

Running Mainframe Z On Distributed Platforms How To Create Robust Cost Efficient Multiplatform Z Environments

IBM® CICS® Transaction Server (CICS TS) has been available in various guises for over 40 years, and continues to be one of the most widely used pieces of commercial software. This IBM Redbooks® publication helps application architects discover the value of CICS Transaction Server to their business. This book can help architects understand the value and capabilities of CICS Transaction Server and the CICS tools portfolio. The book also provides detailed guidance on the leading practices for designing and integrating CICS applications within an enterprise, and the patterns and techniques you can use to create CICS systems that provide the qualities of service that your business requires.

LinuxONE is a portfolio of hardware, software, and solutions for an enterprise-grade Linux environment. It has been designed to run more transactions faster and with more security and reliability specifically for the open community. It fully embraces open source-based technology. Two servers are available for LinuxONE: The IBM® LinuxONE III LT1 and IBM LinuxONE III LT2. We describe these servers in "IBM LinuxONE servers" on page 5. Aside from still running SUSE Linux Enterprise Server and Red Hat Enterprise Linux Servers, LinuxONE runs Ubuntu, which is popular on x86 hardware. Ubuntu, which runs the cloud, smartphones, a computer that can remote control a planetary rover for NASA, many market-leading companies, and the Internet of Things, is now available on IBM LinuxONE servers. Together, these two technology communities deliver the perfect environment for cloud and DevOps. Ubuntu 16.04 on LinuxONE offers developers, enterprises, and Cloud Service Providers a scalable and secure platform for next generation applications that include OpenStack, KVM, Docker, and JuJu. The following are reasons why you would want to optimize your servers through virtualization using LinuxONE: Too many distributed physical servers with low utilization A lengthy provisioning process that delays the implementation of new applications Limitations in data center power and floor space High total cost of ownership (TCO) Difficulty allocating processing power for a dynamic environment This IBM Redbooks® publication provides a technical planning reference for IT organizations that are considering a migration from their x86 distributed servers to LinuxONE. This book walks you through some of the important considerations and planning issues that you might encounter during a migration project. Within the context of a pre-existing UNIX based or x86 environment, it presents an end-to-end view of the technical challenges and methods necessary to complete a successful migration to LinuxONE.

Today, organizations are responding to market demands and regulatory requirements faster than ever by extending their applications and data to new digital applications. This drive to deliver new functions at speed has paved the way for a huge growth in cloud-native applications, hosted in both public and private cloud infrastructures. Leading organizations are now exploiting the best of both worlds by combining their traditional enterprise IT with cloud. This hybrid cloud approach places new requirements on the integration architectures needed to bring these two worlds together. One of the largest providers of application logic and data services in enterprises today is IBM Z, making it a critical service provider in a hybrid cloud architecture. The primary goal of this IBM Redpaper publication is to help IT architects choose between the different application integration architectures that can be used for hybrid integration with IBM Z, including REST APIs, messaging, and event streams.

In this IBM® Redbooks® publication, we describe the role Cognos® plays in an Information On Demand (IOD) solution for IBM System z® and detail the functions of IBM Cognos 8 BI for Linux® on System z in current deployment scenarios. We show typical deployment architectures that show how to access disparate data sources both on and off the System z platform and show how the functions of the Cognos family of products provides a way to consolidate different BI solutions on System z. We provide examples of Cognos functions for resolving business requirements using reporting and OLAP capabilities as well as general deployment considerations of IBM Cognos 8 BI for Linux on System z. This publication is meant to help the Cognos Business Intelligence professional understand the strong points of System z architecture and the database specialist appreciate the Cognos family of products.

System administration is about the design, running and maintenance of human-computer systems. Examples of human-computer systems include business enterprises, service institutions and any extensive machinery that is operated by, or interacts with human beings. System administration is often thought of as the technological side of a system: the architecture, construction and optimization of the collaborating parts, but it also occasionally touches on softer factors such as user assistance (help desks), ethical considerations in deploying a system, and the larger implications of its design for others who come into contact with it. This book summarizes the state of research and practice in this emerging field of network and system administration, in an anthology of chapters written by the top academics in the field. The authors include members of the IST-EMANICS Network of Excellence in Network Management. This book will be a valuable reference work for researchers and senior system managers wanting to understand the essentials of system administration, whether in practical application of a data center or in the design of new systems and data centers. - Covers data center planning and design - Discusses configuration management - Illustrates business modeling and system administration - Provides the latest theoretical developments

Achieve streamlined, rapid production with enterprise-level DevOps Awarded DevOps 2017 Book of the Year, The DevOps Adoption Playbook provides practical, actionable, real-world guidance on implementing DevOps at enterprise scale. Author Sanjeev Sharma heads the DevOps practice for IBM; in this book, he provides unique guidance and insight on implementing DevOps at large organizations. Most DevOps literature is aimed at startups, but enterprises have unique needs, capabilities, limitations, and challenges; "DevOps for startups" doesn't work at this scale, but the DevOps paradigm can revolutionize enterprise IT. Deliver high-value applications and systems with velocity and agility by adopting the necessary practices, automation tools, and organizational and cultural changes that lead to innovation through rapid experimentation. Speed is an advantage in the face of competition, but it must never come at the expense of quality; DevOps allows your organization to keep both by intersecting development, quality assurance, and operations. Enterprise-level DevOps comes with its own set of challenges, but this book shows you just how easily they are overcome. With a slight shift in perspective, your organization can stay ahead of the competition while keeping costs, risks, and quality under control. Grasp the full extent of the DevOps impact on IT organizations Achieve high-value innovation and optimization with low cost and risk Exceed traditional business goals with higher product release efficiency Implement DevOps in large-scale enterprise IT environments DevOps has been one of IT's hottest trends for the past decade, and plenty of success stories testify to its effectiveness in organizations of any size, industry, or level of IT maturity, all around the world. The DevOps Adoption Playbook shows you how to get your organization on board so you

can slip production into the fast lane and innovate your way to the top.

Targeted at management, the first six chapters of Secrets of SOA focus on the business impact of service-oriented architecture technological decisions with an emphasis on cost, flexibility, and the ability to maintain business objectives. Each of the six chapters explores a different topic that illustrates the value of a physically integrated SOA infrastructure organized at the enterprise level. Taken together, they demonstrate why enterprise-level planning, backed by a centralized deployment strategy, is essential to the success of SOA. Aimed at the IT executive, the second half of the book deals with specific IT issues raised by SOAs and why these issues are best dealt with on an enterprise level. Among the topics covered in these eight chapters are virtualizing resources, managing heterogeneous workloads, maintaining data and transactional integrity, and the value of proximity.

Linux on System z offers many advantages to customers who rely on the IBM® mainframe systems to run their businesses. Linux on System z makes use of the qualities of service in the System z® hardware and in z/VM®, making it a robust industrial strength Linux. This provides an excellent platform for hosting Oracle solutions that run in your enterprise. This IBM Redbooks® publication is divided into several sections to share the following experiences that are gained while Oracle Database 11gR2 is installed and tested: Setting up Red Hat Enterprise Linux 6 for Oracle Managing an Oracle on Linux on System z environment Provisioning Linux guests using several tools It also includes many general hints and tips for running Oracle products on IBM System z with Linux and z/VM. Interested readers include database consultants, installers, administrators, and system programmers. This book is not meant to replace Oracle documentation but to supplement it with our experiences while Oracle products are installed and used.

The ABCs of IBM® z/OS® System Programming is a 13-volume collection that provides an introduction to the z/OS operating system and the hardware architecture. Whether you are a beginner or an experienced system programmer, the ABCs collection provides the information that you need to start your research into z/OS and related subjects. Whether you want to become more familiar with z/OS in your current environment, or you are evaluating platforms to consolidate your online business applications, the ABCs collection will serve as a powerful technical tool. Volume 1 provides an updated understanding of the software and IBM zSeries architecture, and explains how it is used together with the z/OS operating system. This includes the main components of z/OS needed to customize and install the z/OS operating system. This edition has been significantly updated and revised.

This book will teach you the basic information and skills you need to develop applications on IBM mainframes running z/OS. The instruction, examples and sample programs in this book are a fast track to becoming productive as quickly as possible in VSAM, IMS DB and intermediate level DB2. The content is easy to read and digest, well organized and focused on honing real job skills. IBM z/OS Quick Start Training for Application Developers is a key step in the direction of mastering IBM application development so you'll be ready to join a technical team.

This IBM® Redbooks® publication documents the strength and value of the IBM security strategy with IBM zTM Systems hardware and software. In an age of increasing security consciousness and more and more dangerous advanced persistent threats, IBM z Systems™ provides the capabilities to address the needs of today's business security challenges. This publication explores how z Systems hardware is designed to provide integrity, process isolation, and cryptographic capability to help address security requirements. We highlight the features of IBM z/OS® and other operating systems, which offer a variety of customizable security elements. We discuss z/OS and other operating systems and additional software that use the building blocks of z Systems hardware to provide solutions to business security needs. We also explore the perspective from the view of an enterprise security architect and how a modern mainframe has to fit into an overarching enterprise security architecture. This book is part of a three-volume series that focuses on guiding principles for optimized mainframe security configuration within a holistic enterprise security architecture. The series' intended audience includes enterprise security architects, planners, and managers who are interested in exploring how the security design and features of z Systems, the z/OS operating system, and associated software address current issues such as data encryption, authentication, authorization, network security, auditing, ease of security administration, and monitoring.

Many companies have built data warehouses (DWs) and have embraced business intelligence (BI) and analytics solutions. Even as companies have accumulated huge amounts of data, however, it remains difficult to provide trusted information at the right time and in the right place. The amount of data collected and available throughout the enterprise continues to grow even as the complexity and urgency of receiving meaningful information continues to increase. Producing meaningful and trusted information when it is needed can only be achieved by having a proper information architecture in place and a powerful underlying infrastructure. The amounts of data to mine, cleanse, and integrate are becoming so large that increasingly the infrastructure is becoming the bottleneck. This results in low refresh rates of the data in the data warehouse and in not having the information available in time where it is needed. And even before information can become available in a BI dashboard or a report, many preceding steps must take place: the collection of raw data; integration of data from multiple data stores, business units or geographies; transformation of data from one format to another; cubing data into data cubes; and finally, loading changes to data in the data warehouse. Combining the complexity of the information requirements, the growing amounts of data, and multiple layers of the information architecture requires an extremely powerful infrastructure. This IBM® Redguide™ publication explains how you can use IBM System z® as the foundation for your information management architecture. The System z value proposition for information management is fueled by the traditional strengths of the IBM mainframe, the specific strengths of DB2® for z/OS®, and the broad functionality of the IBM information management software portfolio. For decades, System z has proven its ability to manage vast amounts of mission-critical data for many companies throughout the world; your data is safe on System z. The available information management functionality on System z has grown from database management systems to a full stack of solutions including solutions for content management, master data management, information integration, data warehousing, and business intelligence and analytics. The availability of Linux® on System z provides an excellent opportunity to place certain components in an easy-to-manage and scalable virtualized Linux server, while benefitting from the System z hardware strengths. DB2 on z/OS can remain the operational data store and the underlying database for the data warehouse. The next generation of System z is growing into a heterogeneous architecture with which you can take advantage of System z-managed "accelerators" running on IBM System x® or IBM Power Blades. The first of these accelerators is the IBM Smart Analytics Optimizer for DB2 for z/OS V1.1, an "all-in-one" solution in which System z, z/OS, DB2 on z/OS, an IBM BladeCenter®, and IBM storage work together to accelerate certain queries by one to two orders of magnitude. With the IBM Smart Analytics Optimizer, slices of data are periodically offloaded from DB2 on z/OS to the BladeCenter. After a query is launched against that data, it will automatically run against the data kept on the BladeCenter. The BladeCenter will process the query an order of magnitude faster than DB2 on z/OS, because all data is cached in internal memory on the BladeCenter and special compression techniques are used to keep the data footprint small and efficient. As a solid information management architecture ready for the future, System z has it all.

This IBM® Redbooks® publication gives a broad understanding of IBM IMSTM integration and connectivity solutions to access applications and data stores across your enterprise architecture. As an application developer, architect, systems integrator, or systems programmer, there is important information that is available in this book that pertains to your responsibilities to continue to include the proven

performance, data integrity, and workload distribution that is available from IMS in to selected projects that are related to your entire enterprise. This book updates and adds to the information in the following IBM Redbooks publications: IMS e-business Connectors: A Guide to IMS Connectivity, SG24-6514 IMS Connectivity in an On Demand Environment: A Practical Guide to IMS Connectivity, SG24-6794 Powering SOA Solutions with IMS, SG24-7662 IBM IMS Version 12 Technical Overview, SG24-7972 IMS 12: The IMS Catalog, REDP-4812 Rethink Your Mainframe Applications: Reasons and Approaches for Extension, Transformation, and Growth, REDP-4938

"Running Mainframe z on Distributed Platforms is particularly suitable for a more detailed discussion." Bill Ogden, IBM zPDT Redbook, April 2015 "The authors offer very well-reasoned solutions accompanied by case studies, which will be useful to specialists. The book is made even more useful as the System z mainframe-based solutions offer an advanced systems management environment for significant segments of data within large companies." Eugen Petac, Computing Reviews, Oct. 8, 2014 "Should you choose to implement zPDT, RDz UT, or RD&T in your team's arsenal, you will find Barrett and Norris's insights, genius, and hard work illuminating as to how to rationally and economically manage the environment." —Scott Fagen, Chief Architect—System z Business, CA Technologies "A must-read for anyone interested in successfully deploying cost-efficient zPDT environments with agility in an enterprise that requires simple or complex configurations. The case-study-based exposition of the content allows for its easy consumption and use. Excellent!" —Mahendra Durai, SVP & Information Technology Officer, CA Running Mainframe z on Distributed Platforms reveals alternative techniques not covered by IBM for creatively adapting and enhancing multi-user IBM zPDT environments so that they are more friendly, stable, and reusable than those envisaged by IBM. The enhancement processes and methodologies taught in this book yield multiple layers for system recovery, 24x7 availability, and superior ease of updating and upgrading operating systems and subsystems without having to rebuild environments from scratch. Most of the techniques and processes covered in this book are not new to either the mainframe or distributed platforms. What is new in this book are the authors' innovative methods for taking distributed environments running mainframe virtual machine (VM) and multiple virtual storage (MVS) and making them look and feel like other MVS systems. The authors' combined expertise involves every aspect of the implementation of IBM zPDT technology to create virtualized mainframe environments by which the mainframe operations on a z series server can be transitioned to distributed platforms. All of the enhancement methods consecutively laid out in this book have been architected and developed by the authors for the CA Technologies distributed platform. Barrett and Norris impart these techniques and processes to CIOs and CTOs across the mainframe and distributed fields, to zPDT and RDz UT implementers, and to IBM's independent software vendors and customers.

IMSTM Version 11 continues to provide the leadership in performance, reliability, and security that is expected from the product of choice for critical online operational applications. IMS 11 also offers new functions to help you keep pace with the evolving IT industry. Through the introduction of the new IMS Enterprise Suite application developers with minimal knowledge of IMS Connect can start developing client applications to communicate with IMS. With Open Database, IMS 11 also provides direct SQL access to IMS data from programs that run on any distributed platform, unlocking DL/I data to the world of SQL application programmers. In this IBM® Redbooks® publication, system programmers get the steps for installing the new IMS components, and the application programmer can follow scenarios of how client applications can take advantage of SQL to access IMS data. We describe the installation of prerequisites, such as IMS Connect and the Structured Call Interface component of Common Service Layer address space and document the set up of the three new IMS drivers: - Universal DB resource adapter - Universal JDBC driver - Universal DL/I driver Our scenarios use the JDBC driver for type-4 access from Windows® to a remote DL/I database and DB2® tables and extend it to use IBM Mashup Center to provide an effective Web interface and to integrate with Open Database. Important: IMS Enterprise Suite V2.1 is the last release of the IMS Enterprise Suite that includes the DLIModel utility plug-in. Customers should migrate to using IMS Enterprise Suite V2.2 or later, which includes the IMS Enterprise Suite Explorer for Development. DLIModel utility projects can be imported into new IMS Explorer projects. In this book, any references to generating IMS metadata classes by using the DLIModel utility are comparable to the actions used to generate the classes using the IMS Explorer for Development.

Running Mainframe z on Distributed PlatformsHow to Create Robust Cost-Efficient Multiplatform z EnvironmentsApress

This IBM® Redbooks® publication provides both introductory information and technical details about the IBM System z® Personal Development Tool (IBM zPDT®), which produces a small System z environment suitable for application development. zPDT is a PC Linux application. When zPDT is installed (on Linux), normal System z operating systems (such as IBM z/OS®) can be run on it. zPDT provides the basic System z architecture and emulated IBM 3390 disk drives, 3270 interfaces, OSA interfaces, and so on. The systems that are discussed in this document are complex. They have elements of Linux (for the underlying PC machine), IBM z/Architecture® (for the core zPDT elements), System z I/O functions (for emulated I/O devices), z/OS (the most common System z operating system), and various applications and subsystems under z/OS. The reader is assumed to be familiar with general concepts and terminology of System z hardware and software elements, and with basic PC Linux characteristics. This book provides the primary documentation for zPDT.

This IBM Redbooks publication discusses z/VM and Linux operations from the perspective of the z/OS programmer or system programmer. Although other books have been written about many of these topics, this book gives enough information about each topic to describe z/VM and Linux on IBM System z operations to somebody who is new to both environments. This book is intended for z/OS programmers and system programmers who are transitioning to the z/VM and Linux on System z environments and who want a translation guide for assistance. We base this book on our experiences using System z10 Enterprise Edition, z/VM version 5.3 RSU 0701, and Novell SUSE Linux Enterprise Server (SLES) 10 on System z.

This IBM® Redbooks® publication is based on the book Introduction to the New Mainframe: z/OS Basics, SG24-6366, which was produced by the International Technical Support Organization (ITSO), Poughkeepsie Center. It provides students of information systems technology with the background knowledge and skills necessary to begin using the basic facilities of a mainframe computer. For optimal learning, students are assumed to have successfully completed an introductory course in computer system concepts, such as computer organization and architecture, operating systems, data management, or data communications. They should also have successfully completed courses in one or more programming languages, and be PC literate. This textbook can also be used as a prerequisite for courses in advanced topics, or for internships and special studies. It is not intended to be a complete text covering all aspects of mainframe operation. It is also not a reference book that discusses every feature and option of the mainframe facilities. Others who can benefit from this course include experienced data processing professionals who have worked with non-mainframe platforms, or who are familiar with some

aspects of the mainframe but want to become knowledgeable with other facilities and benefits of the mainframe environment. As we go through this course, we suggest that the instructor alternate between text, lecture, discussions, and hands-on exercises. Many of the exercises are cumulative, and are designed to show the student how to design and implement the topic presented. The instructor-led discussions and hands-on exercises are an integral part of the course, and can include topics not covered in this textbook. In this course, we use simplified examples and focus mainly on basic system functions. Hands-on exercises are provided throughout the course to help students explore the mainframe style of computing. At the end of this course, you will be familiar with the following information: Basic concepts of the mainframe, including its usage and architecture Fundamentals of IBM z/VSE® (VSE), an IBM zTM Systems entry mainframe operating system (OS) An understanding of mainframe workloads and the major middleware applications in use on mainframes today The basis for subsequent course work in more advanced, specialized areas of z/VSE, such as system administration or application programming

Mainframe computers play a central role in the daily operations of many of the world's largest corporations, and batch processing is a fundamental part of the workloads that run on the mainframe. A large portion of the workload on IBM® z/OS® systems is processed in batch mode. Although several IBM Redbooks® publications discuss application modernization on the IBM z/OS platform, this book specifically addresses batch processing in detail. Many different technologies are available in a batch environment on z/OS systems. This book demonstrates these technologies and shows how the z/OS system offers a sophisticated environment for batch. In this practical book, we discuss a variety of themes that are of importance for batch workloads on z/OS systems and offer examples that you can try on your own system. The audience for this book includes IT architects and application developers, with a focus on batch processing on the z/OS platform.

For more than 40 years, IBM® mainframes have supported an extraordinary portion of the world's computing work, providing centralized corporate databases and mission-critical enterprise-wide applications. The IBM System z®, the latest generation of the IBM distinguished family of mainframe systems, has come a long way from its IBM System/360 heritage. Likewise, its IBM z/OS® operating system is far superior to its predecessors, providing, among many other capabilities, world-class, state-of-the-art, support for the TCP/IP Internet protocol suite. TCP/IP is a large and evolving collection of communication protocols managed by the Internet Engineering Task Force (IETF), an open, volunteer organization. Because of its openness, the TCP/IP protocol suite has become the foundation for the set of technologies that form the basis of the Internet. The convergence of IBM mainframe capabilities with Internet technology, connectivity, and standards (particularly TCP/IP) is dramatically changing the face of information technology and driving requirements for ever more secure, scalable, and highly available mainframe TCP/IP implementations. In this IBM Redbooks® publication, we begin with a discussion of Virtual IP Addressing (VIPA), a TCP/IP high-availability approach that was introduced by the z/OS Communications Server. We then show how to use VIPA for high availability, both with and without a dynamic routing protocol. We also discuss a number of different workload balancing approaches that you can use with the z/OS Communications Server. We also explain the optimized Sysplex Distributor intra-sysplex load balancing. This function represents improved multitier application support using optimized local connections together with weight values from extended Workload Manager (WLM) interfaces. Finally, we highlight the most important tuning parameters and suggest parameter values that we observed to maximize performance in many client installations.

In this IBM® Redbooks® publication we demonstrate that it is possible to combine the traditional strengths of the mainframe to manage large volumes of data and run business transactions with the Web 2.0 paradigm. We can get simpler interfaces, better integration among different services, lightweight protocols for communication, and much more, together with the availability, security, and reliability of mainframe data. And we will show how mainframe data can be accessed by smartphones such as Android or iPhone. But we can do more to demonstrate how flexible the mainframe platform is. Through the use of pervasive devices it is possible to add new possibilities to mainframe applications, extending System z® capabilities. We can receive notifications in real time, for example, of successful or unsuccessful termination of a TWS job stream, or we can immediately get alerts about abends that occurred in a critical application. This book is another demonstration that the mainframe is alive and kicking and can and should play a key role in modern application architectures.

For more than 20 years, Network World has been the premier provider of information, intelligence and insight for network and IT executives responsible for the digital nervous systems of large organizations. Readers are responsible for designing, implementing and managing the voice, data and video systems their companies use to support everything from business critical applications to employee collaboration and electronic commerce.

This IBM® Redbooks® publication presents an overview of the IBM Geographically Dispersed Parallel Sysplex® (IBM GDPS®) offerings and the roles they play in delivering a business IT resilience solution. The book begins with general concepts of business IT resilience and disaster recovery, along with issues related to high application availability, data integrity, and performance. These topics are considered within the framework of government regulation, increasing application and infrastructure complexity, and the competitive and rapidly changing modern business environment. Next, it describes the GDPS family of offerings with specific reference to how they can help you achieve your defined goals for disaster recovery and high availability. Also covered are the features that simplify and enhance data replication activities, the prerequisites for implementing each offering, and tips for planning for the future and immediate business requirements. Tables provide easy-to-use summaries and comparisons of the offerings. The extra planning and implementation services available from IBM also are explained. Then, several practical client scenarios and requirements are described, along with the most

suitable GDPS solution for each case. The introductory chapters of this publication are intended for a broad technical audience, including IT System Architects, Availability Managers, Technical IT Managers, Operations Managers, System Programmers, and Disaster Recovery Planners. The subsequent chapters provide more technical details about the GDPS offerings, and each can be read independently for those readers who are interested in specific topics. Therefore, if you read all of the chapters, be aware that some information is intentionally repeated.

This IBM® Redbooks® publication documents the strength and value of the IBM security strategy with IBM System z® hardware and software. In an age of increasing security consciousness, IBM System z provides the capabilities to address the needs of today's business security challenges. This publication explores how System z hardware is designed to provide integrity, process isolation, and cryptographic capability to help address security requirements. This book highlights the features of IBM z/OS® and other operating systems, which offer various customizable security elements under the Security Server and Communication Server components. This book describes z/OS and other operating systems and additional software that leverage the building blocks of System z hardware to provide solutions to business security needs. This publication's intended audience is technical architects, planners, and managers who are interested in exploring how the security design and features of System z, the z/OS operating system, and associated software address current issues, such as data encryption, authentication, authorization, network security, auditing, ease of security administration, and monitoring.

This is the comprehensive guide to Linux on the mainframe straight from the IBM Linux experts. The book covers virtualization, security, systems management, and more.

This textbook provides students with the background knowledge and skills necessary to begin using the basic functions and features of z/VM Version 5, Release 3. It is part of a series of textbooks designed to introduce students to mainframe concepts and help prepare them for a career in large systems computing. For optimal learning, students are assumed to be literate in personal computing and have some computer science or information systems background. Others who will benefit from this textbook include z/OS professionals who would like to expand their knowledge of other aspects of the mainframe computing environment. This course can be used as a prerequisite to understanding Linux on System z. After reading this textbook and working through the exercises, the student will have received a basic understanding of the following topics: The Series z Hardware concept and the history of the mainframe Virtualization technology in general and how it is exploited by z/VM Operating systems that can run as guest systems under z/VM z/VM components The z/VM control program and commands The interactive environment under z/VM, CMS and its commands z/VM planning and administration Implementing the networking capabilities of z/VM Tools to monitor the performance of z/VM systems and guest operating systems The REXX programming language and CMS pipelines Security issues when running z/VM This book provides students of information systems with the background knowledge and skills necessary to begin using the basic security facilities of IBM System z. It enables a broad understanding of both the security principles and the hardware and software components needed to insure that the mainframe resources and environment are secure. It also explains how System z components interface with some non-System z components. A multi-user, multi-application, multi-task environment such as System z requires a different level of security than that typically encountered on a single-user platform. In addition, when a mainframe is connected in a network to other processors, a multi-layered approach to security is recommended. Students are assumed to have successfully completed introductory courses in computer system concepts. Although this course looks into all the operating systems on System z, the main focus is on IBM z/OS. Thus, it is strongly recommended that students have also completed an introductory course on z/OS. Others who will benefit from this course include experienced data processing professionals who have worked with non-mainframe-based platforms, as well as those who are familiar with some aspects of the mainframe environment or applications but want to learn more about the security and integrity facilities and advantages offered by the mainframe environment.

Software Telemetry is a guide to operating the telemetry systems that monitor and maintain your applications. It takes a big picture view of telemetry, teaching you to manage your logging, metrics, and events as a complete end-to-end ecosystem. You'll learn the base architecture that underpins any software telemetry system, allowing you to easily integrate new systems into your existing infrastructure, and how these systems work under the hood. Throughout, you'll follow three very different companies to see how telemetry techniques impact a greenfield startup, a large legacy enterprise, and a non-technical organization without any in-house development. You'll even cover how software telemetry is used by court processes--ensuring that when your first telemetry subpoena arrives, there's no reason to panic! "1+1=3. That is the equation that summarizes the theme of this book. The book's message is to integrate the developmental principles of Agile with the result-focused approaches integral to performance consulting. Your outcomes in shaping human performance will be significant--and greater than if you only used one of these models. This is a book for anyone who seeks to work collaboratively with leaders to bring about continuously improving and sustainable organizational change." --Dana Gaines Robinson, coauthor of Performance Consulting Agile Performance Improvement demonstrates the mutual benefits that accrue to the worlds of performance consulting and agile software development when the values and principles of both are blended synergistically under the guidance of practitioners skilled in both. The agile performance improvement model blends the principles of human performance technology with the frameworks and practices of Agile. The result is an approach that maximizes the value of interactions among the consultant, the work team, and the customer. Unlike traditional end-to-end waterfall processes, agile performance improvement delivers value continuously and in small increments, relentlessly focusing on outcomes of value to the customer. Building on structures of Agile that are used in software development, such as Scrum, the agile performance improvement model considers the human component of holistic solutions in establishing a continuous stream of value. Bob Winter, a performance consultant, was the product owner for the corporate education scrum supporting an agile transition initiative for hundreds of engineering teams. From this cross-disciplinary experience, he discovered that the two cultures, two languages, and two methodologies of performance consulting and agile software development are—far from being incongruent, incompatible, or irrelevant to each other—in fact ideally suited to complement and support each other. Being agile improves the effectiveness of the performance consultant, and applying the lessons of human performance technology improves the effectiveness of software development teams. In Agile Performance Improvement, Winter teaches performance consultants how to apply agile principles, values, and methods usefully to the tasks of optimizing human performance in areas of practice not only adjoining but also well beyond the realm of software and IT engineering, such as corporate learning solutions, human resources systems, and non-software products. Conversely, he shows engineering teams immersed in an agile environment how to boost their performance using the principles and techniques taught and cultivated by performance consultants. The author, who has worked extensively on both sides of the traditional divide, recounts entertainingly but informatively how both sparks and fur can fly when geeks encounter people people.

This book was written by IBM® IT specialists who have experience implementing IBM Z® solutions, especially Linux on IBM LinuxONETM (LinuxONE) or IBM Z servers. Therefore, the content 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 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 virtualization 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 of 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 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.

You may have several triggers to investigate the feasibility of moving a workload or set of workloads to the IBM® System z® platform. These triggers could be concerns about operational cost, manageability, or delivering the agreed service levels, among others. Investigating the feasibility of a possible migration or transition to any other platform, including System z, requires a number of basic steps. These steps usually start with an understanding of the current workload and its pain points, and end with a business case to move the workload. It is important to find out how easy a migration is going to be and how much risk will be involved. In this IBM Redbooks® publication we offer thoughts on how to move through these steps. We also include a chapter with a System z technology summary to help you understand how a migrated workload may fit on the platform. Our focus in this book is on workloads that are mission-critical and require a high level of availability, including disaster recovery.

Discussions about server sprawl, rising software costs, going green, or moving data centers to reduce the cost of business are held in many meetings or conference calls in many organizations throughout the world. And many organizations are starting to turn toward System z™ and z/VM® after such discussions. The virtual machine operating system has over 40 years of experience as a hosting platform for servers, from the days of VM/SP, VM/XA, VM/ESA® and especially now with z/VM. With the consolidation of servers and conservative estimates that approximately seventy percent of all critical corporate data reside on System z, we find ourselves needing a highly secure environment for the support of this infrastructure. This document was written to assist z/VM support and security personnel in providing the enterprise with a safe, secure and manageable environment. This IBM® Redbooks® publication provides an overview of security and integrity provided by z/VM and the processes for the implementation and configuration of z/VM Security Server, z/VM LDAP Server, IBM Tivoli® Directory Server for z/OS®, and Linux® on System z with PAM for LDAP authentication. Sample scenarios with RACF® database sharing between z/VM and z/OS, or through Tivoli Directory Integrator to synchronize LDAP databases, are also discussed in this book. This book provides information about configuration and usage of Linux on System z with the System z Cryptographic features documenting their hardware and software configuration. The Consul zSecure Pro Suite is also part of this document: this product helps to control and audit security not only on one system, but can be used as a single point of enterprise wide security control. This document covers the installation and configuration of this product and detailed information is presented on how z/Consul can be used to collect and analyze z/VM security data and how it can be helpful in the administration of your audit data.

Transforming data from operational data models to purpose-oriented data structures has been commonplace for the last decades. Data transformations are heavily used in all types of industries to provide information to various users at different levels. Depending on individual needs, the transformed data is stored in various different systems. Sending operational data to other systems for further processing is then required, and introduces much complexity to an existing information technology (IT) infrastructure. Although maintenance of additional hardware and software is one component, potential inconsistencies and individually managed refresh cycles are others. For decades, there was no simple and efficient way to perform data transformations on the source system of operational data. With IBM® DB2® Analytics Accelerator, DB2 for z/OS is now in a unique position to complete these transformations in an efficient and well-performing way. DB2 for z/OS completes these while connecting to the same platform as for operational transactions, helping you to minimize your efforts to manage existing IT infrastructure. Real-time analytics on incoming operational transactions is another demand. Creating a comprehensive scoring model to detect specific patterns inside your data can easily require multiple iterations and multiple hours to complete. By enabling a first set of analytical functionality in DB2 Analytics Accelerator, those dedicated mining algorithms can now be run on an accelerator to efficiently perform these modeling tasks. Given the speed of query processing on an accelerator, these modeling tasks can now be performed much quicker compared to traditional relational database management systems. This speed enables you to keep your scoring algorithms more up-to-date, and ultimately adapt more quickly to constantly changing customer behaviors. This IBM Redbooks® publication describes the new table type that is introduced with DB2 Analytics Accelerator V4.1 PTF5 that enables more efficient data transformations. These tables are called accelerator-only tables, and can exist on an accelerator only. The tables benefit from the accelerator performance characteristics, while maintaining access through existing DB2 for z/OS application programming interfaces (APIs). Additionally, we describe the newly introduced analytical capabilities with DB2 Analytics Accelerator V5.1, putting you in the position to efficiently perform data modeling for online analytical requirements in your DB2 for z/OS environment. This book is intended for technical decision-makers who want to get a broad understanding about the analytical capabilities and accelerator-only tables of DB2 Analytics Accelerator. In addition, you learn about how these capabilities can be used to accelerate in-database transformations and in-database analytics in various environments and scenarios, including the following scenarios: Multi-step processing and reporting in IBM DB2 Query Management Facility™, IBM Campaign, or Microstrategy environments In-database transformations using IBM InfoSphere® DataStage® Ad hoc data analysis for data scientists In-database analytics using IBM SPSS® Modeler

This IBM® Redbooks® publication helps you install, configure, and use the IBM z/OS® Management Facility (z/OSMF). z/OSMF is a product for z/OS that simplifies, optimizes, and modernizes the z/OS system programmer experience. z/OSMF delivers solutions in a task-oriented, web browser-based user interface with integrated user assistance. The goal of z/OSMF is to improve system programmer productivity, and make functions easier to understand and use. This improvement makes system programmers more productive as quickly as possible with the least amount of training. You can automate tasks, reduce the learning curve, and improve productivity through a modern, simplified, and intuitive task-based, browser-based interface. z/OSMF is aimed at a mixed skills workforce: It is suited to professionals who are new to z/OS and those who are skilled in z/OS. Each professional has their own needs and faces their own

challenges. Novice system programmer might need to understand the "big picture" and how procedures are done. Novices also need access to documentation about procedures and tasks, and implement them according to the rules of the enterprise. Experienced system programmers are familiar with tasks and procedures. Therefore, the goal is to make their work less error-prone and easier. This goal allows them to be more productive and contribute more to their business. Although z/OS delivered simplification since it was introduced, z/OSMF brings a new dimension and focus to simplification. z/OSMF simplifies and modernizes the user experience and helps make pertinent information readily available and easily accessible.

Oracle Database 12c Release 1 is now supported on Linux on IBM® System z®. This platform offers many advantages to customers who rely upon the IBM mainframe systems to run their businesses. Linux on System z takes advantage of the qualities of service in the System z hardware and in IBM z/VM®, making it a robust industrial strength version of Linux. This provides an excellent platform for hosting Oracle solutions that run in an enterprise. This IBM Redbooks® publication shares experiences that are gained while installing and testing Oracle Database 12c Release 1: Recommendations about how to set up an infrastructure Installing an Oracle Grid Infrastructure Installing Oracle 12C R1 Real Application Clusters (RAC) and creating a RAC Database, including a multitenant database Using the Cloud Control Agent to manage Oracle Database 12c Release 1 Installing Oracle WebLogic Server 12c Upgrading from an Oracle Database from 11gR2 to 12c Release 1 The audience for this publication includes database consultants, installers, administrators, and system programmers. This publication is not meant to replace Oracle documentation, but to supplement it with our experiences while installing and using Oracle products.

This IBM® Redbooks® publication describes IBM TXSeries® for Multiplatforms, which is the premier IBM distributed transaction processing software for business-critical applications. Before describing distributed transaction processing in general, we introduce the most recent version of TXSeries for Multiplatforms. We focus on the following areas: The technical value of TXSeries for Multiplatforms New features in TXSeries for Multiplatforms Core components of TXSeries Common TXSeries deployment scenarios Deployment, development, and administrative choices Technical considerations It also demonstrates enterprise integration with products, such as relational database management system (RDBMS), IBM WebSphere® MQ, and IBM WebSphere Application Server. In addition, it describes system customization, reviewing several features, such as capacity planning, backup and recovery, and high availability (HA). We describe troubleshooting in TXSeries. We also provide details about migration from version to version for TXSeries. A migration checklist is included. We demonstrate a sample application that we created, called BigBlueBank, its installation, and the server-side and client-side programs. Other topics in this book include application development and system administration considerations. This book describes distributed IBM Customer Information Control System (IBM CICS®) solutions, and how best to develop distributed CICS applications.

This IBM® Redbooks® publication explores various implementations of z/OS® Identity Propagation where the distributed identity of an end user is passed to z/OS and used to map to a RACF® user ID, and any related events in the audit trail from RACF show both RACF and distributed identities. This book describes the concept of identity propagation and how it can address the end-to-end accountability issue of many customers. It describes, at a high level, what identity propagation is, and why it is important to us. It shows a conceptual view of the key elements necessary to accomplish this. This book provides details on the RACMAP function, filter management and how to use the SMF records to provide an audit trail. In depth coverage is provided about the internal implementation of identity propagation, such as providing information about available callable services. This book examines the current exploiters of z/OS Identity Propagation and provide several detailed examples covering CICS® with CICS Transaction Gateway, DB2®, and CICS Web services with Datapower.

Analytics is increasingly an integral part of day-to-day operations at today's leading businesses, and transformation is also occurring through huge growth in mobile and digital channels. Enterprise organizations are attempting to leverage analytics in new ways and transition existing analytics capabilities to respond with more flexibility while making the most efficient use of highly valuable data science skills. The recent growth and adoption of Apache Spark as an analytics framework and platform is very timely and helps meet these challenging demands. The Apache Spark environment on IBM z/OS® and Linux on IBM z Systems™ platforms allows this analytics framework to run on the same enterprise platform as the originating sources of data and transactions that feed it. If most of the data that will be used for Apache Spark analytics, or the most sensitive or quickly changing data is originating on z/OS, then an Apache Spark z/OS based environment will be the optimal choice for performance, security, and governance. This IBM® Redpaper™ publication explores the enterprise analytics market, use of Apache Spark on IBM z Systems™ platforms, integration between Apache Spark and other enterprise data sources, and case studies and examples of what can be achieved with Apache Spark in enterprise environments. It is of interest to data scientists, data engineers, enterprise architects, or anybody looking to better understand how to combine an analytics framework and platform on enterprise systems.

For more than 40 years, Computerworld has been the leading source of technology news and information for IT influencers worldwide. Computerworld's award-winning Web site (Computerworld.com), twice-monthly publication, focused conference series and custom research form the hub of the world's largest global IT media network.

[Copyright: d537ad3de02a6eddc0ce1875499a6404](https://www.ibm.com/redbooks/pdfs/rd05954.pdf)