than to implement a good security policy.

Although it is technically possible to maintain security on hundreds of physical servers, it is more complicated and expensive to keep the systems safely under your control. As the number of x86 servers increases, it becomes more difficult for system administrators to keep track of everything. Therefore, it is impracticable, cumbersome, and overly expensive to keep this environment compliant with security policies and auditable.

Choosing IBM z/Architecture to reduce x86 server sprawl, power costs, and your data center's footprint by consolidating scale-out servers onto a LinuxONE or IBM Z server is cost saving and helps you control and improve system security.

This platform comes with a key feature called *IBM Secure Service Container* (exclusive for z13 and z14) that represents a significant leap forward in data privacy and security capabilities. This feature increases security by preventing changes from being made at the command-line level.

Additionally, the LinuxONE and IBM Z servers have superior security performance, whether generating digital signatures or hashing encryption for the blockchain. The LinuxONE and IBM Z platforms use hardware accelerators that enable pervasive encryption not found on x86 platforms common to most public clouds.

For more information about this feature, see the IBM Redpaper™ publication, *Security and Linux on z Systems*, REDP-5464.

1.3.8 Software licensing

Software licensing is something that you must consider when choosing your scaling platform solution. Scaling up and scaling out are both viable architectures, but there are significant licensing cost differences between them. So, knowing the differences helps you to cut software license costs.

As previously stated, you need to think about software pricing for your Linux distribution and middleware. Otherwise, It can be more expensive to run (or deploy) your application in distributed systems than for the LinuxONE and IBM Z platform, because software licenses are attached to CPU cores. One of the advantages of this platform is to lower your licensing costs with an Integrated Facility for Linux (IFL) that is usually classified as a single core. There could be significant savings by consolidating many distributed servers to an IFL.

Note: The IFL is a processor dedicated to Linux workloads on IBM Z and expressly designed to reduce costs. It is supported by the Linux operating system, z/VM, IBM Wave for z/VM, and KVM virtualization technology.

The IBM z14 mainframe offers massive processing capacity with up to 170 user-configurable IFLs on one server.

Another important point is that some critical workloads, such as cloud, big data, and mobile, are gaining more space on the global market. It requires more CPU cycles and more complex software to support these new workloads.

As IBM keeps improving technology and continues its commitment to open solutions, IBM has introduced a new way of paying for IBM hardware and software that makes LinuxONE and IBM Z servers cheaper than they have even before. The elastic pricing model is based on how much you use the great capabilities of the IBM Z and LinuxONE system, where distributed systems are usually priced per processor core.

With the LinuxONE and IBM Z servers, you usually use fewer cores than you would use for distributed systems (x86) because of the strengths of these platforms. Therefore, you can save licensing costs by choosing them to run your workloads while having the world's fastest commercially available processor to deliver outstanding processing and response times.

Note: The LinuxONE TCO calculator tool can help to calculate the savings with software licenses and other cost factors.

In summary, you can have significant software pricing savings from core consolidation on the IBM Z platform, because it has flexible pricing where you can take advantage of the platform strengths and capabilities while controlling your IT costs.

Software licensing references

For more information about software licensing, see the following websites:

- Learn about Software licensing
- Base license agreements
- IBM Z Software Pricing reference guide
- IBM Z Software Pricing
- The IBM International Passport Advantage Agreement can be dowloaded from the "Learn about Software licensing" website

1.3.9 Availability

The cost of your application being down might be too high, and in many cases it brings irreversible losses to your business. The costs of downtime have increased over time, so companies are paying more attention to high availability (HA) solutions today than in the past. Also, costs of HA software have decreased considerably in way that makes much more sense to invest in protecting business continuity than to incur the downtime costs.

The challenge of your availability strategy is making it simpler to manage more data without increasing the complexity of managing it. In essence, the scale-out architecture provides the availability using more than one node through a load-balancing software.

From an x86 scale-out perspective, you need to maintain two or more servers with the same resources increasing the complexity and the hidden costs. As discussed earlier in this publication, you must consider the hidden costs before you make the decision about what platform to use. As LinuxONE and IBM Z scale up becomes more common, your application can be packaged in a single box to meet your performance and stability goals.

The LinuxONE and IBM Z scale-up architecture provides high reliability through various redundant hardware options, such as transparent CPU sparing and Redundant Array of Independent Memory (RAIM) memory modules, to prevent hardware from affecting the availability of your critical workload.

Simultaneously, several high availability software programs can run on LinuxONE and IBM Z to improve high availability based on different application requirements. The combination of high availability software, LinuxONE, and IBM Z is a key customer criteria where high availability is required to support your mission-critical workloads.

1.4 Considerations about the LinuxONE and IBM Z platforms

Refurbishing existing data centers can also prove cost-prohibitive, such as installing new cooling units that require reconfigured floors. The cost of power over time must also be considered in data center planning.

With the rising trends in energy costs is an accompanying trend towards high density distributed servers that stress the power capacity of today's environment. However, this trend has been met with rising energy bills, and facilities that do not accommodate new energy requirements. Distributed servers result in power and cooling requirements per square foot that stress current data center power thresholds.

Because these servers have an attractive initial price point, their popularity has increased. At the same time, their heat has created a problem for data centers whose total utility usage is consumed entirely by the energy proliferating servers. The LinuxONE and IBM's virtualization uses the power of many servers using a small hardware footprint. Today's IBM Z reduces the impact of energy cost to a near-negligible value when calculated on a per logical server basis because more applications, several hundred of them, can be deployed on a single machine.

With the LinuxONE and IBM Z servers, fewer physical servers running at a near constant energy level can host multiple virtual software servers. This setup allows a company to optimize the utilization of hardware, and consolidate physical server infrastructure by hosting servers on a few powerful servers. With server consolidation onto a LinuxONE system, often using open source software, companies can achieve better hardware utilization, and reduce floor space and power consumption, thus driving down costs.

Because of the many benefits offered by the LinuxONE and IBM Z platforms (including low floor space requirements, unmatched virtualization capabilities, energy efficiency, security, and reliability), mission-critical enterprise Linux applications are increasingly being deployed to LinuxONE and IBM Z servers, which are arguably the most reliable and available servers on the market.

As in any IT environment, elements such as sizing, capacity planning, and total cost of ownership (TCO) have always been important factors in supporting a decision to run or consolidate new applications onto a given platform. These same elements are essential in minimizing the possibility of critical situations, and they are key to enabling IBM to earn from its clients the high levels of satisfaction that we strive to deliver.

Part of the advice given in this chapter specifically targets the LinuxONE and IBM Z platforms, which are flexible platforms that offer alternatives in terms of hardware and software. For instance, KVM hosted Linux images running on the LinuxONE and IBM Z platforms are an efficient and cost-effective alternative to distributed servers.

Organizations are realizing that moving their workloads to the LinuxONE and IBM Z platforms provides significant advantages to their virtualized Linux infrastructure, with the power and flexibility they need for Linux workloads.

Figure 1-6 shows the advantages when moving x86 workloads onto a single IBM Z server.

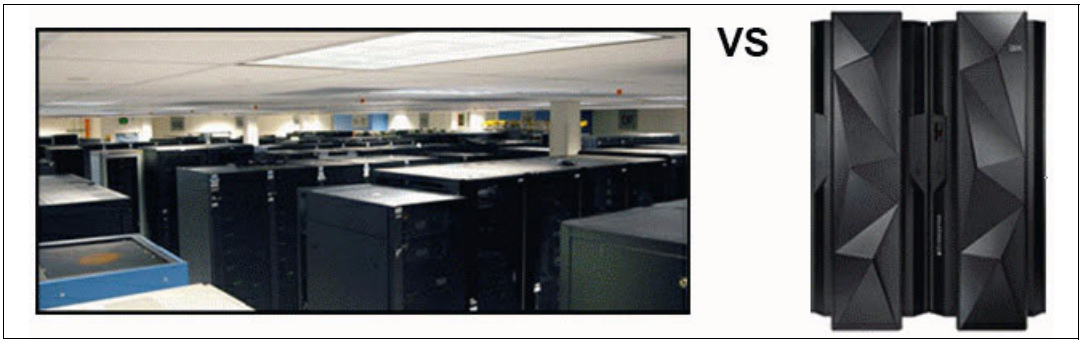

Figure 1-6 Scale out (x86) versus scale up (IBM Z and LinuxONE)

The movement of applications to the LinuxONE and IBM Z platforms can be caused by consolidation initiatives due to any of the following reasons:

- The desire to eliminate an island of technology (resources and technology used only for a few applications)
- Standardization of IT operations and development resources
- Applications running on antiquated hardware
- Applications running on proprietary software
- Limitations and high cost of supporting proprietary programming languages
- Using existing resources (hardware, software, and labor)
- Reduction in IT costs (reduction of data center space, power, cooling, and so on)

From a technical perspective, new applications (cloud, big data, blockchain, and so on) have many characteristics that need to be addressed when deciding on the design of the future architecture. From a business perspective, the main criteria is the cost and risk factor. You need to determine what the cost is for installing and maintaining an infrastructure to provide the necessary resources for the business.

These are some examples of the decisions that need to be made for all levels of the architecture to provide the proper performance and recovery times while, factoring in the proper investments that need to be made. In fact, virtualization is the key for a successful consolidation, and the IBM Z and LinuxONE servers are excellent at that. With the correct platform, you can deliver exceptional services and keep your environment secure, fast, and ready for growth.

There is a compelling reason for consolidation. Software (from many vendors) can be expensive and typically costs more than hardware. It is usually less expensive (and sometimes much less expensive) to replace multiple software licenses (for smaller machines) with 1 - 2 licenses (for larger machines).

Software license costs are often linked to the power of the system, but the pricing curves favor a few large machines. Software license costs have become a dominant factor in the growth and direction of the power industry. There are several nonlinear factors that make software pricing difficult. One such factor, the exponential growth of mainframe processing power, has been problematic in recent years.

Chapter 2, “Virtualization on the LinuxONE and IBM Z platforms” on page 19 of this publication contains detailed information about server consolidation onto LinuxOne and IBM Z.

2

Virtualization on the LinuxONE and IBM Z platforms

The base for a secure system is tightly related to the way the architecture, and, more specific, virtualization has been implemented on the IBM LinuxONE (LinuxONE) and IBM Z platforms. Since its inception more than 50 years ago, the architecture has been continuously developed to meet the increasing demands for a more secure and stable platform.

In this chapter, the following topics are discussed:

- The IBM LinuxONE and IBM Z servers
- Virtualization
- Security

Terminology: The terms *virtual server*, *guest*, and *virtual machine* are interchangeable. These terms are used throughout this book, depending on the component being used.

2.1 The IBM LinuxONE and IBM Z servers

Digital is transforming the industry. To be competitive, enterprises must deliver trusted services to their clients while accelerating value. This requires an open source platform that speeds your developers' creative genius and a highly secure cloud infrastructure that provides instantaneous data delivery any day of the year, whether you have thousands, or millions, of simultaneous users.

With the launch of IBM LinuxONE Emperor™ II, IBM is adding major enhancements to the platform, making it an efficient and cost-effective system for hosting various enterprise-grade Linux workloads.

2.1.1 Benefits

Linux is available on a large variety of computing platforms, from set-top boxes and handheld devices to the largest servers. This flexibility means that after your applications are running on Linux, you are no longer tied to a specific hardware platform. You have control over the choice of hardware platform that supports your application. Workloads running on LinuxONE benefits from a hardware platform that includes specialized processors, cryptographic cards with dedicated processors.

A major benefit of Linux is that it is open source. The software is weighed down by licensing fees and its source code is freely available. Hundreds of Linux distributions are available for almost every computing platform.

All of the Linux distributions, such as Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server, and UBUNTU, provide customers who use Linux with various support options. These options include 24 x 7 support with one-hour response time worldwide for customers running production systems. In addition to the Linux operating system, all the major Linux distributions offer a number of other open source products that they also fully support.

To simplify problem determination, IBM customers can contact IBM in the first instance and, if it is a new problem with Linux, IBM works with the distributors to resolve the problem. The increased interest and usage of Linux resulted from its rich set of features, including virtualization, security, Microsoft Windows interoperability, development tools, a growing list of independent software vendor (ISV) applications, performance, and, most importantly, its multiplatform support.

This multiplatform support allows customers to run a common operating system across all computing platforms, which means significantly lower support costs and, for Linux, no incremental license charges. It also offers customers the flexibility of easily moving applications to the most appropriate platform. For example, many IT organizations choose Linux for the ability to scale databases across highly scalable hardware.

2.1.2 Hardware strengths

IBM LinuxONE is a highly engineered system for both data serving and premium cloud services. It accomplishes this by delivering industry-leading security, faster data through application performance, and the highest levels of scalability to handle the most unpredictable data and transaction growth.

With IBM LinuxONE Emperor II, you can protect your valued data:

- Whether in flight or at rest
- With low resource usage
- Without changing your applications
- With minimal impact to service level agreements (SLAs)

IBM LinuxONE Emperor II delivers IBM Secure Service Container technology that can be used to create isolated partitions, helping to protect data and applications automatically helping to keep them safe from insider threats and external cybercriminals.

IBM LinuxONE Emperor II provides workload isolation through logical partitions certified at EAL5+. IBM LinuxONE Emperor II provides encryption of data at rest without application change and with minimal performance impact. In addition, IBM LinuxONE Emperor II uses the Crypto Express6S for trusted multi-tenant Hardware Security Module service designed for FIPS142 level 4.

IBM LinuxONE Emperor II delivers the performance and scale needed to address today's business demands. IBM LinuxONE Emperor II has up to 170 cores using the fastest commercially available processor (5.2 GHz). IBM LinuxONE Emperor II offers dedicated I/O processors to enable data to be moved quickly without impacting performance. IBM LinuxONE Emperor II delivers industry-leading Java performance that runs applications up to 50% faster than on x86 alternatives

LinuxONE supports a wealth of new open source software products, such as Go, Python, Scala, Node.js, Docker, Spark, MongoDB, PostgreSQL, and MariaDB. Vertical scale allows IBM LinuxONE Emperor II to scale up to 2 million Docker containers in a single system, or to support 17 TB MongoDB databases.

IBM LinuxONE server strengths

The following sections describe some of the strengths of the IBM LinuxOne platform.

Reliability

LinuxONE is extremely reliable, because there are redundant processors, I/O, and memory. It also has error detection and correction, and Remote Support Facility helps in this area.

Availability

At near 0% downtime, LinuxOne has fault tolerance, automated failure detection, and nondisruptive hardware and software changes.

Scalability

IBM is committed with scalability and offers the performance and scale needed to address today's business demands:

- IBM LinuxONE Emperor II scales up to 170 physical processors and 2.2 times more of improved memory.
- IBM LinuxONE Emperor scales up to 141 physical processors and up to 10 TB of memory.
- IBM LinuxONE Rockhopper™ scales up to 20 physical processors and up to 4 TB of memory.

Security

The IBM Z and LinuxOne servers have a rich history of delivering a secure infrastructure. It is built on a set of hardware security capabilities that include multi-state operation modes, storage key isolation, high-speed standards-based encryption, logical partitions, and many other features and benefits.

The IBM LinuxONE Emperor II is designed for Common Criteria Evaluation Assurance Level 5+ (EAL5+) certification for security of logical partitions. This means that IBM LinuxONE Emperor II is designed to prevent an application running on one operating system image on one LPAR from accessing application data running on a different operating system image on another LPAR on the same hardware server.

The IBM point of view is that encrypting as much of your data in flight and at rest as possible can help reduce potential data breach risks and financial losses.

Note: For further details, see *Security and Linux on z Systems*, REDP-5464.

Just-in-time deployment of resources

When needed, more capacity can be added due to a specific workload in a certain seasonal period of time, temporarily or permanently, IBM Z family and IBM LinuxOne have the following options:

- **On/Off Capacity on Demand** provides temporary processing capacity to meet short-term requirements, or for testing new applications.
- **Capacity Backup (CBU)** enables you to replace model capacity or specialty engines to a backup server in case of an unforeseen loss of server capacity because of an emergency. CBU ensures that customers can access extra capacity during a disaster recovery situation without having to purchase more capacity.

Power and cooling savings

With its low power and cooling requirements, the IBM LinuxONE and IBM Z family are an ideal platform for the consolidation of distributed servers.

Virtualization

A combination of hypervisors configured to allow unparalleled flexibility in Linux deployment. IBM PR/SM™, z/VM, and KVM employ hardware and firmware innovations that make virtualization part of the basic fabric of the IBM LinuxONE platform:

- PR/SM

  PR/SM is a standard component of all of the IBM LinuxONE and IBM Z family servers, which enables LPARs to share system resources. PR/SM divides physical system resources, both dedicated and shared, into isolated logical partitions. Each partition is like an independent system running its own operating environment. It is possible to add and delete resources, such as processors, I/O, and memory, across partitions while they are actively in use.

- KVM

  It is an open source hypervisor kernel-based virtual machine (KVM) and is supported on all IBM LinuxONE and IBM Z family servers.

- z/VM

  z/VM V6.4 virtualization technology is designed to run hundreds to thousands of Linux servers on a single IBM Z or IBM LinuxONE server with the highest degrees of efficiency and elasticity. It is a very stable and mature IBM virtualization operating system with more than 50 years of development. Over this period, an increasing set of new features build on the microcode and new functions of the LinuxOne and IBM Z servers.

  z/VM enables you to bring new virtual servers online in a matter of minutes (or less) to accommodate guest servers growth, including Linux on IBM Z or other guests.

With the improved memory management support of IBM LinuxONE Emperor II, which enables greater concurrency in address translation, z/VM has improved workload performance compared to the original Emperor, particularly when z/VM is configured to use multithreading.

2.1.3 How it works

LinuxONE and IBM Z family have the world's fastest commercially available processor. Built for speed, both platforms support simultaneous multithreading for Linux and Java workloads. It helps deliver outstanding transaction processing and data serving performance.

The LinuxONE and IBM Z family offers excellent economies of scale, supporting up to 141 Linux cores, up to 85 logical partitions for secured workload isolation, and more efficient use of critical data. Coupled with up to 10 TB of memory, and large memory pools with four levels of cache, it is ideally suited for consolidating large-scale distributed environments and the introduction of new in-memory workloads.

It can support up to thousands of virtual Linux servers on a single footprint. That means that the virtualization capabilities in a single LinuxONE system or IBM Z family can result in a less complex Linux infrastructure with fewer components, less management, less space requirements, and optimized core workload efficiency.

What the z/Architecture brings to Linux

z/Architecture and Linux provide the following advantages:

- The most reliable hardware platform available
- Centralized Linux servers can be easier to manage
- Designed to support mixed workloads:
  - Enables consolidation while maintaining one server per application
  - Complete workload isolation
  - High-speed inter-server connectivity
- Utilization can usually exceed 90% and handle peak loads around 100%
- Non-disruptive scalability:
  - IBM Z and LinuxONE can scale processors and memory depending on the model
  - Dedicated I/O processors (no additional costs)
  - Hundreds to thousands of Linux server guests

For more information about how LinuxONE (Figure 2-1) provides the perfect platform for running open source, mission-critical Linux applications for fast deployment and cost savings, see the IBM LinuxONE website.

Figure 2-1 LinuxONE servers

2.2 Virtualization

In the late 1960s, hardware was expensive and there was not much of it available. Therefore, virtualization technology was developed to address the business needs to better use limited resources. In one single server, we could run in parallel several environments that kept full isolation and gave the illusion of owning the hardware to each user.

Virtualization has re-emerged as a solution to cost issues. The benefits to running X86 servers in a virtual environment are immense. IT managers are mainly driven by the need to reduce the costs of IT labor, and the savings are most pronounced when companies choose LinuxOne or IBM Z server to run the virtualized applications.

Virtualization has become a key solution for simplifying service management and reducing energy costs in Data Centers. Virtualization also reduces the other costs mentioned previously in this publication. The benefits from virtualization can be profound, and there are good results in virtualization and server consolidation when comparing with non-consolidated systems. There are even more opportunities for efficiency gains and cost savings in an x86 server world.

This chapter focuses on the use of virtualization to solve many data center issues and costs, using all the resources of LinuxOne or IBM Z servers, with more efficiency and productivity.

2.2.1 Benefits of virtualization

The cost of administering IT systems is growing faster than the cost of new hardware for those systems, because the complexity of those systems requires growing numbers of people to manage them. The primary concern of management is to contain cost, while increasing revenue levels.

Introducing virtualization can be a critical first step in managing computing infrastructures in the following ways:

- By lowering the cost of existing infrastructure
- By reducing the complexity of adding resources to that infrastructure
- By building heterogeneous infrastructure across multiple data centers, making those centers more responsive to business needs

The benefits of virtualization vary, depending on the objectives, the specific virtualization technologies selected, and the existing IT infrastructure. Not all users obtain the same benefits from implementing a particular virtualization solution. However, users realize many of the following benefits to some degree, even when using virtualization for server consolidation.

Higher resource utilization

Virtualization enables the dynamic sharing of physical resources and resource pools, resulting in higher resource utilization, especially for variable workloads where the average needs are much less than an entire dedicated resource.

Lower management costs

Virtualization can improve staff productivity:

- Reducing the number of physical resources that must be managed
- Hiding some of the resource complexity
- Simplifying management tasks through automation, better information, and centralization
- Enabling workload management automation

Virtualization also enables common tools to be used across multiple platforms.

Usage flexibility

Virtualization enables resources to be deployed and reconfigured dynamically to meet changing business needs.

Improved security and guest isolation

Virtualization enables separation and compartmentalization that is not available with simpler sharing mechanisms, and that provides controlled, secure access to data and devices. Each virtual machine can be completely isolated from the host machine and other virtual machines. If one virtual machine crashes, none of the others is affected.

Note: Virtualization prevents data from leaking across virtual machines, and ensures that applications communicate only over configured network connections.

Higher availability

Virtualization enables physical resources to be removed, upgraded, or changed without affecting users.

Increased scalability

Resource partitioning and aggregation enable a virtual resource, depending on the product, to be much smaller or much larger than an individual physical resource, so you can make scale adjustments without changes to the physical resource configuration.

Interoperability and investment protection

Virtual resources can provide compatibility with interfaces and protocols that are unavailable in the underlying physical resources. This is increasingly important for supporting existing systems and ensuring compatibility with earlier versions, as with z/VM.

Improved provisioning

Virtualization can enable resource allocation to a finer degree of granularity than individual physical units. Virtualized resources, because of their abstraction from hardware and operating system issues, are often capable of recovering much more rapidly after a crash than a physical resource.

Consolidation

Virtualization enables multiple applications and operating systems to be supported in one physical system, in addition to consolidating servers into virtual machines on either a scale-up or scale-out architecture. It also enables systems to treat computing resources as a uniform pool that can be allocated to virtual machines in a controlled manner.

Note: X86 architecture was not designed for virtualization. Consequently, high-performance virtualization is difficult to achieve.

2.2.2 Levels of virtualization on the LinuxONE and IBM Z servers

IBM Z and LinuxONE servers feature unparalleled levels of virtualization at the hardware and software levels. Virtualization provides resource optimization and introduces the granularity that is required by organizations to create and separate environments of resources that might provide an outstanding solution.

2.2.3 Hardware virtualization

These facilities are available in IBM Z and LinuxONE servers, and they are used to partition the hardware into LPARs. They are classified as type 1 hypervisors (Figure 2-2) and run directly on the hardware. They are known as *native* or “bare metal”.

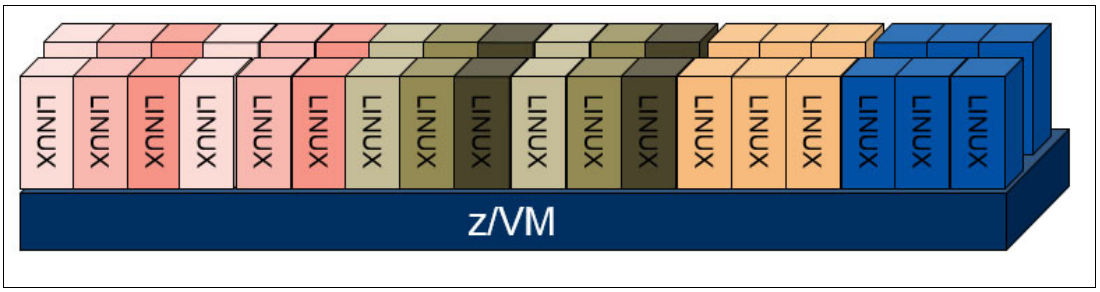

Figure 2-2 Hardware and software hypervisors

PR/SM

PR/SM is the facility in IBM Z and LinuxONE servers that provides another layer of virtualization. It is a hypervisor classified as type 1 and enables multiple logical partitions (LPARs) to share physical resources, such as processors, memory, channel paths, DASDs, and so on. It is extensively used in IBM Z and LinuxONE servers, as mentioned in “Virtualization” on page 22.

DPM

Dynamic Partition Manager (DPM) is a new administrative mode introduced to LinuxONE servers. A system can be configured in either DPM mode or PR/SM mode, as shown in Figure 2-3. The mode is enabled before system power-on reset (POR).

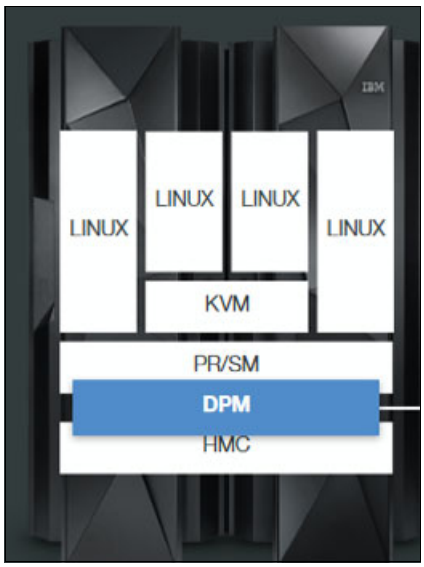

Figure 2-3 Dynamic Partition Manager (DPM) + PR/SM = easy to use

IBM Dynamic Partition Manager (DPM) provides a greatly simplified way to configure a LinuxONE server. It supports Linux and KVM systems with FCP-attached SCSI storage. It removes the need for IBM Z configuration tools, such as Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM).

The entire system configuration is performed from the HMC. It eliminates the need to build a stand-alone IOCP input deck for I/O devices. It also eliminates the need for the second power-on reset (POR) to enable dynamic I/O, which is required in non-DPM configurations. DPM provides for fully dynamic reconfiguration of CPU, memory, and I/O resources. Adding and dynamically reconfiguring I/O devices is greatly simplified with DPM.

A simplified vision of DPM

DPM provides the following advantages:

- **Fast**. Much faster than managing with HCD and HCM. From hours to minutes
- **Easy**. Intuitive user interface. No need for multiple administrators with different skills or tools. Do *not* expect First In Enterprise Linux clients to adopt the previous way.
- **Powerful**. The same efficient PR/SM hardware virtualization without the complexity. It supports dynamic configuration changes with just a few clicks of the mouse. It provides a foundation for “bare metal” Cloud.

Note: Review the linked HMC publication for current details about DPM operation and capabilities.

2.2.4 Software virtualization (hypervisors)

Hypervisors type 2 run on a host operating system. They control the entire hardware or only a partition of this hardware called LPAR. The LinuxOne and IBM Z servers have two hypervisors for virtualization.

2.2.5 KVM

IBM is changing how KVM is delivered. Instead of IBM offering our KVM product, the Linux distributions can offer a KVM hypervisor integrated into their Linux distribution.

Delivery of KVM hypervisor and Linux on IBM Z can be done more efficiently through our Linux distribution partners. This delivery method helps simplify the delivery of open source infrastructure, and makes it easier to obtain and install KVM distributions, including open source cloud infrastructure components. These are the KVM distributions:

- SUSE Linux Enterprise Server 12 SP2 - KVM
- Ubuntu 16.04 long-term support (LTS) - KVM
- Ubuntu 17.04 - KVM

KVM for IBM Z and LinuxONE servers can manage and administer multiple virtual machines, enabling a large number of Linux-based workloads to run simultaneously on the IBM Z servers. IBM has a long history of providing security for applications and sensitive data in virtual environments with z/VM. It is the most securable platform in the industry, with security integrated throughout the stack in hardware, firmware, and software.

Advantages of using KVM for IBM Z and LinuxONE servers

KVM for IBM Z offers enterprises a cost-effective alternative to other hypervisors. It has simple and familiar standard user interfaces, offering easy integration of the LinuxONE servers and IBM Z servers into any IT infrastructure.

In addition, KVM for IBM Z can help make platform mobility easier. Its live relocation capabilities enable you to move virtual machines and workloads between multiple instances of KVM for IBM Z without incurring downtime.

Main features and benefits of KVM

Table 2-1 describes the major benefits of KVM.

Table 2-1 KVM main features

Feature	Benefits
KVM hypervisor	Supports running multiple Linux VMs on a single system
CPU sharing	Enables the sharing of CPU resources by VMs
I/O sharing	Enables the sharing of I/O resources among VMs
Memory and CPU overcommitment	Supports the overcommitment of CPU, memory, and swapping of inactive memory
Live VM relocation	Enables workload migration with minimal impact
Dynamic addition and deletion of virtual I/O devices	Reduces downtime to modify I/O deviceconfigurations for VMs
Thin-provisioned VMs	Enables copy-on-write virtual disks to save on storage

Feature	Benefits
Hypervisor performance management	Supports policy-based, goal-oriented management and monitoring of virtual CPU resources
Installation and configuration tools	Supplies tools to install and configure KVM for IBM z Systems®
Transactional execution use	Provides improved performance for running multi-threaded applications

For further details, see: *Getting Started with KVM for IBM z Systems*, SG24-8332.

2.2.6 z/VM

Virtualization has become a key technology for simplifying service management and reducing energy costs in data centers. z/VM is the most mature solution in the market, because it was developed in the late 1960s to make more efficient use of hardware, and IBM invests heavily to improve its functionalities whenever a new release is launched.

z/VM continues to support more virtual machines in a single footprint with more excellent service levels than any other solution, and you can scale up your capacity without requiring additional support personnel.

As a major component of the TCO story, IBM introduced subcapacity pricing for z/VM V6, along with subcapacity terms for select z/VM-based programs running in a z/VM system. This can help customers to improve alignment between software use and software pricing:

- Add hardware capacity without necessarily increasing software pricing
- Enable greater flexibility with software licensing
- Achieve better price/performance for z/VM products

The z/VM subcapacity program licensing requirements

Software pricing at less than full machine capacity can provide more flexibility and improved TCO as a client manages the volatility and growth of new workloads. You can see more about licensing costs in 1.3.8, “Software licensing” on page 14.

The marriage of z/VM and Linux on IBM Z is perfect because z/VM can have hundreds of Linux servers running harmonically when z/VM takes care of all of the resources of the LPAR, with the benefit from improved price performance as workloads grow.

A fundamental strength of z/VM is the ability for virtual machines to share system resources (Figure 2-4) with high levels of resource utilization. z/VM V6.4 provides even greater levels of extreme scalability, security, and efficiency to create opportunities for cost savings, while providing a robust foundation for cognitive computing on the IBM Z and LinuxONE platforms.

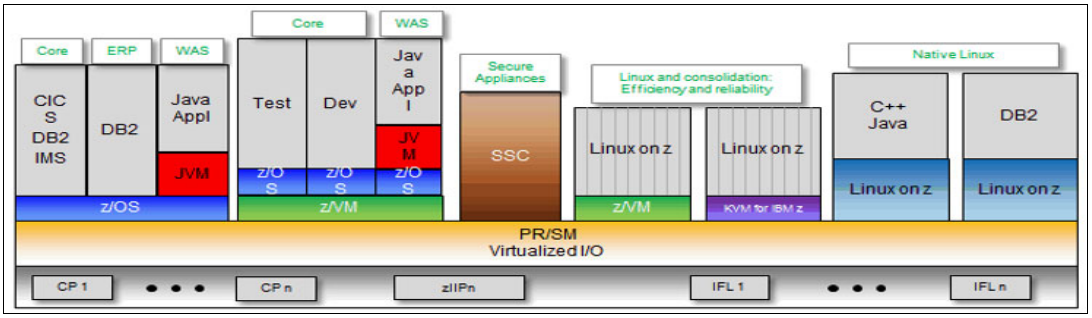

Figure 2-4 Several options of delivering virtualization

When multiple Linux servers run on IBM Z and LinuxONE, each Linux system acts as though it has dedicated access to a defined portion of the IBM Z machine, using a technique known as *timesharing*. Each Linux server runs in its own virtual machine whose characteristics (for example, memory size and number of CPUs) define the hardware that Linux sees. The allocation and tuning controls in z/VM specify how real hardware resources are allocated to the virtual machine.

Note: The most important z/VM capabilities are rapid deployment of Linux guests and high server consolidation ratio.

Main features and benefits of z/VM

Table 2-2 describes the major benefits of z/VM.

Table 2-2 Main features of z/VM

Feature	Benefits
z/VM hypervisor	Supports running hundreds of Linux on Z servers with high performance.
Resources sharing	Allows sharing resources as processors, DASDs, VSwitches, and so on.
I/O support	z/VM is well-integrated with IBM Z servers, using the SAP processors to manipulate the I/O subsystem.
Memory and CPU overcommits	Supports the overcommit of CPU, memory with CPU utilization without any service degradation.
Dynamic memory adding	z/VM has the option to change the memory size of Linux servers automatically (if they are defined on the user directory as standby memory).
System Single Image (SSI)	Supports an SSI cluster with 4 z/VM members, which helps to avoid planned outages.
Live Guest Relocation (LGR)	Supports LGR when using SSI, moving the Linux servers to another z/VM member without any interruption.
Virtual disks	Supports the usage of virtual disks for Linux servers to be in the real memory to fast access.

Feature	Benefits
Performance management	z/VM has several commands to monitor performance. It supports also Performance Toolkit product that has more monitoring and reports about z/VM performance.
SMT	Provides improved performance for running multithread applications. Supports simultaneous multithreading (SMT), enabling up to two threads to be used in IFL core.
NSS and DCSS segments	z/VM supports NSS and DCSS segments, so you can have only one copy of the kernel or product loaded in real memory once, and other Linux guests can use the same copy, improving the usage of real memory.
VSWITCH	z/VM supports a virtual switch that enables connectivity to external LAN segments without requiring a router, cables, and so on.
Real and virtual memory enhancements	z/VM 6.4 supports real memory of 2 TB, and Linux guests limit remains at 1 TB.
Linux on Z cloning	A Linux guest under z/VM can be quickly cloned to make another Linux image easily, in a few minutes.
Upgrade in place	Enables a smoother upgrade of existing z/VM 6.2 and z/VM 6.3 systems to z/VM 6.4.

2.2.7 z/VM 6.4 overview

z/VM V6.4 supports IBM Z, IBM LinuxONE servers, as well as Red Hat, SUSE, and Ubuntu Linux distributions. Support for simultaneous multithreading (SMT) technology extends per-processor, core capacity growth beyond single-thread performance for Linux on IBM Z running on an IBM Integrated Facility for Linux (IFL) specialty engine on IBM Z and LinuxONE servers.

z/VM multithreading technology support provides additional price/performance benefits over previous hardware generations and can meet workload requirements transparently. Improvements made in the areas of reliability, availability, and serviceability allow low-end devices, such as IBM Storwize® V7000, V840, and V9000, to be attached to a z/VM host, removing the need for a SAN Volume Controller.

z/VM V6.4 is a supported environment using IBM Dynamic Partition Manager for Linux-only systems with SCSI storage. This simplifies system administration tasks for a more positive experience by those with limited mainframe skills. IBM Wave Version 1 Release 2 (V1.2) is now included in z/VM V6.4 as a priced feature. It can greatly simplify the task of administering a z/VM environment.

z/VM V6.4 can help you extend the business value of IBM Z and IBM LinuxONE technology across the enterprise by integrating applications and data, while providing exceptional levels of availability, security, and operational ease.

World-class virtualization technology offered by z/VM enables you to host many virtual servers running different operating systems on IBM Z and LinuxONE servers, as shown in Figure 2-5.

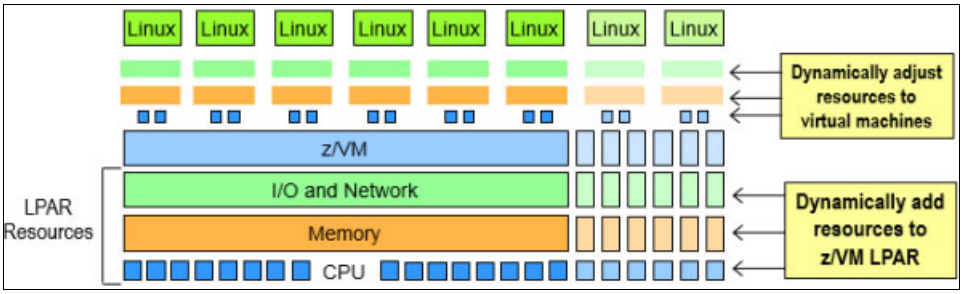

Figure 2-5 IBM Z and LinuxONE hypervisor virtualization

The ability of z/VM to support multiple machine images and architectures provides a highly flexible production and test environment for IBM Z and LinuxONE operating systems to:

- Simplify migration from one release to another
- Facilitate the transition to newer applications
- Provide a test environment whenever one is needed
- Consolidate many systems onto one physical server

For further details, see the IBM z/VM link.

2.2.8 Single System Image feature

Available since z/VM V6.2, Single System Image (SSI) is an optional feature and connects up to four z/VM systems as members of an SSI cluster, as shown in Figure 2-6. Each member can run on the same or different IBM Z or LinuxONE servers, which simplifies management of a multi-z/VM environment:

- Enables live guest relocation (LGR) of running Linux guests among z/VM members
- A single user directory of all virtual machines (VMs) is defined
- Cluster management from any member:
  - Apply maintenance to all members in the cluster from one location
  - Issue commands from one member to operate on another
- Built-in cross-member capabilities
- Resource coordination and protection of network and disks

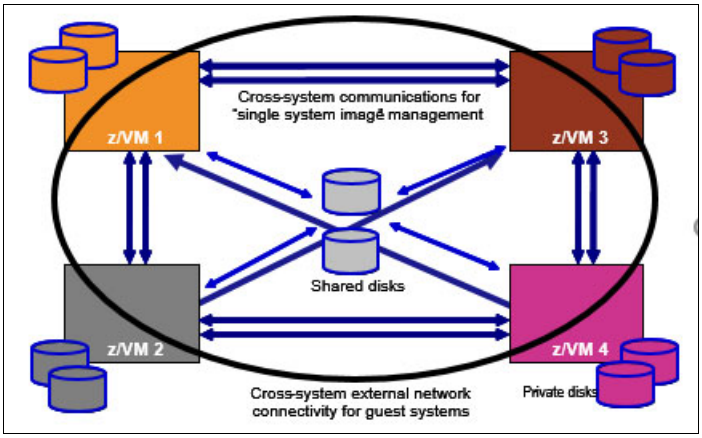

Figure 2-6 z/VM SSI configuration

SSI cluster management has greater reliability:

- Cross-checking of configuration details as members join to the cluster and as resources are used:
  - SSI membership definition and identity
  - Consistent definition of shared spool volumes
  - Compatible virtual network configurations (MAC address ranges, VSwitch definitions)
- Cluster-wide policing of resource access:
  - Volume ownership marking to prevent dual use
  - Coordinated minidisk link checking
  - Autonomic minidisk cache management
  - Single logon enforcement

2.2.9 Live guest relocation (LGR)

With the IBM z/VM SSI, a running Linux on IBM Z virtual machine can be relocated from one member system to any other, a process known as live guest relocation (LGR). LGR occurs without disruption to the business. It provides application continuity across planned z/VM and hardware outages and flexible workload balancing that enables work to be moved to available system resources.

Note: Linux on IBM Z is currently the only operating system supported for relocation.

There are several reasons why you might need to relocate a running virtual server:

- Availability through flexibility for planned outages
- Maintenance of hardware or software
- Fixing performance problems
- Administrative savings
- Increased isolation capability
- Management and balancing of workloads

Relocating virtual servers can be useful for load balancing and for moving workload off of a physical server or member system that requires maintenance. After maintenance is applied to a member, guests can be relocated back to that member, thereby allowing you to maintain z/VM as well as keeping your Linux on IBM Z virtual servers available.

For details about SSI cluster and LGR, see:

- *z/VM CP Planning and Administration version 6 release 4*, SC24-6178.
- *An introduction to z/VM Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8006
- *Using z/VM 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039

2.2.10 z/VM networking

Most servers need network connectivity. On IBM Z and LinuxONE servers, network connectivity is provided by Open Systems Adapter (OSA) features that present direct, industry-standard local area network (LAN) connectivity and communication in a networking infrastructure. OSA adapters use IBM z/Architecture called *Queued Direct Input/Output* (QDIO) that has a highly efficient data transfer mechanism, using system memory queues to directly exchange data between the OSA microprocessor and the network stack running in z/VM.

The LinuxONE and IBM Z platforms have different networking methods available, which are provided by either the I/O subsystem that controls the IBM Z LPARs, such as IBM PR/SM, or by the z/VM hypervisor.

It is important to ensure that you have an optimal network configuration.The virtual network that is provided by z/VM for Linux guests communication offers high throughput and better reliability (failure tolerance).

Typically, z/VM provides three networking options:

- Virtual switches
- IBM HiperSockets™
- Guest LANs

These options give Linux on z guests the ability to communicate over the network. The Linux guests use virtual devices as their own physical network adapters. For complex environments that require outside LAN communication, one of the best choices is virtual switches.

Virtual switches

The virtual switch (VSWITCH) method allows Linux on System z guests to connect over the network. The VSWITCH definition in z/VM enables mapping of physical networks into virtualized networks, which offers a secured network simplification and centralized management of a network topology. Multiple VSWITCH can be defined in a single z/VM environment that represents hubs in isolated networks in the same z/VM virtualization layer.

This method is both efficient and secure as it is provided by z/VM to Linux guests. VSwitches are software switches that offer many of the capabilities that are provided by a real switch, avoiding the need for real switches, cables, and so on.

Note: A virtual switch allows connectivity to external LAN segments without requiring a router.

VLAN aware and unaware

Using a VSwitch is important during network communications when running Linux guests under z/VM. Two options are available for virtual switches: Aware and unaware. Carefully select your option, because it changes the behavior of how virtual switches handle and process packages and frames. Depending on what option you select, the virtual switch ignores or processes the VLAN tags:

- **VLAN aware**. In this mode, the virtual switch reads and handles VLAN tags. The switch port that is connected to the OSA-Express port must be configured as a trunk port (check with your network administrator). The trunk port carries traffic from all VLANs.
- **VLAN unaware**. In this mode, the virtual switch ignores VLAN tags. The switch port that is connected to the OSA-Express port must be configured as an access port (check with your network administrator). The access port carries traffic for a single VLAN.

Guest LANs and HiperSockets

Guest LANs are virtual networks that are used to connect Linux guests in the same z/VM LPAR. They facilitate the communication between these guests without any additional hardware, using QDIO.

Although the Guest LAN method is still available and used in some scenarios, for complex network environments where there is intense network traffic activity and external connectivity is required, virtual switches are the best choice.

IBM HiperSockets

IBM HiperSockets is a networking option that is defined in the I/O configuration of IBM Z and LinuxONE servers. This type of network connection provides only an internal connection among LPARs.

The main usage for this type of network option is as an internal connection between two IBM Z or LinuxONE LPARs on the same box (CPC). An example is the connectivity between an IBM z/OS® LPAR that serves IBM DB2® and a z/VM LPAR where one or more Linux guests access the DB2 instance on z/OS.

Characteristics of HiperSockets

HiperSockets have the following characteristics:

- They are used to configure the connection between two IBM Z or LinuxONE LPARs.
- It is a direct memory-to-memory pipeline and operates at memory transfer speed.
- The interconnection between Linux and z/OS is configured with Layer 3 HiperSockets.
- As with an OSA connection, a triplet of devices must be dedicated.
- Unlike an OSA connection, offloading is not possible.

2.2.11 z/VM cloud components

An enterprise cloud might be composed of various components, depending on what the main purpose of it is. Implementing an Infrastructure as a Service (IaaS) cloud in z/VM demands the integration of some important components. The components that are listed here play a role in the integration of a cloud in z/VM:

- **The z/VM Directory Manager (IBM DirMaint™)** or a supported equivalent provides a command driven interface to manage z/VM directory entries. For more information about DirMaint, see “DirMaint” on page 41.
- **The z/VM Systems Management Application Programming Interface (SMAPI)** provides programmatic access for managing many virtual images running on a single z/VM image by using a standard, platform-independent client interface, reducing the number of z/VM-specific programming skills that are required. It is already part of z/VM.
- **The z/VM Cloud Manager Appliance (CMA)** provides an easy method to deploy z/VM OpenStack enablement. OpenStack products and solutions can be constructed to use as many or as few of the services as is appropriate, whether that means that the CMA runs cloud controller services, compute node services, or only services that are needed by OpenStack z/VM drivers running in other virtual machines (VMs) or on other platforms.

z/VM Cloud Manager Appliance (CMA)

CMA allows the usage of OpenStack to deploy Linux guests on z/VM, and for the integration of z/VM into larger environments. The CMA version is upgraded to OpenStack Liberty for z/VM 6.3 and OpenStack Newton for z/VM 6.4, and is fully supported as a z/VM component without extra license costs.

CMA only manages z/VM platforms and it does not deploy guests onto non-z/VM platforms. The CMA changes provide several different options for using CMA, either as stand-alone cloud or integrated with another OpenStack environment.

Note: z/VM V6.3 and later natively supports OpenStack Cloud Manager Appliance (CMA).

A Cloud on z/VM environment might use several components and it is important to protect each of the them. It is important to have your company’s security policy job roles relate to the cloud, such as a cloud administrator and a cloud auditor.

Make sure the job roles that are related to the cloud also have their accesses described in the security policy, and that those accesses are implemented across the cloud environment.

For more information, see the following publications:

- *Enabling z/VM for OpenStack (Support for OpenStack Newton release)*, SC24-6253.
- Chapter 5 of *Securing your Cloud: IBM z/VM Security for IBM Z and LinuxOne*, SG24-8353.

2.2.12 Reasons to use the LinuxONE and IBM Z platforms together with z/VM

Using an infrastructure with the LinuxONE or IBM Z servers together with z/VM to virtualize the resources to Linux servers, you can accomplish the following tasks:

- Start small and scale up the servers
- Quickly deploy various workloads
- Consolidate servers
- Implement cloud services that can rapidly deploy a trusted, scalable, and OpenStack-based Linux cloud environment with mainframe qualities of service

z/VM uses and was designed to take advantage of all hardware functionalities of IBM Z and LinuxOne servers that provide explicit support for virtualization as SIE instruction that runs a virtual processor for a given virtual machine (VM). In other words, z/VM creates virtual machines adhering to z/Architecture of the IBM LinuxONE and IBM Z servers, which allows for a high level of trust. Furthermore, z/VM enables many such virtual machines to exist and operate simultaneously on one instance of real hardware and, in the most interesting cases, z/VM overcommits the real hardware without any extra financial or resource costs.

The heart of z/VM is a multi-programming, multi-processing operating system kernel known as the Control Program (CP). CP is the main component of z/VM that creates and dispatches virtual machines (VM) on the real hardware (the IBM LinuxONE and IBM Z servers).

Main reasons of using z/VM

There are some major advantages to z/VM:

- Infrastructure simplification:
  - Consolidate distributed, discrete servers and their networks
  - IBM Z and LinuxOne qualities of service
  - Use built-in z/VM system management
- Speed to Market:
  - Quickly deploy servers, networks, and solutions
  - React quickly to challenges and opportunities
  - Allocate server capacity when needed and dynamically
- Making the most of technology:
  - Total integration with the LinuxONE and IBM Z servers, efficiently using the z/Architecture
  - Linux guests with z/VM offer more functionalities than Linux servers alone
  - Linux uses unique z/VM technology features
  - Build innovative on-demand solutions

Note: z/VM continues to support more virtual machines in a single footprint with excellent service levels.

2.2.13 Server consolidation

Server consolidation increases the effective usage of a hardware server by enabling IBM LinuxOne or IBM Z server to host multiple virtual machines (VMs) by using z/VM. This ability can lead to lower power and cooling costs, and to greater computing efficiency and the flexibility to migrate workloads.

- Consolidate heterogeneous workloads into a single Linux server as a virtual machine (VM) running in z/VM or KVM.
- Consolidate multiple applications into a single physical server like IBM LinuxONE.

The most important step is to have the correct strategy to perform consolidation, virtualization or both. You need to undergo a deep analysis of the applications that you have:

- Are they I/O or processor intensive?
- Are they heavily accessed?
- Consider other, similar aspects of your app usage.

When this analysis is complete, you can decide to use one or both strategies to address your datacenter issues, lowering your costs and decreasing server sprawl, as shown in Figure 2-7.

Note: IBM has consolidated around 3900 distributed internal servers to 30 IBM Z servers with a consolidation called Project Green, in 2006.

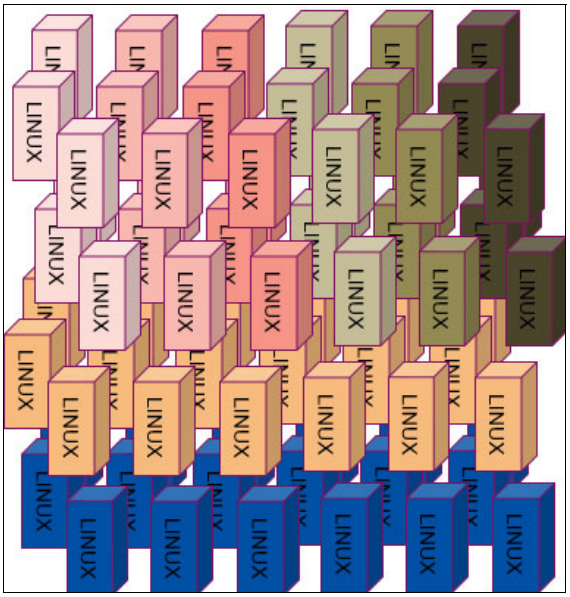

Figure 2-7 Server sprawl in distributed environment

Using z/VM clusters with SSI, with up to 4 z/VM LPARs connected, you can avoid planned outages with maintenance. Also, if you have some performance issues, you can move the Linux guests from one LPAR to another rapidly, in a matter of seconds. Doing so, you extend the application's availability with more flexibility, without any disruption.

You can consolidate your distributed servers in the virtual world of z/VM, as shown in Figure 2-8. TCO studies spanning more than 3 years repeatedly show the value of consolidation on z/VM, especially when the following conditions exist:

- A non-trivial number of servers is involved
- All operating expenses are included
- Software licensing is considered

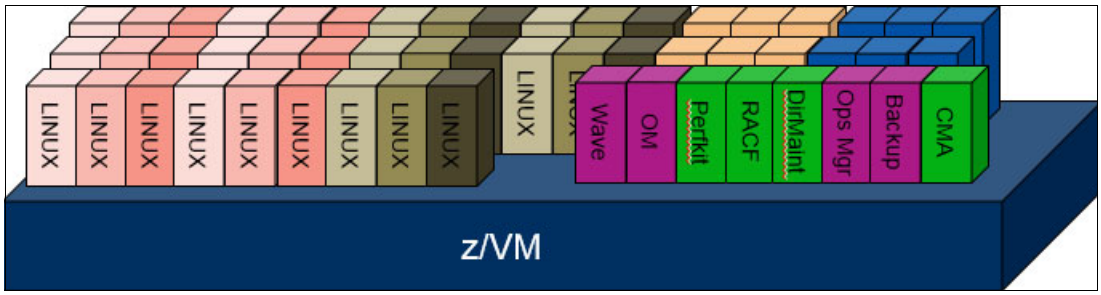

Figure 2-8 z/VM consolidation on the IBM LinuxONE and IBM Z servers

There is a compelling reason for consolidation. Software (from many vendors) can be expensive, and typically costs more than hardware. It is usually less expensive (and sometimes much less expensive) to replace multiple software licenses of smaller machines with 1 or 2 licenses for larger machines.

Note: Software license costs are often linked to the power of the system, but the pricing curves favor a few large machines.

For a complex environment, server consolidation helps reduce power consumption and cooling needs and reduces data center rack space requirements and server costs. It helps data centers better manage resources and resiliency. In addition, z/VM has a powerful mechanism to clone servers that enables existing servers to be cloned in a few minutes. The process leads to increased administration, system controls, and network complexity for your environment.

Lowering total cost of ownership (TCO) is another advantage of using a virtualized environment on the IBM LinuxONE and Z servers.

Think about the following costs:

- Hardware
- Software licensing
- Ongoing support
- Space usage
- Machine room running and cooling

Considering all these factors, the IBM Z and LinuxONE options deliver the overall lowest TCO for installations with a currently large investment composed of many distributed servers. The reliability, availability, and serviceability that are proven strengths of IBM Z and LinuxONE servers are additional positive factors.

2.3 Security

Security is essential in many ways. This is true for physical security, but because electronic services are more prevalent, it is evident that companies must secure and protect these services too.

2.3.1 Why security matters

Every company that handles customer information or offers services through Internet platforms must make sure that processed data is secured against all threats.

All precautions to prevent data leakage and to assure system and data integrity must be taken. It is no longer sufficient to state that data processing is secure today. You also must offer proof to auditors and comply with regulations to establish trust in your services.

Most important, you must prevent a loss of revenue and reputation due to security exposures. Therefore, it is a preferred practice to establish the strongest security mechanisms at all levels of data processing, including the physical security of the machine rooms at your data center (controlling access to the facilities) and implementing appropriate access levels to applications, programs, data, archives, and so on. The principle of least privilege should be met at all levels.

Protecting information from unintended use is one key element of a secure IT environment. Basically, there are two different methods to ensure privacy of information:

- Access control
- Encryption methods

Access control mechanisms determine who has the right to access particular information or data. The access control mechanisms then verify who accesses the information (authentication) and whether they have the right to access this information (authorization).

There are cases where proper access control cannot be ensured in all situations, especially if data is stored on movable media and also when data is transferred through a network that might not be protected. It is not possible to ensure that there is no unintended access to data while it is stored or transferred through a network. The only way to protect such information is by using encryption methods.

For more information about security, see *Security and Linux on IBM Z*, REDP-5464.

2.3.2 A brief overview of hardware security features

The hardware security features provide a fundamental part of the security definitions of software techniques and solutions, and each of the operating systems for IBM Z and LinuxOne servers (z/VM, Linux on IBM Z, z/VSE®, and z/OS) use these hardware features to some degree.

Despite being different classes of IBM hardware, IBM Z and LinuxONE both adhere to z/Architecture. Security features on the mainframe are integrated into the hardware. The following list provides some of the available hardware security features:

- With the Hardware Management Console (HMC), logical partitions (LPARs) can be defined and isolated from each other. Additionally, all the resources that are needed to run the operating systems are defined through LPAR profiles by the HMC. These resources are storage, processors, and Direct Access Storage Device (DASD) and tape units.

- Crypto Express Cards can encrypt both session traffic and physical data on DASDs and tape. For better performance, cryptographic coprocessors are used.
- Signed microcode is applied to the hardware to ensure microcode authenticity.

z/VM provides a host of features that isolates virtual machines (VMs) (also called guests) from one another. This isolation is implemented in the z/VM Control Program (CP), which can be considered the kernel of the hypervisor. Separation of guest workloads is a vital component of system integrity, and it provides the foundation of the security context on which the IBM Z Integrity Statement is based.

For more information about the z/VM CP, see *z/VM CP Planning and Administration,* SC24-6178.

2.3.3 z/VM Hypervisor

z/VM is considered a hypervisor operating system and its security does not differ from the security of any other operating system on a server. However, the virtual infrastructure relies on the security of the hypervisor, so protecting the z/VM hypervisor typically prevents attempts to breach the security of the operating system and compromises to the integrity of the operating system and data.

Although each guest can have its own security configuration and faces threats particular to it, it is essential to protect the hypervisor itself as an equally important part of an overall end-to-end security policy, because actions such as creating, changing, and removing virtual machines (VMs) are performed at the hypervisor level. Protecting the guests and not the hypervisor would be like locking all the windows to your home and then leaving the front door open. Access to the virtualization management system should be restricted to authorized administrators only.

Performing z/VM maintenance is part of the system administrator role. It is important to apply service to your z/VM system to ensure that the latest security measures are in place. Installing the corrections when they are released decreases the time frame that the vulnerability can be used.

Besides operating system setup and customization for security, monitoring the hypervisor for signs of compromise helps you promptly respond to a threat. Use monitoring tools to help monitor the hypervisor and look at the hypervisor logs for suspicious activities, both of which make the work of the hypervisor system administrator easier.

2.3.4 Security settings in an SSI cluster

A z/VM system is secured by using the security features of the IBM Z or LinuxONE hardware by maintaining compliance to security policy within operating practices. The system administrator must lead the way in following security standards and guidelines.

External security management

A preferred practice to extend the z/VM environment is installing an external security manager (ESM), such as IBM RACF® for z/VM or another ESM product to maximize your security. An ESM is an external software product designed to manage user identities and control access to system resources. When the Resource Access Control Facility (RACF) feature for z/VM is installed, it can be configured to control functions normally being checked in the directory for authorization. RACF can control the password field, the minidisk access, spool files, and commands privileges.

Enabling RACF/VM complements and extends basic z/VM security features while providing more granularity and accountability of all accesses and operations events. Its access control includes user verification, resource authorization, and logging capabilities. Enabling RACF/VM requires customization so that it can comply with your information security policy. Furthermore, RACF/VM allows for auditing of any security-relevant event on the system.

For more information about Security, you can see the following publications:

- *Securing your Cloud IBM z/VM Security for IBM Z and LinuxONE*, SG24-8353
- *Security and Linux on IBM Z*, REDP-5464
- *End to End Security with IBM Z*, REDP-5153

DirMaint

z/VM Directory Maintenance Feature (DirMaint) is a CMS application that helps manage your VM directory. Directory statements can be added, deleted, or altered using the DirMaint directory statement-like commands. DirMaint provides automated validation and disk allocation routines to reduce the chance of operator error. DirMaint is an interactive, multi-user application.

To add a user to z/VM, you must create a directory entry for a new virtual machine. The default method is through a manual process where you update a file that is called `USER DIRECT`, which is the z/VM system directory. The `USER DIRECT` file is a CMS file that contains the configuration values for each guest virtual machine as Linux servers. A virtual machine definition is a grouping of directory statements that begin with the term **`USER`** or **`IDENTITY`**.

For more details, a description of virtual machine types, and the **`USER`** and **`IDENTITY`** statements, see *z/VM Getting Started with Linux on System z*, SC24-6194.

Reasons for using DirMaint

The following list describes the high-level reasons for using DirMaint:

- Manually administering the user directory is needlessly complex, time-consuming, and error prone. Your time is better spent boarding new workloads quickly and optimizing performance versus counting disk cylinders manually and performing data entry.
- The user directory is a fundamentally critical part of z/VM; a corrupted or invalid online user directory can be disastrous. For example, if you inadvertently overlap minidisk definitions, it can cause serious and permanent data loss.
- If your z/VM system has more than a few virtual machines or belongs to a z/VM SSI cluster, it is illogical to attempt manual user management when automation exists.

DirMaint features

DirMaint provides the following features:

- Automatic disk management handles the management of minidisk extents.
- Error checking ensures that only valid changes are made to the user directory.
- Continuous synchronization of changes occurs to all member nodes in an SSI cluster.
- Change authorization permits only authorized personnel to make changes.
- Increased efficiency and productivity are possible through prototypes for Linux servers cloning purpose (and IBM FlashCopy®, if available).
- Control of all user-initiated transactions occurs through passwords.
- Logging transactions tracks changes, satisfies governance, and assists with auditing.

Smart economics: Non-disruptively scale your z/VM environment by adding hardware assets that can be shared with every virtual server.

2.3.5 Summary

It is widely known that IBM Z and LinuxONE are the platforms for consolidation, where multiple services are combined on the platform using its proven virtualization infrastructure. End-to-end solutions based on multi-platform workloads can be deployed across the IBM LinuxONE structure and benefit from IBM Z's traditional qualities of service, including high availability, and simplified and improved management of the virtualized infrastructure.

IBM z/VM uses the architecture of IBM LinuxOne and Z servers, and IBM is investing more each year in new z/VM releases and features to offer the most stable platform for Linux on IBM Z. z/VM and KVM provide flexibility, availability, and security capabilities for Linux instances, while it creates an isolated and protected environment for critical applications.

- Clients can start small with Linux on IBM Z or LinuxONE, and non-disruptively grow their environment as business dictates.
- Users can dynamically add CPUs, memory, I/O adapters, devices, and network cards to a running z/VM LPAR.

2.3.6 Dynamic synergy of virtualization in the LinuxONE and IBM Z platforms

Virtualization can help reduce the number of footprints by increasing server usage and reallocating resources, bringing value to customer challenges. It supports diverse operating systems and workloads with a high degree of isolation/security.

The LinuxONE and IBM Z servers are a large-scale, robust consolidation platform, and have built-in virtualization capability supporting hundreds of virtual Linux servers with z/VM or KVM. IBM is committed to those hypervisors and responsible for the architecture and use of the IBM Z and LinuxONE servers in the Linux upstream code.

IBM virtualization solutions address the need to increase usage of information assets, simplify the IT infrastructure, and reduce operating costs across servers, storage, networking, and grid computing. By extracting some administrative costs out of the infrastructure (through increased resource use and improved productivity and flexibility), these virtualized IT assets can help fuel business growth, control costs, and in doing so, increase staff productivity.

Factors that make the LinuxONE server compelling for a client

Data volumes, business demands, and the complexity of applications are increasing dramatically. Clients are looking for systems that provide security, speed, and scalability – on a server able to run today's leading open source software.

Note: The IBM LinuxONE and IBM Z platforms can deliver all the strength of the architecture in a single server or more, depending on the business requirements.

Organizations seek to improve the effectiveness of the IT infrastructure, consolidating workloads onto a single larger system becomes an attractive proposition. IBM Z and LinuxONE technologies are designed to enable the reduction in overall total cost of ownership (TCO) and increase in business flexibility to meet anticipated and unanticipated processing capacity demands with a more streamlined system infrastructure.

The more components an IT environment has, such as servers, cables, switches, and routers, the more difficult managing, identifying potential problems, and ensuring compliance become. IT complexity can drain budgets and hinder a company's ability to maximize the business value of its IT investment. For this reason, the LinuxONE and IBM Z servers were designed specifically to maximize resource use, minimize energy consumption, simplify virtualization management, and provide an economically attractive alternative to IT complexity

Note: Virtualization and consolidation are key for IT simplification.

Whatever you feel is best for you and your company, weigh the pros and cons of scaling up and scaling out, see the IBM LinuxONE and IBM Z servers versus benefits, and you might discover that this solution offers the best fit for your company's environment.

3

Scaling up on the LinuxOne and IBM Z platforms

This chapter describes the options and factors that need to be considered to make sound architectural decisions to deploy and scale up services with the LinuxOne and IBM Z applications and data.

The following topics are discussed here:

- Database scalability
- Performance
- Application workload
- Scale in with the LinuxONE and IBM Z platforms
- Final considerations

3.1 Database scalability

Database scalability is a common technique used to enable system administrators to grow existing systems to handle new business demands or growth of work. Details about scale out and scale up architectures can be found in 1.2, “Scale-up and scale-out definitions” on page 3.

The need for scaling up is covered in 1.1, “The x86 sprawl problems” on page 2, which “refers to the number of underutilized computer systems running in a data center. Usually, server sprawl is associated with typical x86 scale-out environments because they had an uncontrolled proliferation in past years, and also because x86 technology was the first choice for many organizations.”

Traditionally, databases were notorious culprits in hindering scalability, and were considered incapable of meeting the needs of a high-performance distributed computing design. However, databases have risen to the scalability challenge with new advancements in technology:

- The hardware that hosts the databases
- Networks (10 GigE, and so on)
- Disk access technologies (Fibre Channel, and so on)
- Advancements in multi-processing database management system (DBMS) technologies

The scale out of a database is illustrated in Figure 3-1. The database tier is scaled out by having a shadow database that contains log shipping capability to support reports, analysis, and so forth. As illustrated, the database complexity has increased and poses challenges in keeping the database cluster members synchronized.

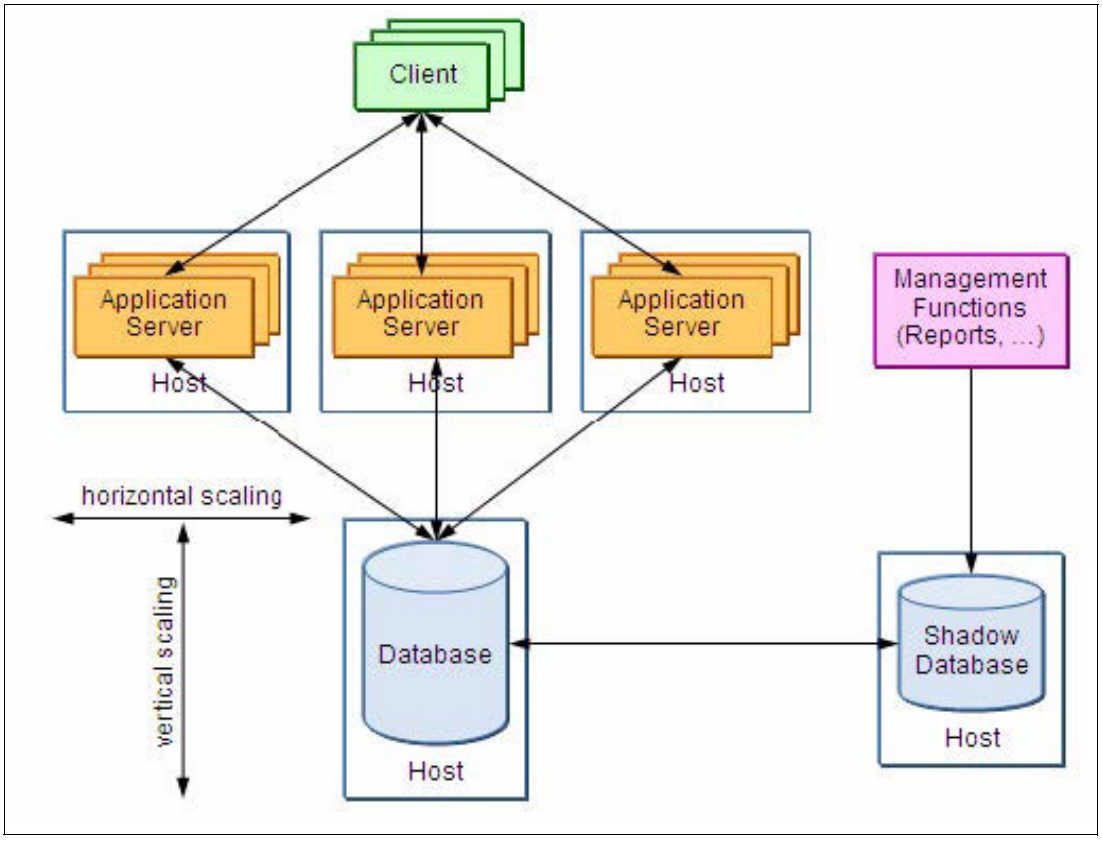

Figure 3-1 Database scalability challenge

Managing hundreds of database instances and the complex distributed architecture are other major challenges of the scale-out architecture. As discussed previously in this publication, the operational costs and the amount of effort to control this large environment tend to be more expensive and less efficient than the scale up architecture.

From the availability point of view, a scalability approach cannot lead to an application downtime. Otherwise, it can affect application users and sometimes, if the application is critical, the image of an organization. This point is important due to emerging applications, such as mobile and social media, which demand 24 x 7 availability.

Cloud applications introduced substantial resource demands, which can require large databases. So, they need more memory, CPU, and disk I/O to handle increased load. The fact is that the vertically scaled systems are the most cost-effective and the best-performing platform available for large, high-volume, transaction-intensive environments.

The LinuxONE and IBM Z platforms are designed to enable organizations to keep pace with this growth. Their hardware technology and open source support were improved, while at the same time shrinking the need for electricity and data center space when building your infrastructure with 1 or 2 LinuxONE or IBM Z servers.

To prevent you from needing to modify your existing applications and increase the environment complexity, the scaling up architecture brings several benefits over the distributed scaling out approach:

- Special Cluster software is not required
- Resources can be instantly available with a simple push of a button (added when needed)
- Less management
- No need to split your database in small pieces
- No application code changes
- Improve performance
- Lower infrastructure costs
- Eliminate x86 server sprawl

In summary, the operational requirements and complexity of scaling out systems are often far greater than scaling up architecture. The use of multiple systems to meet application needs is another major consideration to take into account when choosing scaling out architecture for your database or application. It is because the system will be responsible to process requests coming from different sources, and they must maintain the activities synchronized between a large group of servers. This results in increased complexity and higher costs.

Using IBM Z and LinuxONE scaling up architectures have significant advantages over other x86-based solutions, while helping to eliminate performance bottlenecks, environment complexity, and sprawl effects. In fact, database servers are one of the most highly recommended services to be moved to Linux on IBM Z.

The z/Architecture was designed with the concept of sharing. Sharing starts in the hardware components and ends with the data that is being used by the platform. The ability to share everything is based on one of the major strengths of the IBM Z system: virtualization.

As it is commonly used in computing systems, *virtualization* refers to the technique of hiding the physical characteristics of the computing resources from users of those resources. Virtualizing the IBM Z environment involves creating virtual systems (logical partitions and virtual machines), and assigning virtual resources (such as processors, memory, and I/O channels) to them. Resources can be dynamically added or removed from these logical partitions through operator commands, without the need to take the system down.

The IBM Z platform capabilities represent some of the most mature and sophisticated virtualization technologies in the industry today. For example, a single IBM Z system can scale up to millions of transactions per day or scale out to manage tens to hundreds of virtual servers. It can also redistribute system resources dynamically to manage varying server demands on the system resources automatically. The IBM Z servers can run the toughest workloads quickly and securely.

The following options are available when moving your databases from shared-nothing distributed data servers to the shared environment of the IBM Z platform:

- *Traditional* refers to a database that runs on a virtual machine.
- *Database as a Service* (DaaS) refers to a public or private cloud-based service offering databases on demand. The DaaS solution on LinuxONE and on IBM Z offers a flexible and scalable platform that enables you to provision and manage your database of choice quickly and easily.

Note: For information about installing the DaaS solution, see *Database as a Service*, SC34-2780.

Because organizations are focusing on keeping capital and operational expenses to a minimum, it is crucial to build a flexible data center infrastructure to accommodate the large databases.

There is no doubt that businesses can reap huge benefits from the LinuxONE and IBM Z platforms scaling up architectures to grow capacity and attend new business demands.

3.2 Performance

Cost and performance often are the main elements for selecting a scale architecture. So, organizations have to be ready for today's challenge but also for the bigger challenges of tomorrow by readying an infrastructure that can handle it. There are many ways to scale systems, but scale up is capturing the attention of decision-makers that want a solution to solve the scaling challenges and performance problems. The LinuxONE and IBM Z scale ups have a significant role in today's and tomorrow's data center and even has significant advantages over x86-based scale-out systems.

In fact, these architectures deliver the flexibility an organization needs to meet for their business demands and high-performance requirements, while also providing the ability to scale up or scale out later on as necessary. That means you have a choice today, and an alternative to the x86-based servers.

Today, the LinuxONE and IBM Z servers play a central role in the daily operations of most of the world's largest corporations, including many Fortune 1000 companies. Although other forms of computing are used extensively in business in various capacities, the IBM Z platform occupies a coveted place in today's e-business environment. In banking, finance, healthcare, insurance, utilities, government, and a multitude of other public and private enterprises, the z/Architecture continues to form the foundation of modern business.

Much of its popularity and longevity is due to its inherent richness in reliability and stability, a result of continuous technological advances since the introduction of the IBM System/360 in 1964. No other computer architecture in existence can claim as much continuous, evolutionary improvement, while maintaining compatibility with existing applications.

In the same way, the evolution of emerging technologies, such as cloud, big data, and blockchain, over the past years, creates a strong case for organizations looking to deploy open source derived software running on Linux into mission-critical environments.

In terms of hardware, the platform must easily scale up to improve performance and manage large masses of application data.

As mentioned, organizations across the world are increasing focus on their infrastructure to support these new initiatives. It continues to be a strong focus area in the overall strategic plan of most organizations across the world. Therefore, it is critical to explore the scaling, performance, and management of very large databases, complex applications, and the IBM Z platforms.

The IBM Z platform scale is designed to allow you to manage more workloads at less cost with increases in capacity per core and increases in single system capacity. The IBM Z platform delivers exceptionally fast application performance, millisecond response time, and extreme reliability thanks to the technical innovations.

Customers have expressed a strong interest in understanding how scale up solutions built on the IBM Z and LinuxONE platforms can be the answer to their growing requirements and performance problems. The configuration, scalability, and management solutions on the IBM Z platform create a robust and flexible infrastructure where organizations would like to see themselves.

In addition to that, the LinuxONE and IBM Z platforms are equipped with some of the fastest general-purpose processors in the world, ideally suited for data processing throughput. The large number of cores available and their high input/output bandwidth means that open source solutions can both scale up and scale out. They can also bring with it an exceptional total cost of ownership (TCO), highest performance, and enormous processor scalability. Indeed, the platform can scale up to 8,000 virtual machines or thousands of containers, which would be the most for any single x86-based server.

The concern around availability, performance, scalability, and security is a common one for mission-critical workloads, and a major differential factor when deciding the better infrastructure architecture to load vital applications.

The selection criteria between the x86-based platform, the IBM Z platform, and the LinuxONE platform is in their capabilities to handle data workloads such as databases, big data, data encryption, and certainly in analytics. Comparing these three platforms using complex workloads demonstrates a significant advantage for the LinuxONE and IBM Z servers.

3.3 Application workload

The new set of applications, such as analytics, mobile, and social, just to list a few, has shown beyond any doubt that there is a better, faster and easier way to build infrastructure for any type of workload.

Applications and data become common and they are increasingly hosted on cloud infrastructure. With an appropriate infrastructure, you can start small and grow within the server, deploy various workloads, consolidate servers, and implement cloud services that can rapidly deploy a trusted, scalable, and OpenStack-based Linux cloud environment with excellent qualities of service.

A key platform for the emerging digital services is the cloud. Many of the emerging digital applications and microservices that provide these innovative new capabilities are built to use the agility of cloud platforms, such as IBM Bluemix®. However, for the applications and services to function effectively, they must harness the application logic and business data that is held within enterprise systems, such as the LinuxONE and IBM Z platforms. The ability to use the IBM Z platform provides a powerful way to fuel the emerging digital applications by abstracting away the underlying infrastructure complexity.

Every computing platform offers specific areas of strength, and the aim of moving the workloads should be to select applications that take advantage of the strengths of the target platform. The classic strengths of the LinuxONE and IBM Z platforms include high availability, high I/O bandwidth capabilities, the flexibility to run disparate workloads concurrently, and excellent disaster recovery capabilities.

Another critical factor for choosing the LinuxONE and IBM Z platforms is the reduced total cost of ownership (TCO):

- Environmental savings: Single footprint versus hundreds of servers.
- Consolidation savings: Less storage, fewer servers, fewer software licenses, less server management/support.

These platforms were built to take optimum advantage of open source workloads that need higher levels of scalability, security, and availability from the underlying hardware infrastructure. Combined with the other factors already mentioned, the LinuxONE and IBM Z platforms can deliver more than x86 solutions. In particular, the ability to get your workloads processed at a lower overall cost and less complexity than the x86 scale-out architecture.

Each of the following sections contains information about how the LinuxONE and IBM Z platforms run more workloads with a consistent approach to improve service levels and meet demands with ease.

3.3.1 Cloud

Cloud is a method for delivering computing services that allows users to self-service provision environments with the agility that their businesses need. Systems running mission-critical applications are the backbone of an organization. When they go down, there is a major impact to the business.

As the computing world moves to the cloud, you need to decide the best scaling architecture to meet your needs and ensure that your cloud solution can benefit exceptional system uptime, excellent data security and privacy, and a powerful vertical scale architecture.

The following are some characteristics of cloud computing system:

- Rapid elasticity
- Broad network access
- Resource pooling
- Measured service
- On-demand self-service

Today it is imperative to reduce IT cost and respond quickly to business demands. Therefore, you need to consider the use of the IBM Z system scaling up architecture for large cloud data centers, environments running transaction-heavy databases, and other mission-critical applications. This architecture is more reliable and robust at the hardware level than any other platform, and fundamental to run mature and highly reliable solutions. Many advanced features, such as virtualization and logical partitioning for efficiency, make the IBM Z platform the correct choice for cloud.

The popularity of the x86 scale-out model can sometimes lead to the belief that cheap hardware and software can cut costs. In practice, this architecture can require more hardware, more cooling, more energy, and more networking infrastructure, resulting in increased costs. Additionally, it typically requires more skilled staff to handle the server operation and to lead with the complex clustering software. The effective scaling with reduced resource demands can be achieved through the IBM Z platform scaling up.

The traditional strengths of the LinuxONE and IBM Z platforms, such as architecture and software infrastructure, can offer a differentiated cloud platform to deploy and run your open source cloud-native workloads. The distributed servers deployed over time to run various application programs (scale out) can now be consolidated onto a single LinuxONE or IBM Z system to reduce operational costs, complexity, and administrative efforts. Consolidation also increases availability while eliminating the x86 server sprawl problems.

With the adoption of cloud, security remains a serious problem for most organizations today. Choosing what kind of security will be in place in the data center is an important decision, especially considering that information is one of the most valuable assets of an organization.

Data encryption is becoming common, and is a critical security method to protect and keep information secret. Organizations are seeking for solutions that ensure that their sensitive data in databases, virtualized platforms, and cloud environments are secure.

Some x86 options to accomplish data encryption solutions can dramatically degrade performance, be more complex, and are more expensive to manage. However, the LinuxONE and IBM Z platforms provide a way to encrypt every level of a network, from applications to local databases and cloud services. The persuasive encryption keeps data encrypted at all times unless it is being actively processed, and even then it is only briefly decrypted during those actual computations, before being encrypted again.

Note: For more information about the persuasive encryption, see *Security for Linux on System z*, SG24-7728.

The LinuxONE and IBM Z platforms encompass all of the various types of business software solutions, including database management software, such as IBM DB2. They can deliver infrastructure, such as web services, software, and platform services, to your entire organization.

Figure 3-2 shows the IBM LinuxONE platform ready for each layer of cloud.

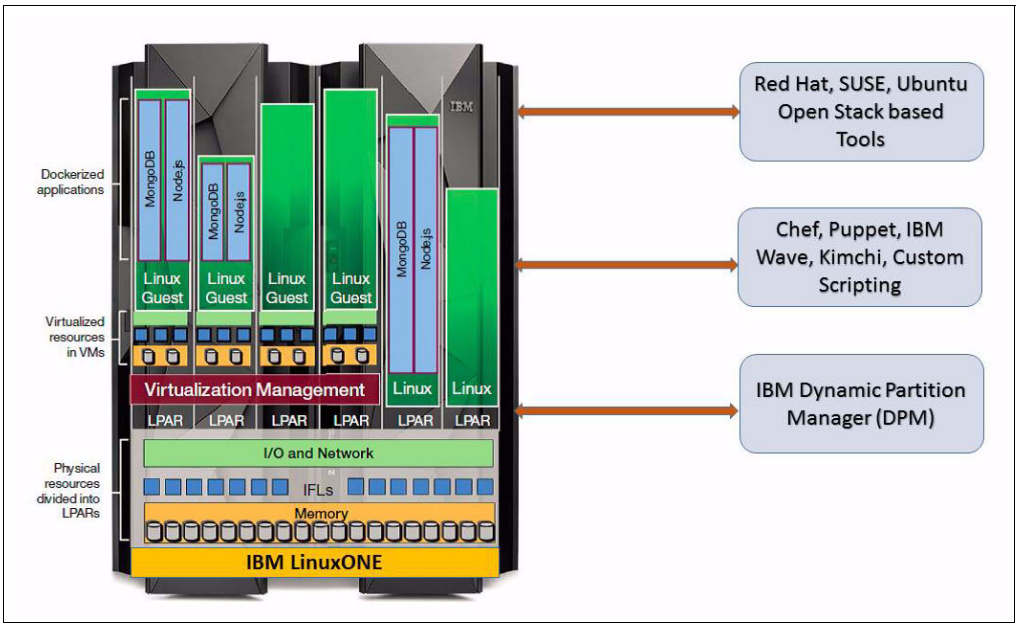

Figure 3-2 IBM LinuxONE is “cloud ready”

In summary, the LinuxONE and IBM Z platforms are designed to encrypt data, and have the best encryption performance, when compared to x86 solutions. Therefore, moving your applications onto Z can improve security with an entire application, cloud service, or database, in flight or at rest with one click.

The LinuxONE and IBM Z platforms enable those delivering cloud services to rapidly deploy a trusted, scalable OpenStack-based Linux cloud environment that can start small and scale up to 8000 virtual machines in a single footprint on any IBM LinuxONE or IBM Z family.

Note: For more information, see *Practical Migration from x86 to LinuxONE*, SG24-8377.

3.3.2 Blockchain

Blockchain is a technology for a new generation of transactional applications that fundamentally changed the way businesses created and captured value. Blockchain is a shared, distributed ledger that facilitates the process of recording transactions and tracking assets in a business network. Virtually anything of value can be tracked and traded on a blockchain network, reducing risk and cutting costs for all involved.

Because of the unique functionality and ability to handle the most sensitive data and critical applications, the LinuxONE and IBM Z platforms have the potential to run blockchain workloads better than any other platform. The IBM Z platform is optimized to deal with massive transactions and memory operations, enabling you to scale the system up to accommodate the new work while increasing security with the new cryptocard.

Blockchain can use special hardware in the IBM Z family and LinuxONE to deploy this disruptive technology at scale, with performance, availability, and security built-in.

Also, these platforms can deliver the following benefits:

- Blockchain peer-to-peer nodes realize optimized communication with z/OS, speeding up access to colocated business data
- Isolated partitions in memory keep ledgers separate and secure
- Availability and scalability of the IBM Z servers as an environment for both blockchain development/testing and production
- Vertical (scale up) scalability offers unmatched processing power
- Reduced data center footprint, simplified management, and energy savings
- Hardware encryption with built-in accelerators for blockchain hashing, signing, and security
- Faster responses with HiperSockets
- Global Security Standards compliant
- Tamper-proof crypto keys in firmware/crypto cards
- Unlimited random keys to encode transactions

If your organization is thinking about exploring blockchain technology, you might want to look into how the LinuxONE and IBM Z platforms can support your blockchain initiatives.

Additionally, you can explore blockchain technology by downloading and running the blockchain Z Docker image.

Note: You can find more details about Hyperledger installation in *Practical Migration from x86 to LinuxONE*, SG24-8377.

3.3.3 Analytics (IBM Watson)

IBM Watson is a technology to analyze and interpret all of your data, including unstructured text, images, audio, and video. With Watson, you can provide personalized recommendations by understanding a user's personality, tone, and emotion.

To analyze large amounts of structured and unstructured operational data, you need a powerful system environment. This cognitive infrastructure can be used to detect anomalous systems behavior in near real-time and then recommend corrective actions. Additionally, it can help to improve decision quality.

With all the benefits of the IBM Z platform, it makes sense to put your analytics initiatives into this platform. The reason is because this kind of application has specific heavy workloads that can take advantage of the IBM Z power to extract hidden value from large data volumes.

The IBM Z technology is designed to help propel cognitive workloads and to drive greater data center efficiency. Additionally, the system can be dynamically scaled up to accommodate workload peeks and handle structured and unstructured information that continues to grow rapidly.

Indeed, the ability to enhance applications with the IBM Z capabilities can vastly improve efficiency and performance without requiring hundreds of distributed servers to handle the application workload. The IBM Z platform is driving this efficiency by consolidating x86 workloads onto a single machine to save energy, cooling, and space.

You have more choices today, and you have more information to select a better system architecture to meet the data demands of the cognitive era with more flexibility than x86-based servers.

3.3.4 Mobile and social

The world is becoming more digitized and interconnected, which opens the door to emerging mobile applications. Analytics, mobile, social, and cloud computing all have one thing in common: They need a platform that has integrated security and can be flexible to easily scale up to meet growing demands.

These digital applications are more than just emerging technologies, because their capabilities are disrupting traditional business models. They also provide organizations with new sources of information that can lead to improved decision making and the potential for an increase in business results. Existing core applications on the LinuxONE and IBM Z platforms, such as transactional processing or database services, are often characterized as secure, reliable, and highly scalable.

These properties are also important for a mobile application. In addition to the application requirements, the challenges for a mobile workload are the expected fast response times and the use of interfaces on the mobile device, based on touchscreens and voice.

These applications have a high rate of growth and a large number of active users. So, the system architecture plays an important role to ensure excellent levels of performance and elastic scalability to meet unexpected surges in mobile and social requests.

The data that these applications require access to must be available continuously. It needs to be secure at the source, and it must be transmitted rapidly and securely. These factors are what make the LinuxONE and IBM Z platforms the primary choice for mobile and social architecture.

The platform offers everything in availability, scalability, and accessibility. The platform also provides the domain of an extensive set of tools for development, and Business Analytics can be integrated into any chosen solution.

The IBM Z platform has evolved over the years to process exactly that situation, not only in a single instance, but by many thousands of requests, running concurrently.

The IBM Z transaction processing systems have been responsible for the control of this huge repository of data for some time. The IBM Z platform offers the ability to operate continuously, with secure access being provided by the latest technological advances in cryptography and digital signatures. Therefore, building a mobile and social solution onto the IBM Z infrastructure makes perfect sense.

The IBM Z platform has a renowned capability to have a linear and automated scalability, and can handle the unpredictable user workload. Additionally, high availability and disaster recovery are further requirements for every organization that embraces social and mobile applications. Implementing and maintaining these requirements has always been one of the strongest capabilities of an IBM Z environment. The IBM Z platform can help you deliver an IT infrastructure that can keep pace with the increased transaction workload that results from mobile and social engagements, while ensuring the highest levels of protection and privacy.

For mobile and social application transactions, end-to-end security is paramount. It is crucial for the overall transactional confidentiality, integrity, and auditability of the requests initiated by a mobile device.

Understanding all of the advantages available between the IBM Z and x86-based platforms can help you to select the appropriate system architecture that can better deal with the challenges of these emerging digital applications.

3.3.5 Docker

Docker is gaining popularity among many businesses, enabling them to package an application software and all its dependencies. Now, an application can be moved between environments and run without needing an application update. As virtualization provided efficiencies, capabilities, reduced costs, and other benefits a few years ago, containers are significantly improving the way applications are developed and deployed across the IT infrastructure.

As discussed earlier in this publication, the LinuxONE and IBM Z servers are prepared to handle business-critical applications. And now they can provide a standard platform to run your Docker instances. These systems can support up to two million Docker containers on a single machine, maintaining the highest levels of security and performance and not requiring application server farms.

Additionally, they can facilitate transparent use of redundant processor execution steps and integrity checking, which are necessary in financial services industries. The LinuxONE and IBM Z platforms typically enable hot-swapping of hardware, such as processors and memory. This swapping is typically transparent to the operating system, enabling routine repairs to be performed without shutting down the system.

Figure 3-3 illustrates the differences between server virtualization and Docker.

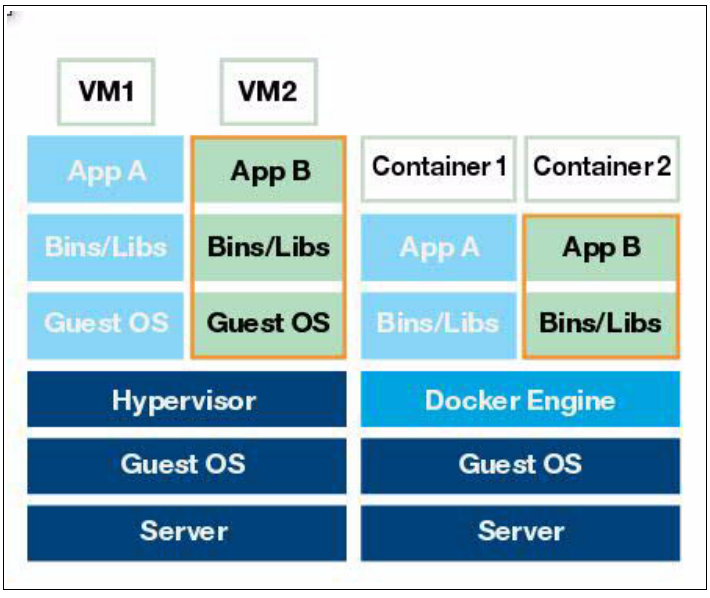

Figure 3-3 Server Virtualization versus Docker

Organizations are struggling to find ways to lower costs, keep systems and data secure, and provision applications with an increasing number of moving parts.

Virtual machines do a great job at abstracting from the underlying hardware. This lowers costs and makes it possible to automate provisioning of a complete software stack, including the operating system, the application, and all its dependencies. However, virtual machines are not ideal for every use case.

There are many ways to increase application flexibility. In fact, the easiest way is to create a virtual machine and then deploy it to run on this specific environment. However, this approach is not efficient and makes the scalability of several virtual machines impossible to do for microservices architectures.

Lifecycle management, of course, is one of the areas in which Docker really stands out from server virtualization, due to both the nature of how Docker rolls out the applications, and to the excellent advantage to remove the administrative resources for patching and lifecycle management that are required for every virtual machine.

Docker, an open source technology, was built to make it easier to create, deploy, and run applications by using Linux containers. From a system resources point of view, Docker containers are much more efficient than hypervisors, and leave your application small and neat.

Apparently, Docker looks like a virtual machine. However, unlike a virtual machine, rather than create a whole virtual operating system, Docker enables applications to use the same Linux kernel as the system that they're running on, and only requires applications to be sent with parts not already running on the host computer.

Docker containers provide you application portability, simplify security, and help achieve greater application mobility. That is the reason some organizations are choosing this technology to bring standardization, reduce costs, and speed up deployments for their environments.

As Docker and containers continue to gain attention and adoption in many organizations, It makes sense to provide details about how Docker with the LinuxONE and IBM Z platforms can take advantage of the Z scaling up architecture to increase efficiency with reduced costs.

Virtualization

On the LinuxONE and IBM Z servers, there is an elaborate system of virtualization:

- Multilevel virtualization, composed of logical partitions (LPARs)
- IBM z/VM, which is a tried and tested hypervisor technology
- In the future, the open source kernel-based virtual machine (KVM) for Linux on IBM Z

Docker is an ideal addition to those virtualization approaches because, within the virtual entity, you can set up your application landscape, which is composed of multiple services. You are able to provision the same kind of applications using the same kind of deployment paradigm on the LinuxONE and IBM Z servers that you do on an x86 or other architecture, yet retain the sophisticated virtualization management used by many of the IBM Z clients.

Application portability

Docker facilitates the use of all of the applications that are now on distributed platforms, including the LinuxONE and IBM Z servers. The same packaging structure makes it simpler to deploy the IBM Z servers edition of the application in a container and run it on the LinuxONE or IBM Z servers. Eventually, it will enable the z/Architecture to also play a role in environments that are dominated by and only considered for the distributed space today.

Consolidation

It is widely known that the LinuxONE and IBM Z platforms are the best for consolidation, where multiple services are combined on the platform using its proven virtualization infrastructure. The containers used by Docker can accommodate larger density than virtual machines and enable more applications in one system, which further expands the consolidation capabilities of Z servers by making it possible to introduce new types of applications.

On the LinuxONE and IBM Z platforms, you can shape your environment with system virtualization and Docker containers according to your landscape and requirements, without performance constraints.

Security

IBM understands the requirements of enterprise organizations as well as public sector and government entities with mission-critical environments for security and isolation. LPAR technology provides the highest non-military certification of virtualization isolation. In addition, the z/VM hypervisor also provides an extremely high level of isolation, and likewise for KVM on IBM Z. Security and isolation of workloads are often critical concerns to large enterprise clients. We are bringing this high level of security into the partnership with Docker.

From our many client interactions, we know well that when you get into production in an enterprise environment, virtualization suddenly plays a major role and workload isolation becomes an issue. Docker running in virtualized environments provides workload and tenant isolation with mature and sophisticated management combining simple deployment with enterprise-grade compute environments.

3.3.6 MongoDB

MongoDB is a NoSQL document-oriented database. It stands out from relational databases by using dynamic schemas, with which records can be inserted without creating an initial schema to define data structure. Therefore, fields and their values can be easily modified to map application changes without interruption.

This OSDBM is widely used for mobile apps, real-time analytics, product catalogs, and content management systems. It can also be used for many other use cases, such as storing streams of data from IOT.

MongoDB stores data in a JSON-like documents (aligning data storage formats with modern programming languages that are used by developers) instead of columns and rows in relational databases. Each document can contain different fields, which contain a value that belongs to the same data type including subdocuments and arrays.

Instead of having related data that is represented by different tables, the cost of joining separate tables is eliminated by storing the linked objects in the same document. This configuration reduces complexity and simplifies data access. Similar documents are organized into collections in the database.

Figure 3-4 shows the difference between MongoDB documents and traditional relational data models.

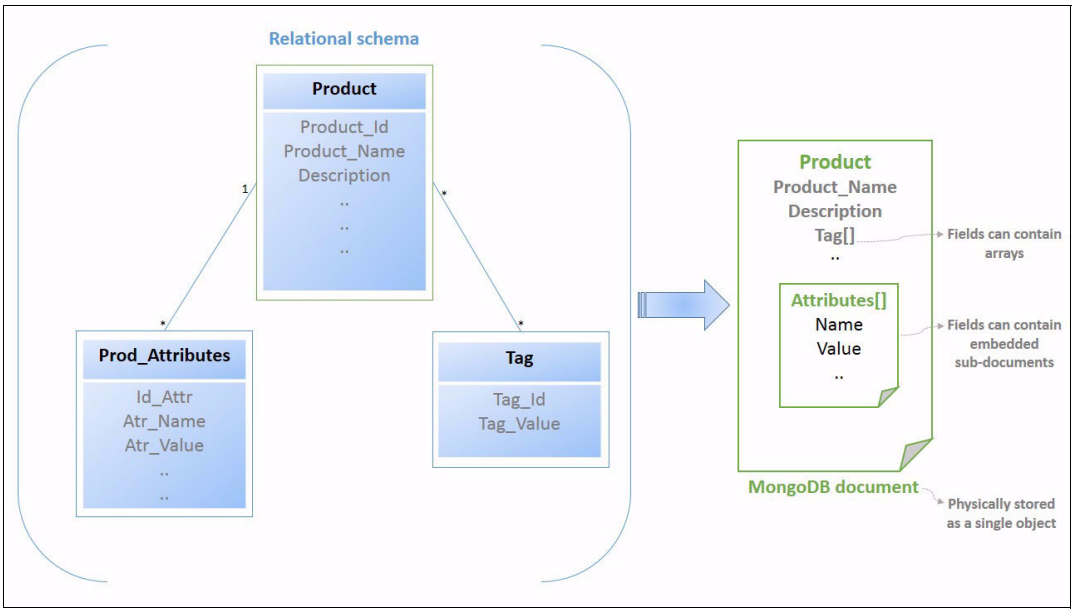

Figure 3-4 MongoDB document model versus relational data model

Table 3-1 lists some of the concepts that are used in MongoDB and their counterparts in relational databases; for example, DB2 database.

Table 3-1 MongoDB versus a traditional SQL database

MongoDB	DB2
JSON-like document	Row
Collection	Table
Embedded subdocuments and Linking	Join
Index	Index

Added value of running MongoDB on LinuxONE

With the LinuxONE platform, many cores can be used (scale 1 - 170 cores), which allows MongoDB to scale vertically (up to 17 TB single-node MongoDB) without the use of the sharding technique, which is also known as horizontal scaling (or scale out), to split a large amount of data across different servers. This method is the best way to remove the resource use required by sharding, and to reduce the complexity of managing different MongoDB shards.

The network communication between the MongoDB shards is also drastically reduced (near zero latency) by using IBM HiperSockets technology that is used in LinuxONE and IBM Z servers to accelerate the communication between the different LPARs.

The enterprise edition of MongoDB (called MongoDB Enterprise Advanced) is supported on LinuxONE and IBM Z.

In addition, the capabilities of the LinuxONE and IBM Z platforms in terms of security (EAL5+ certification,) high availability, and resiliency can be used to efficiently enhance MongoDB's capabilities and achieve the most stringent market requirements.

The EAL5+ (Evaluation Assurance Level certification) isolation certification means that a single Z machine can create up to 85 logical partitions, each with correct separation and isolation.

3.4 Scale in with the LinuxONE and IBM Z platforms

The LinuxONE and IBM Z platforms have proven to be utterly reliable to handle large amounts of data and mission-critical applications, more efficiently than any other platform. These systems were specifically designed to accommodate applications with heavy on storage or heavy on compute. They can respond to a sudden surge in demand extremely quickly, preventing complex applications from having slow responses times.

In the emergence of a new software platform, applications such as cloud-based, MongoDB, and Docker prove that the LinuxONE and IBM Z platforms can run open source software with excellent flexibility, security, and scalability. The reason is that the IBM Z platform combines the best of both scaling worlds and bring with them effective costs and availability advantages.

The architecture is flexible enough to allow applications running on cloud services to easily scale up in whatever increments, while offering high-performance for applications requiring processing power. Also, it addresses application scale out by adding new virtual servers within a few minutes and make them part of the cluster.

Note: For more details about virtualization, see “Software virtualization (hypervisors)” on page 28.

Scale out or scale up alone do not adequately serve today’s applications, which are responsive, always available, and can service many users, including cloud, mobile, analytics, and big data. That makes Z the best platform to meet different demands, from small to complex. In summary, the platform remains the powerhouse for tasks that computers have performed for decades and technologically redesigned to attend emerging demands.

The following scaling information refers to performance tests performed by IBM on a LinuxONE server.

Scale out

Scale out has the following capabilities:

- Run 1344 concurrent databases executing a total of 377 billion database transactions per day on a single LinuxONE Emperor II server.
- Run 25% more MongoDB guests with the same throughput under z/VM 6.4 on the LinuxONE Emperor II server compared to the LinuxONE Emperor server.
- Use up to 170 cores on the LinuxONE Emperor II server to scale-out MongoDB databases under z/VM 6.4, each with a constant throughput and not more than 10 µs latency increase per additional MongoDB instance.
- Scale-out to 2 million Docker containers in a single LinuxONE Emperor II system, no application server farms necessary.
- Run 41.8 billion web transactions per day on a single Emperor II server.

Scale up

Scale up has the following capabilities:

- Scale up a single MongoDB instance to 17 TB in a single system without database sharding and get 2.4 times more throughput and 2.3 times lower latency on a LinuxONE Emperor II server leveraging the additional memory available compared to a LinuxONE Emperor server.
- Run MongoDB under z/VM 6.4 on a LinuxONE Emperor II server and get 4.8 times better performance leveraging additional memory available per z/VM instance compared to a LinuxONE Emperor server.

3.5 Final considerations

In a data-driven world, your infrastructure needs flexibility and stability to run a wide range of open source products next to Oracle databases or DB2 databases. Organizations are looking for solutions with constant innovation that stems from open source world and incurs lower costs.

Many organizations are under the false impression that the LinuxONE and IBM Z servers are servers that are accompanied by higher overall software, hardware, and people costs. Most organizations do not accurately calculate the total costs of their server proliferation, largely because chargeback mechanisms do not exist, because only the incremental IBM Z investment costs are compared to incremental distributed costs, or because total shadow costs are not included.

Many organizations also fail to recognize the path length delays and context switching of running workloads across many servers, which typically adds up to a performance penalty that is non-existent on the IBM Z platform. Also, the autonomic capabilities of the LinuxONE and IBM Z platforms (reliability, scalability, and self-managing design) might not be considered.

Distributed servers encounter an efficiency barrier where adding incremental servers after a certain point fails to add efficiency. The total diluted cost of the LinuxONE and IBM Z platforms is not used correctly in calculations; rather, the delta costs attributed to an added workload often make the comparisons erroneous. In distributed servers, the cost per unit of work never approximates the incremental cost of a mainframe.

However, over time, it is unlikely that a server farm could achieve the economies of scale associated with a fully loaded LinuxONE or IBM Z server, regardless of how many devices are added. In effect, there is a limit to the efficiencies realizable in a distributed computing environment. These inefficiencies are due to shadow costs, execution of only one style of workload versus a balanced workload, underutilization of CPUs, people expenses, and real estate cost of a distributed operations management.

Following this further, server sprawl and x86 distributed systems can cost you in more ways than one. Cash, complexity, and capacity can all limit your ability to say “yes” to the business. Consolidating your servers to a LinuxONE and IBM Z server gives you the competitive edge so that you do not have to make either choice for your data center.

In particular, the value that the IBM Z scaling up architecture brings to running Linux applications on a secure and reliable platform makes it the most efficient and economical environment when compared with the x86-based scale-out approach. It means taking advantage of the power of the IBM Z platform to run multiple workloads and open source applications giving you a cost-effective alternative to x86 distributed servers.

Scaling up is certainly easier, and most widely used applications, such as MongoDB, analytics, big data, and so on, lend themselves well to that approach. In fact, these applications were designed for scale up to make the system bigger.

Organizations have a preference for open source software, an environment that is flexible, and cost-effective solutions. These aspects help you decide to use the LinuxONE and IBM Z servers that were engineered for the industry's most secure data serving, with faster throughput to get more from the open source infrastructure, databases, and heavy applications.

Table 3-2 shows some technical enhancements in the IBM Z (z14) platform over its predecessor platforms, as also shown in Figure 3-5 on page 62. The IBM Z platform delivers unique functionality and scale to meet all of the challenges of modern deployment for mission-critical applications.

Table 3-2 Technical highlights

<table>
<tr><td>Vertical scaling up to 170 user-accessible cores running at 5.2 GHz.</td><td>32 TB of shared memory enabling larger in-memory applications.</td><td>Multi-core, single-chip modules running to help improve the running of processor-intensive workloads.</td></tr>
<tr><td colspan="3">More real memory per system, ensuring high availability in the memory subsystem through use of proven redundant array of independent memory (RAIM) technology.</td></tr>
<tr><td colspan="3">Improved cryptographic functions and performance, achieved by having one dedicated cryptographic co-processor per processor unit.</td></tr>
<tr><td colspan="3">Proven technology (fifth-generation high frequency a nd third-generation out-of-order design) with a single-instruction, multiple-data (SIMD) processor that increases parallelism to accelerate analytics processing. In addition, simultaneous multithreading (SMT) increases processing efficiency and throughput and raises the number of instructions in flight.</td></tr>
<tr><td colspan="2">The channel subsystem is built for I/O resilience. The number of logical channel subsystems (LCSS), subchannel sets, and I/O devices are consistent with its predecessor platform, as is the number of logical partitions (LPARs).</td><td>► Six LCSS
► 85 LPARs
► Four subchannel sets
► 32,000 I/O devices per channel</td></tr>
</table>

Note: You can find more details about the z14 platform in *IBM z14 Technical Introduction*, SG24-8450. For an in-depth description of the IBM z14 functions and features, also see *IBM z14 Technical Guide*, SG24-8451.

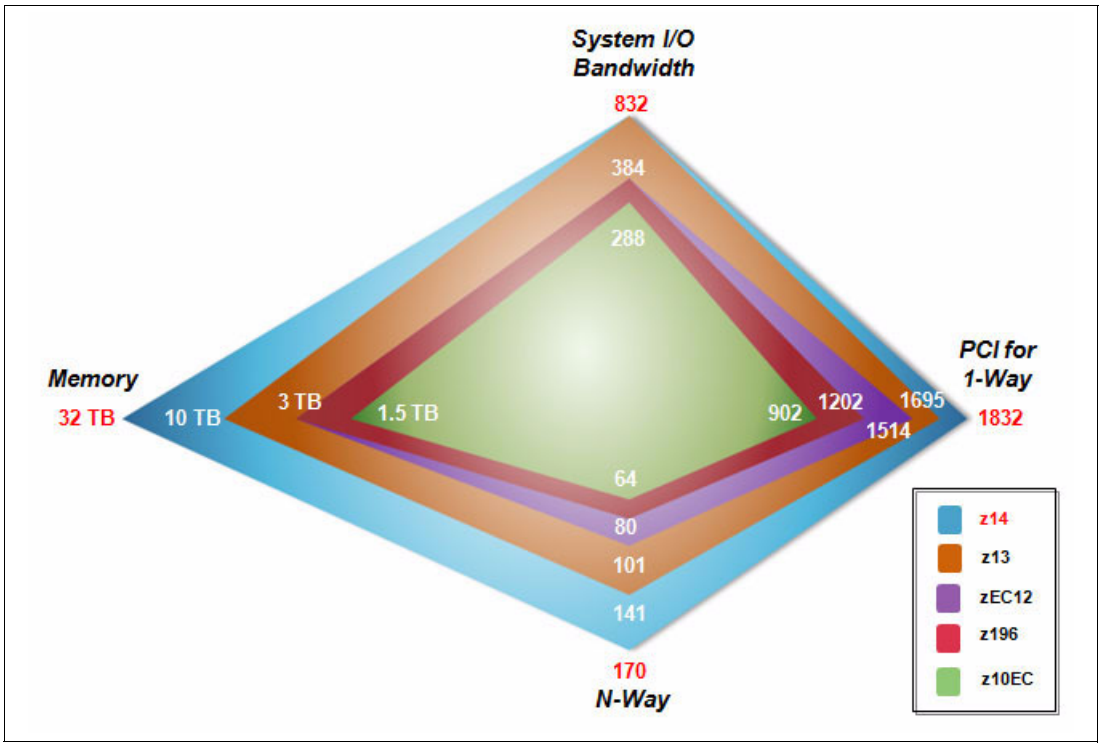

Figure 3-5 Balanced system design: z14 versus its predecessors

The latest LinuxONE and IBM Z servers have a tried-and-true architecture to support your digital transformation, create a strong cloud infrastructure, and make back-end services available through secure APIs. They can also streamline your ability to integrate disparate data center systems and create a single, cohesive IT shop.

The infrastructure built on the LinuxONE and IBM Z platforms provides the following benefits:

- Less data center space
- Less power consumption
- Fewer administrators
- More efficient: Can be scale up or scale out
- Fewer servers: Reduced Complexity
- Fewer software licenses
- Fewer intrusion points: Tighter security
- Fewer points of failure: Greater availability

By choosing the LinuxONE and IBM Z platforms, organizations can benefit from a secure, highly available, and easy-to-manage environment to consolidate their data while optimizing overall costs.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Getting Started with KVM for IBM z Systems*, SG24-8332
- *An introduction to z/VM Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8006
- *Using z/VM 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039
- *Practical Migration from x86 to LinuxONE*, SG24-8377
- *IBM z14 Technical Introduction*, SG24-8450
- *IBM z14 Technical Guide*, SG24-8451
- *Securing your Cloud IBM z/VM Security for IBM Z and LinuxONE*, SG24-8353
- *Security for Linux on System z*, SG24-7728
- *End to End Security with IBM Z*, REDP-5153
- *Security and Linux on z Systems*, REDP-5464

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *Database as a Service*, SC34-2780
- *z/VM CP Planning and Administration version 6 release 4*, SC24-6178

Online resources

These websites are also relevant as further information sources:

- https://www.ibm.com/linuxone/solutions
- https://www.redbooks.ibm.com/Redbooks.nsf/domains/linuxone?Open&page=trending
- https://www.ibm.com/software/info/eagletco/iteconomics/
- https://www.ibm.com/software/passportadvantage/about_software_licensing.html
- http://www.ibm.com/software/sla/sladb.nsf/sla/bla/
- https://www.ibm.com/it-infrastructure/z/software/pricing-resources
- https://www.ibm.com/it-infrastructure/z/software/pricing
- ftp://ftp.software.ibm.com/software/passportadvantage/PA_Agreements/PA_Agreement_International_English.pdf
- https://www.ibm.com/linuxone
- https://www.ibm.com/servers/resourcelink/svc03100.nsf?Opendatabase
- http://www.vm.ibm.com/

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5461-00

ISBN 073845656X

Printed in U.S.A.