

## Db2 For Z Os Disaster Recovery For The Rest Of Us

The ABCs of z/OS® System Programming is an eleven-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 you need to start your research into z/OS and related subjects. If you would like to become more familiar with z/OS in your current environment, or if you are evaluating platforms to consolidate your e-business applications, the ABCs collection will serve as a powerful learning tool. The contents of the volumes are:

Volume 1: Introduction to z/OS and storage concepts, TSO/E, ISPF, JCL, SDSF, and z/OS delivery and installation  
Volume 2: z/OS implementation and daily maintenance, defining subsystems, JES2 and JES3, LPA, LNKLST, authorized libraries, Language Environment®, and SMP/E  
Volume 3: Introduction to DFSMS, data set basics, storage management hardware and software, VSAM, System-Managed Storage, catalogs, and DFSMStvs  
Volume 4: Communication Server, TCP/IP and VTAM®  
Volume 5: Base and Parallel Sysplex®, System Logger, Resource Recovery Services (RRS), global resource serialization (GRS), z/OS system operations, Automatic Restart Management (ARM), Geographically Dispersed Parallel Sysplex™ (GPDS), availability in the zSeries® environment  
Volume 6: Introduction to security, RACF®, Digital certificates and PKI, Kerberos, cryptography and z990 integrated cryptography, zSeries firewall technologies, LDAP, Enterprise identity mapping (EIM), and firewall technologies  
Volume 7: Printing in a z/OS environment, Infoprint Server and Infoprint Central  
Volume 8: An introduction to z/OS problem diagnosis  
Volume 9: z/OS UNIX® System Services  
Volume 10: Introduction to z/Architecture®, zSeries processor design, zSeries connectivity, LPAR concepts, HCD, and HMC  
Volume 11: Capacity planning, performance management, RMF, and SMF  
Volume 12: WLM  
Volume 13: JES3  
Disaster Recovery with DB2 UDB for z/OS  
IBM Db2 Analytics Accelerator V7 High Availability and Disaster Recovery  
IBM Redbooks

The IBM® Smart Analytics System 9600 is a single, end-to-end business analytics solution to accelerate data warehousing and business intelligence initiatives. It provides integrated hardware, software, and services that enable enterprise customers to quickly and cost-effectively deploy business-changing analytics across their organizations. As a workload-optimized system for business analytics, it leverages the strengths of the System z® platform to drive:

Significant savings in hardware, software, operating, and people costs to deliver a complete range of data warehouse and BI capabilities  
Faster time to value with a reduction in the time and speed associated with deploying Business Intelligence  
Industry-leading scalability, reliability, availability, and security  
Simplified and faster access to the data on System z

This IBM Redbooks publication is a companion to IBM System Storage Business Continuity: Part 1 Planning Guide, SG24-6547 . We assume that the reader of this book has understood the concepts of Business Continuity planning described in that book. In this book we explore IBM System Storage solutions for Business Continuity, within the three segments of Continuous Availability, Rapid Recovery, and Backup and Restore. We position these solutions within the Business Continuity tiers. We describe, in general, the solutions available in each segment, then present some more detail on many of the products. In each case, the reader is pointed to sources of more information.

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 IBM Redbooks publication describes and demonstrates common, prescriptive scenarios for setting up disaster recovery for common workloads using IBM WebSphere Application Server, IBM DB2, and WebSphere MQ between two IBM PureApplication System racks using the features in PureApplication System V2. The intended audience for this book is pattern developers and operations team members who are setting up production systems using software patterns from IBM that must be highly available or able to recover from a disaster (defined as the complete loss of a data center).

Today, organizations engage with customers, business partners, and employees who are increasingly using mobile technology as their primary general-purpose computing platform. These organizations have an opportunity to fully embrace this new mobile technology for many types of transactions, including everything from exchanging information to exchanging goods and services, from employee self-service to customer service. With this mobile engagement, organizations can build new insight into the behavior of their customers so that organizations can better anticipate customer needs and gain a competitive advantage by offering new services. Becoming a mobile enterprise is about re-imagining your business around constantly connected customers and employees. The speed of mobile adoption dictates transformational rather than incremental innovation. This IBM® Redbooks® publication has an end-to-end example of creating a scalable, secure mobile application infrastructure that uses data that is on an IBM mainframe. The book uses an insurance-based application as an example, and shows how the application is built, tested, and deployed into production. This book is for application architects and decision-makers who want to employ mobile technology in concert with their mainframe environment.

This IBM® Redbooks® publication documents the strength and value of the IBM security strategy with IBM z™ 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.

As organizations strive to do more with less, IBM® DB2® for Linux, UNIX, and Windows provides various built-in high availability features. DB2 further provides high availability solutions by using enterprise system resources with broad support for clustering software, such as IBM PowerHA® SystemMirror®, IBM Tivoli® System Automation for Multiplatforms (Tivoli SA MP), and Microsoft Windows Cluster Server. This IBM Redbooks® publication describes the DB2 high availability functions and features, focusing on High Availability Disaster Recovery (HADR) in the OLTP environment. The book provides a detailed description of HADR, including setup, configuration, administration, monitoring, and preferred practices. This book explains how to configure Cluster software PowerHA, Tivoli SA MP, and MSCS with DB2 and show how to use these products to automate HADR takeover. DB2 also provides unprecedented enterprise-class disaster recovery capability. This book covers single system view backup, backup and restore with snapshot backup, and the db2recovery command, in detail. This book is intended for database administrators and information management professionals who want to design, implement, and support a highly available DB2 system.

The IBM® DB2® Analytics Accelerator for IBM z/OS® is a high-performance appliance that integrates the IBM zEnterprise® infrastructure with IBM PureData™ for Analytics, powered by IBM Netezza® technology. With this integration, you can accelerate data-intensive and complex queries in a DB2 for z/OS highly secure and available environment. DB2 and the Analytics Accelerator appliance form a self-managing hybrid environment running online transaction processing and online transactional analytical processing concurrently and efficiently. These online transactions run together with business intelligence and online analytic processing workloads. DB2 Analytics Accelerator V4.1 expands the value of high-performance analytics. DB2 Analytics Accelerator V4.1 opens to static Structured Query Language (SQL) applications and row set processing, minimizes data movement, reduces latency, and improves availability. This IBM Redbooks® publication provides technical decision-makers with an understanding of the benefits of version 4.1 of the Analytics Accelerator with DB2 11 for z/OS. It describes the installation of the new functions, and the advantages to existing analytical processes as measured in our test environment. This book also introduces the DB2 Analytics Accelerator Loader V1.1, a tool that facilitates the data population of the DB2 Analytics Accelerator.

The IBM® DB2® pureScale® feature offers clustering technology that helps deliver high availability and exceptional scalability transparent to applications. The DB2 pureScale feature helps you to meet your business needs around availability and scalability, and is also easy to configure and administer. This IBM Redbooks® publication addresses the DB2 pureScale feature that is available in IBM DB2 10.1 for Linux, UNIX, and Windows operating systems. It can help you build skills and deploy the DB2 pureScale feature. This book bundles all the information necessary for a in-depth analysis into the functions of the DB2 pureScale feature, including the actual hardware requirements. It includes validated step-by-step hardware and software installation instructions. In addition, this book provides detailed examples about how to work effectively with a DB2 pureScale cluster and how to plan and run an upgrade for all DB2 related components to DB2 10.1. This book is intended for database administrators (DBAs) who use IBM DB2 10.1 for Linux, UNIX, and Windows operating systems who want to explore and get started with the DB2 pureScale feature.

IBM® Z has a close and unique relationship to its storage. Over the years, improvements to the Z processors and storage software, the disk storage systems, and their communication architecture consistently reinforced this synergy. This IBM Redpaper publication summarizes and highlights the various aspects, advanced functions, and technologies that are often pioneered by IBM, and that make the IBM Z® and the IBM DS8000 products an ideal combination. This paper is intended for users who have some familiarity with IBM Z and the IBM DS8000® series and want a condensed but comprehensive overview of the synergy items up to the IBM z15™ server with z/OS v2.5 and the IBM DS8900 Release 9.2 firmware.

This IBM Redbooks publication addresses the challenges posed by monitoring high availability, scalability, and performance in an SAP sysplex data sharing environment. It introduces the motivations for utilizing a design based on DB2 data sharing. It includes the principal SAP-DB2 data sharing architecture options and trade-offs used in the industry today and issues that play a role in both high availability and scalability, such as failover design, database connectivity design, workload splitting and load balancing, MCOB, and coupling facility design. The book discusses single point of failure, important failover scenarios and outage avoidance, automation of high availability constructs, and backup and recovery considerations in data sharing environments. Performance issues are detailed in the order you would approach them at planning and implementation time. First, it discusses tuning the sysplex, which is the base for a well-performing DB2 data sharing system, then tuning the DB2 data sharing system, which is the base for a well-performing SAP system, and finally, tuning the SAP system. The book focuses on initial planning for performance and monitoring it afterward, and explains the key points to look for to health-check your system and maintain high performance.

DB2® 10 for z/OS can reduce the total DB2 CPU demand from 5-20%, compared to DB2 9, when you take advantage of all the enhancements. Many CPU reductions are built in directly to DB2, requiring no application changes. Some enhancements are implemented through normal DB2 activities through rebinding, restructuring database definitions, improving applications, and utility processing. The CPU demand reduction features have the potential to provide significant total cost of ownership savings based on the application mix and transaction types. Improvements in optimization reduce costs by processing SQL automatically with more efficient data access paths. Improvements through a range-list index scan access method, list prefetch for IN-list, more parallelism for select and index insert processing, better work file usage, better record identifier (RID) pool overflow management, improved sequential detection, faster log I/O, access path certainty evaluation for static SQL, and improved distributed data facility (DDF) transaction flow all provide more efficiency without changes to applications. These enhancements can reduce total CPU enterprise costs because of improved efficiency in the DB2 10 for z/OS. DB2 10 includes numerous performance enhancements for Large Objects (LOBs) that save disk space for small LOBs and that provide dramatically better performance for LOB retrieval, inserts, load, and import/export using DB2 utilities. DB210 can also more effectively REORG partitions that contain LOBs. This IBM Redbooks® publication® provides an overview of the performance impact of DB2 10 for z/OS discussing the overall performance and possible impacts when moving from version to version. We include performance measurements that were made in the laboratory and provide some estimates. Keep in mind that your results are likely to vary, as the conditions and work will differ. In this book, we assume that you are somewhat familiar with DB2 10 for z/OS. See DB2 10 for z/OS Technical Overview, SG24-7892-00, for an introduction to the new functions.

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.

This IBM® Redbooks® publication helps you plan, install, configure, and manage Copy Services on the IBM DS8000® operating in an IBM Z® or Open Systems environment. This book helps you design and implement a new Copy Services installation or migrate from an existing installation. It includes hints and tips to maximize the effectiveness of your installation, and information about tools and products to automate Copy Services functions. It is intended for anyone who needs a detailed and practical understanding of the DS8000 Copy Services. This edition is an update for the DS8900 Release 9.1. Note that the Safeguarded Copy feature is covered in IBM DS8000 Safeguarded Copy, REDP-5506.

As business cycles speed up, many customers gain significant competitive advantage from quicker and more accurate business decision-making by using real data. For many customers, choosing the path to co-locate their transactional and analytical workloads on System z® better leverages their existing investment in hardware, software, and skills. We created a project to address a number of best practice questions on how to manage these newer, analytical type workloads, especially when co-located with traditional transactional workloads. The goal of this IBM® Redbooks® publication is to provide technical guidance and performance trade-offs associated with resource management and potentially DB2® data-sharing in a variety of mixed transactional / data warehouse System z topologies. The term co-location used here and in the rest of the book is specifically defined as the practice of housing both transactional (OLTP) and data warehouse (analytical) workloads within the same System z configuration. We also assumed that key portions of the transactional and data warehouse databases would reside on DB2 for z/OS®. The databases may or may not reside in a DB2 data-sharing environment; we discuss those pros and cons in this book. The intended audience includes DB2 data warehouse architects and practitioners who are facing choices in resource management and system topologies in the data warehouse arena. This specifically includes Business Intelligence (BI) administrators, DB2 database administrators (DBAs) and z/OS performance administrators / systems programmers. In addition, decision makers and architects can utilize this book to assist in making platform and database topology decisions. The book is divided into four parts. Part I, "Introducing the co-location project" covers the System z value proposition and why one should consider System z as the central platform for their data warehousing / business analytics needs. Some topics are risk avoidance via data consolidation, continuous availability, simplified disaster recovery, IBM Smart Analytics Optimizer, reduced network bandwidth requirements, and the unique virtualization and resource management capabilities of System z LPAR, z/VM® and WLM. Part I also provides some of the common System z co-location topologies along with an explanation of the general pros and cons of each. This would be useful input for an architect to understand where a customer is today and where they might consider moving to. Part II, "Project environment" covers the environment, products, workloads, workload drivers, and data models implemented for this study. The environment consisted of a logically partitioned z10™ 32way, running z/VM, Linux®, and z/OS operating system instances. On those instances we ran products such as z/OS DB2 V9, IBM Cognos® Business Intelligence Version 8.4 for Linux on System z, InfoSphere™ Warehouse for System z, InfoSphere Change Data Capture, z/OS WebSphere® V7, Tivoli® Omegamon for DB2 Performance expert. Utilizing these products we created transactional (OLTP), data warehouse query, and data warehouse refresh workloads. All the workloads were based on an existing web-based transactional Bookstore workload, that's currently utilized for internal testing within the System p® and z labs. While some IBM Cognos BI and ISWz product usage and experiences information is covered in this book, we do not go into the depth typically found in IBM Redbooks publications, since there's another book focused specifically on that IBM® DB2® Tools for z/OS® support and take advantage of the latest versions of DB2 for z/OS. These tools are integral for the administration of the DB2 for z/OS environment and for optimization of data performance. In addition, the IBM portfolio addresses additional client requirements in the areas of data governance and version upgrade acceleration. Underlying the operation of any database management system are the utilities. With the number of database objects growing exponentially, managing utility jobs, meeting service level agreements (SLAs), and ensuring recoverability can be overwhelming. IBM offers DB2 Tools solution packs that assist in the DB2 utilities management process. Solution packs combine several products into a single consolidated solution providing everything necessary to ensure the execution of a set of database administration functions. The goals are to reduce the operational complexity and reduce cost. The objective of this IBM Redbooks® publication is to document the added value in terms of productivity and performance for database administrators when using the IBM DB2 Utilities Solution Pack and the IBM DB2 Fast Copy Solution Pack. We show the functions of the tools provided by the solution packs as used in real-life scenarios and adopting utilities best practices.

This IBM® Redbooks® publication is intended to make System Programmers, Operators, and Availability Managers aware of the enhancements to recent releases of IBM z/OS® and its major subsystems in the area of planned outage avoidance. It is a follow-on to, rather than a replacement for, z/OS Planned Outage Avoidance Checklist, SG24-7328. Its primary objective is to bring together in one place information that is already available, but widely dispersed. It also presents a different perspective on planned outage avoidance. Most businesses care about application availability rather than the availability of a specific system. Also, a planned outage is not necessarily a bad thing, if it does not affect application availability. In fact, running for too long without an IPL or subsystem restart might have a negative impact on application availability because it impacts your ability to apply preventive service. Therefore, this book places more focus on decoupling the ability to make changes and updates to your system from IPLing or restarting your systems.

DB2 Developer's Guide is the field's #1 go-to source for on-the-job information on programming and administering DB2 on IBM z/OS mainframes. Now, three-time IBM Information Champion Craig S. Mullins has thoroughly updated this classic for DB2 v9 and v10. Mullins fully covers new DB2 innovations including temporal database support; hashing; universal tablespaces; pureXML; performance, security and governance improvements; new data types, and much more. Using current versions of DB2 for z/OS, readers will learn how to: \* Build better databases and applications for CICS, IMS, batch, CAF, and RRSAP \* Write proficient, code-optimized DB2 SQL \* Implement efficient dynamic and static SQL applications \* Use binding and rebinding to optimize applications \* Efficiently create, administer, and manage DB2 databases and applications \* Design, build, and populate efficient DB2 database structures for online, batch, and data warehousing \* Improve the performance of DB2 subsystems, databases, utilities, programs, and SQL stat DB2 Developer's Guide, Sixth Edition builds on the unique approach that has made previous editions so valuable. It combines: \* Condensed, easy-to-read coverage of all essential topics: information otherwise scattered through dozens of documents \* Detailed discussions of crucial details within each topic \* Expert, field-tested implementation advice \* Sensible examples

A practical guide to DB2 z/OS database administration that is 100 percent focused on running DB2 in z/OS environments The only comprehensive preparation guide for the IBM Certified Database Administrator for DB2 Universal Database V8 z/OS certification Covers database planning, design, implementation, operation, recovery, security, performance, installation, migration, and more Sample test questions help you prepare for both IBM DB2 DBA Tests 700 and 702 IBM DB2 Universal Database Version 8 for z/OS offers enterprises unprecedented opportunities to integrate information, deliver it on demand, and manage it simply and cost-effectively. Now, one of the world's leading DB2 consultants presents the definitive guide to administering DB2 UDB V8 databases in z/OS environments. DB2 for z/OS Version 8 DBA Certification Guide also serves as a key tool for anyone preparing for IBM Certified Database Administrator for DB2 Universal Database V8 for z/OS certification. IBM Gold Consultant Susan Lawson presents hundreds of practical techniques, expert guidelines, and useful tips for every facet of DB2 UDB database administration, including database implementation, operation, recovery, security, auditing, performance, installation, migration, SQL, and more. Coverage includes Understanding the DB2 product family, architecture, attachments,

and the DB2 z/OS environment Securing enterprise-class DB2 installations and applications Using SQL to create and manage database objects, and manipulate and retrieve information Mastering key DBA tasks, including loading, reorganizing, quiescing, repairing, and recovering data; recovering and rebuilding indexes; and gathering statistics Implementing data sharing in Parallel Sysplex environments Learning the fundamentals of DB2 application development from the DBA's perspective Leveraging advanced DB2 functions, including stored procedures and other object-relational extensions Optimizing DB2 applications and the DB2 engine for maximum performance Whether you are administering DB2 UDB V8 in z/OS environments, planning to do so, or preparing for DB2 UDB V8 DBA certification, DB2 for z/OS Version 8 DBA Certification Guide will be your single most valuable resource.

Over the last few years, IBM® IMSTM and IMS tools have been modernizing the interfaces to IMS and the IMS tools to bring them more in line with the current interface designs. As the mainframe software products are becoming more integrated with the Windows and mobile environments, a common approach to interfaces is becoming more relevant. The traditional 3270 interface with ISPF as the main interface is no longer the only way to do some of these processes. There is also a need to provide more of a common looking interface so the tools do not have a product-specific interface. This allows more cross product integration. Eclipse and web-based interfaces being used in a development environment, tooling using those environments provides productivity improvements in that the interfaces are common and familiar. IMS and IMS tools developers are making use of those environments to provide tooling that will perform some of the standard DBA functions. This book will take some selected processes and show how this new tooling can be used. This will provide some productivity improvements and also provide a more familiar environment for new generations DBAs. Some of the functions normally done by DBA or console operators can now be done in this eclipse-based environment by the application developers. This means that the need to request these services from others can be eliminated. This IBM Redbooks® publication examines specific IMS DBA processes and highlights the new IMS and IMS tools features, which show an alternative way to accomplish those processes. Each chapter highlights a different area of the DBA processes like: PSB creation Starting/stopping a database in an IMS system Recovering a database Cloning a set of databases This IBM® Redbooks® publication can help you tailor and configure DFSMS constructs to be used in an IBM DB2® 9 for z/OS® environment. In addition, it provides a broad understanding of new disk architectures and their impact in DB2 data set management for large installations. This book addresses both the DB2 administrator and the storage administrator. The DB2 administrator can find information about how to use DFSMS for managing DB2 data sets; the storage administrator can find information about the characteristics of DB2 data sets and how DB2 uses the disks. This book describes optimal use of disk storage functions in DB2 for z/OS environments that can best make productive use of the synergy with I/O subsystem on IBM System z®. This book covers the following topics: - Using SMS to manage DB2 catalog, log, data, indexes, image copies, archives, work files - Taking advantage of IBM FlashCopy® for DB2 utilities, striping, copy pools - Setting page sizes and using sliding allocation - A description of PAV, MA, MIDAW, EF, EA, EAV, zHPF and why they are helpful - Compressing data and the use disk and tape for large data sets - Backup and restore, and remote copy services

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.

DB2 9 for z/OS is an exciting new version, with many improvements in performance and little regression. DB2 V9 improves availability and security, as well as adds greatly to SQL and XML functions. Optimization improvements include more SQL functions to optimize, improved statistics for the optimizer, better optimization techniques, and a new approach to providing information for tuning. V8 SQL procedures were not eligible to run on the IBM System z9 Integrated Information Processor (zIIP), but changing to use the native SQL procedures on DB2 V9 makes the work eligible for zIIP processing. The performance of varying length data can improve substantially if there are large numbers of varying length columns. Several improvements in disk access can reduce the time for sequential disk access and improve data rates. The key DB2 9 for z/OS performance improvements include reduced CPU time in many utilities, deep synergy with IBM System z hardware and z/OS software, improved performance and scalability for inserts and LOBs, improved SQL optimization, zIIP processing for remote native SQL procedures, index compression, reduced CPU time for data with varying lengths, and better sequential access. Virtual storage use below the 2 GB bar is also improved. This IBM Redbooks publication provides an overview of the performance impact of DB2 9 for z/OS, especially performance scalability for transactions, CPU, and elapsed time for queries and utilities. We discuss the overall performance and possible impacts when moving from version to version. We include performance measurements that were made in the laboratory and provide some estimates. Keep in mind that your results are likely to vary, as the conditions and work will differ. In this book, we assume that you are familiar with DB2 V9. See DB2 9 for z/OS Technical Overview, SG24-7330, for an introduction to the new functions.

IBM® continues to enhance the functionality, performance, availability, and ease of use of IBM DB2® utilities. This IBM Redbooks® publication is the result of a project dedicated to the current DB2 Version 9 Utilities Suite product. It provides information about introducing the functions that help set up and invoke the utilities in operational scenarios, shows how to optimize concurrent execution of utilities and collect information for triggering utilities execution, and provides considerations about partitioning. It also describes the new functions provided by several utilities for SHARE LEVEL CHANGE execution, which maximize availability and the exploitation of DFSMS constructs by the BACKUP and RESTORE SYSTEM utilities. This book concentrates on the enhancements provided by DB2 UDB for z/OS Version 8 and DB2 for z/OS Version 9. It implicitly assumes a basic level of familiarity with the utilities provided by DB2 for z/OS and OS/390® Version 7.

The IBM® DB2® Analytics Accelerator Version 3.1 for IBM z/OS® (simply called Accelerator in this book) is a union of the IBM System z® quality of service and IBM Netezza® technology to accelerate complex queries in a DB2 for z/OS highly secure and available environment. Superior performance and scalability with rapid appliance deployment provide an ideal solution for complex analysis. In this IBM Redbooks® publication, we provide technical decision-makers with a broad understanding of the benefits of Version 3.1 of the Accelerator's major new functions. We describe their installation and the advantages to existing analytical

processes as measured in our test environment. We also describe the IBM zEnterprise® Analytics System 9700, a hybrid System z solution offering that is surrounded by a complete set of optional packs to enable customers to custom tailor the system to their unique needs..

IBM DB2® for z/OS® is a high-performance database management system (DBMS) with a strong reputation in traditional high-volume transaction workloads that are based on relational technology. IBM WebSphere® Application Server is web application server software that runs on most platforms with a web server and is used to deploy, integrate, execute, and manage Java Platform, Enterprise Edition applications. In this IBM® Redbooks® publication, we describe the application architecture evolution focusing on the value of having DB2 for z/OS as the data server and IBM z/OS® as the platform for traditional and for modern applications. This book provides background technical information about DB2 and WebSphere features and demonstrates their applicability presenting a scenario about configuring WebSphere Version 8.5 on z/OS and type 2 and type 4 connectivity (including the XA transaction support) for accessing a DB2 for z/OS database server taking into account high-availability requirements. We also provide considerations about developing applications, monitoring performance, and documenting issues. DB2 database administrators, WebSphere specialists, and Java application developers will appreciate the holistic approach of this document. IBM® Db2® Analytics Accelerator is a workload optimized appliance add-on to IBM DB2® for IBM z/OS® that enables the integration of analytic insights into operational processes to drive business critical analytics and exceptional business value. Together, the Db2 Analytics Accelerator and DB2 for z/OS form an integrated hybrid environment that can run transaction processing, complex analytical, and reporting workloads concurrently and efficiently. With IBM DB2 Analytics Accelerator for z/OS V7, the following flexible deployment options are introduced: Accelerator on IBM Integrated Analytics System (IIAS): Deployment on pre-configured hardware and software Accelerator on IBM Z®: Deployment within an IBM Secure Service Container LPAR For using the accelerator for business-critical environments, the need arose to integrate the accelerator into High Availability (HA) architectures and Disaster Recovery (DR) processes. This IBM Redpaper™ publication focuses on different integration aspects of both deployment options of the IBM Db2 Analytics Accelerator into HA and DR environments. It also shares best practices to provide wanted Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO). HA systems often are a requirement in business-critical environments and can be implemented by redundant, independent components. A failure of one of these components is detected automatically and their tasks are taken over by another component. Depending on business requirements, a system can be implemented in a way that users do not notice outages (continuous availability), or in a major disaster, users notice an outage and systems resume services after a defined period, potentially with loss of data from previous work. IBM Z was strong for decades regarding HA and DR. By design, storage and operating systems are implemented in a way to support enhanced availability requirements. IBM Parallel Sysplex® and IBM Globally Dispersed Parallel Sysplex (IBM GDPS®) offer a unique architecture to support various degrees of automated failover and availability concepts. This IBM Redpaper publication shows how IBM Db2 Analytics Accelerator V7 can easily integrate into or complement existing IBM Z topologies for HA and DR. If you are using IBM Db2 Analytics Accelerator V5.1 or lower, see IBM Db2 Analytics Accelerator: High Availability and Disaster Recovery, REDP-5104.

The importance of business continuity and disaster recovery remains at the forefront of thought for many executives and IT technical professionals. This IBM® Redpaper™ describes the lessons learned from recent disasters and how IBM storage technology can help businesses address many of the issues related to protecting their storage infrastructures and business-critical IT applications. Two principal disaster recovery metrics, Recovery Time Objective and Recovery Point Objective, are defined and, along with the associated cost tradeoffs, are discussed from the vantage point of various IBM storage technology solutions. Two IBM Business Continuance/Disaster Recovery (BC/DR) automation solutions, known as GDPS/PPRC with HyperSwap® and GDPS/PPRC HyperSwap Manager, are described and shown how they can help an installation move closer to attaining a goal of continuous operation. GDPS/PPRC with HyperSwap operates in z/OS® environments. For z/OS installations operating two or more sites, in the event of a storage subsystem, host, network or communications facility failure, a switch to processing at an alternate site can be made in almost real time by using GDPS/PPRC with HyperSwap. Additionally, many Clustered Open Systems that are integrated with IBM Remote Copy technology can be configured to switch to a second site in almost real time. In these situations, when a site switch is executed, applications that have been cloned at both sites can continue running with minimal impact to the user.

Conducting business via the Web and remaining open for business 24 hours a day, seven days a week is now commonplace. Customers come in with unexpected workloads through the Web and operate in a self-service fashion with mostly context-sensitive metadata to guide them. The strong requirement is availability. However, even with new Web applications, most of the core business systems considerations still apply, and performance is critical. Technology has been accelerating for mainframe systems. They had become adept at providing business resiliency accommodating strategic software that has been around for the last several decades such as IMSTM , DB2® , and CICS® , and they have also become a host for developing and running Web applications built in Java™ accommodating the latest business requirements. Businesses need to leverage, extend and integrate the strategic assets which represent multi-year investments to support leading edge technology. DB2 for z/OS® has come a long way and provides facilities to exploit the latest hardware and software technologies, accommodating a majority of user requirements. However, special considerations are required to build high performance applications. If you want to achieve high performance or high availability, you must use the design, programming, and operational techniques applicable to DB2. In this IBM Redbooks publication we discuss many of these techniques and provide guidelines for database and application design. We expect the best practices described in this book will help DB2 professionals design high-performance and high-availability applications. Please note that the additional material referenced in the text is not available from IBM.

As Linux on System z becomes more prevalent and mainstream in the industry, the need for it to deliver higher levels of availability is increasing. This IBM Redbooks publication starts with an explanation of high availability (HA) fundamentals such as HA concepts and terminology. It continues with a discussion of why a business needs to consider an HA solution and then explains how to determine your business single points of failure. We outline the components of a high availability solution and describe these components. Then we provide some architectural scenarios and demonstrate how to plan and decide an implementation of an end-to-end HA solution, from Linux on System z database scenarios to z/OS, and include storage, network, z/VM, Linux, and middleware. This implementation includes the IBM Tivoli System Automation for Multiplatforms (TSA MP), which monitors and automates applications distributed across Linux, AIX®, and z/OS® operating systems, as well as a GDPS based solution. It

