

DB2 Information Management Software

IBM DB2 Universal Database version 8.2 for Linux on zSeries

HIGHLIGHTS

- **Superior Value**
Combines the reliability of zSeries, the freedom of Linux, and the power of DB2 Universal Database.
- **Simplified IT Infrastructure**
Consolidates existing Linux servers on one physical machine. Eliminates unnecessary access layers with DB2 Connect.
- **Outstanding High Availability**
Leverages automated failover options and hardware resiliency to protect you from downtime and data loss.
- **Easily Managed**
Reduces costs with autonomic database capabilities and advanced server management features.
- **Exploits 64-bit and 2.6 Kernel Features**
Improves scalability and performance with larger bufferpools and page sizes.
- **World-Class Support**
Backed with years of Linux experience, expertise and world-class support services from IBM.

Overview

With the release of version 8.2, IBM® DB2® Universal Database™ (DB2 UDB) cements its position as the most flexible and robust choice for a database environment. By leveraging the reliability of the IBM eServer® zSeries® platform and the openness of Linux®, DB2 UDB helps you transform your business into a true on demand environment.

Simplified IT Infrastructure

DB2 Connect™ and DB2 UDB deployed on Linux for zSeries can help reduce network sprawl and eliminate unnecessary hardware.

Eliminate Mid-tier Hardware with DB2 Connect

In a traditional DB2 for z/OS® environment, a middle layer of hardware was needed to act as a gateway to the database. This hardware layer can be completely removed using DB2 Connect for Linux on a zSeries IFL or LPAR. With DB2 Connect on the same zSeries or S/390® server where the target data resides, you can take advantage of shortened data path length.

Consolidate Servers with DB2 UDB

A number of underutilized database servers can be consolidated with DB2 UDB for Linux on zSeries. Moving workloads to zSeries and utilizing its advanced virtualization capabilities allows you to reduce costs and improve availability.


High Availability

DB2 UDB complements the resiliency of the zSeries platform. The High Availability Disaster Recovery (HADR) and automated client-reroute features give you 24-hour-a-day, 7-day-a-week information flow and availability that is required by an on demand business.

DB2 UDB v8.2's HADR feature allows you to quickly and easily set up a highly available environment.

The HADR environment backs up the database with a replica on another physical or virtual machine in standby mode at a local or remote site. In the event of a database failure, the database on the standby system takes on the role of the primary database. The client-reroute feature allows client threads to detect the outage and automatically reconnect to the database on the standby system with minimal downtime.

To automate failure detection and initiate failover operations between two servers, Tivoli® System Automation for Linux on zSeries is included with DB2 UDB V8.2 at no additional cost.



Exploit 64-bit and 2.6 Kernel

With DB2 UDB V8.2, you can now take advantage of 64-bit memory addressability, and exploit some features of the new Linux 2.6 kernel. The larger memory capacity makes DB2 UDB more vertically scalable. For example, you can now have shared memory areas for bufferpools, heaps, and caches that are larger than the 1 GB or so limitation that exists with 31-bit databases. This new functionality results in improved performance for BI and OLTP workloads.

Application Portability

With Linux on zSeries, you can host applications in a familiar environment. You can easily port applications that have been designed and built for DB2 UDB on other Linux, UNIX, and Windows environments to DB2 UDB for Linux on zSeries.

Ease of Use

DB2 UDB version 8.2 builds on the extensive capabilities delivered in earlier versions of the product by leveraging industry-leading autonomic technologies that make DB2 UDB more self-configuring, self-optimizing, and self-healing. For example, the Design Advisor can reduce the amount of time spent fine-tuning complex databases from days and weeks to a matter of minutes. Similarly, the automated database maintenance features improve administrator productivity by automating tasks such as database backups, data defragmentation (table or index reorganization), and data access optimization (collecting statistics).



© Copyright IBM Corporation 2004
IBM Toronto Lab
8200 Warden Avenue
Markham, ON
L6G 1C7
Canada

10-04
All Rights Reserved.

Neither this documentation nor any part of it may be copied or reproduced in any form or by any means or translated into another language, without the prior consent of the IBM Corporation.

IBM, the IBM logo, DB2, DB2 Connect, DB2 Universal Database, eServer, S/390, Tivoli, z/OS and zSeries are trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

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

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

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

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT.

The information contained in this document is subject to change without any notice. IBM reserves the right to make any such changes without obligation to notify any person of such revision or changes. IBM makes no commitment to keep the information contained herein up to date.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

For more information about IBM DB2 UDB for Linux, please visit ibm.com/db2/linux