DB2 UDB 8.2 (64 bit) für Linux on zSeries

IBM Software Group

Günter Weber
DM Technical Sales
weberg@de.ibm.com

DB2 Information Management Software

e business on demand software

Disclaimer

- The information contained in this document has not been submitted to any formal IBM review and is distributed on an "As Is" basis without any warranty either express or implied. The use of this information is a customer responsibility.
- The measurement results presented here were run in a controlled laboratory environment using specific workloads. While the information here has been reviewed by IBM personnel for accuracy, there is no guarantee that the same or similar results will be obtained elsewhere. Performance results depend upon the workload and the environment. Customers attempting to adapt this data to their own environments do so at their own risk.
- In addition, the materials in this document may be subject to enhancements or Programming Temporary Fixes (PTFs) subsequent to the level used in this study.

Agenda

- Überblick DB2 & DB2 Connect für LUW
- Überblick Linux on zSeries
- DB2 V8.2 aka. „Stinger“ on Linux for zSeries
- DB2 V8.2 (Optional)
- Additional Information

IBM Software Group

Überblick DB2 & DB2 Connect für LUW

DB2. Data Management Software

e business software

The DB2 Evolution

- Universal Application
- Universal Extensibility
- Universal Scalability
- Universal Access
- Universal Reliability
- Universal Management

8.1
- MDC
- Common Client
- Developement Center
- SMART
- ...

7.2
- Integrated Warehousing
- Integrated OLAP
- Federated databases
- XML extender
- Spatial extender
- Abstract Data Types
- SQL stored procedures

6
- DB2 for Linux
- DB2 Satellite Edition
- Ease of use
- Ease of application development
- Performance

5
- SMP and MPP parallelism
- Data Links Manager
- SQLJ
- Initial OLAP SQL

2
- Multiple platforms
- Capacity and performance
- JDBC
- UDTs, UDFs, LOBs

1
- OS/2 and AIX platform
- Low-end OLTP

DB2 V8.1 Packaging

- DB2 UDB Enterprise Server Edition with Database Partitioning Feature (DPF) – nicht für DB2 on zLinux
- ESE – Enterprise Server Edition
  - Serveredition incl. Hostconnectivity (DB2 Connect)
  - ESE = EE + EEE
  - DPF notwendig, wenn part. DB2 Datenbanken verwendet werden
- WSE – Workgroup Server Edition (xSeries, p/iSeries)
  - Für Datenbankserver mit bis zu 4 CPUs
  - keine Hostconnectivity beinhaltet
- DB2 Connect
  - Connectivity zum DB2 for z/OS / DB2 for iSeries / DB2 for VM/VSE

DB2 V8.1 Key Enhancements

- New levels of integrated information (XML Extender)
  - Store, compose, validate, decompose, and manipulate XML data
- Innovative manageability (SMART)
  - Configuration Advisor puts the knowledge of a seasoned DBA at your fingertips
  - Health Center/Monitor keeps your database functioning
  - Memory Visualizer let's you dynamically see and control DB2's memory usage
  - Advisors to deliver expert advice on index and materialized query tables
  - Simplified management of large scale partitioned databases
  - Federated database and Web Services allows combination of data from web service providers
- Robust e-business foundation
  - Connection Concentrator for more user scalability
  - Dynamic configuration
  - In-place online reorganization
  - Online load
  - Online storage management
  - Null and default compression
  - Replication enhancements
  - New client architecture
- Integrated business intelligence
  - Multidimensional data clustering improves performance of complex queries
  - Real-time and bulk scoring of data
- Enhanced application development productivity
  - Development Center
  - WebSphere integration
  - Microsoft integration

DB2 Connect V8 - Summary of offerings - 1/2

- DB2 Connect Personal Edition (PE)
  - Licensed per workstation install
  - Targeted for single-user running workstation-based applications
  - Applicable for customers that do not want a gateway server
- DB2 Connect Enterprise Edition (EE)
  - Licensed per server (where EE is installed) and user (registered or concurrent)
  - Targeted for workstation or server based applications
  - Applicable for customers that desire a gateway server
  - Most applicable for non-web based applications
  - For web-based applications, can only license DB2 Connect EE with registered users (must be able to count ALL users)
- DB2 Connect Unlimited Edition (UE)
  - Licensed per host (S/390 or zSeries)server and per MSU of host server (or sysplex)
  - License allows unlimited installation of DB2 Connect PE and EE code with unlimited users
  - Targeted for customers with widespread use of DB2 Connect
  - Applicable for web-based applications
  - Applicable for customers that cannot identify individual users or do not want to count/track users, servers, etc.

DB2 Connect V8 - Summary of offerings - 2/2

- DB2 Connect Application Server Edition (ASE)
  - Licensed per processor
  - Processors are counted on the server(s) running the application(s) which use DB2 Connect
  - DB2 Connect ASE code may be installed on application server or on separate gateway
  - Targeted for customers using DB2 Connect with a single or few applications
  - Applicable for web-based applications: Peoplesoft, Siebel, WebSphere, Weblogic, etc.
  - Not applicable for customers that do not want to track or count processors
- How is DB2 Connect ASE different than DB2 Connect UE?
  - DB2 Connect UE license entitles you to unlimited users and applications (per host server); DB2 Connect ASE license entitles you to all users of an application(s) running on the licensed processors
  - DB2 Connect UE license entitles you to install DB2 Connect PE or DB2 Connect EE code on unlimited workstations/servers; DB2 Connect ASE entitles you to install DB2 Connect ASE code on the licensed application server or a gateway server
  - DB2 Connect UE is licensed by the capacity of the host server(S/390, zSeries)or sysplex that you connect to; DB2 Connect ASE is licensed by the processors on the server which runs the application that uses DB2 Connect
  - DB2 Connect UE charges increase as the host server capacity grows; DB2 ConnectASE charges increase as the application server grows (# of processors) and as additional application servers which use DB2 Connect are added

What is DB2 Connect

- DB2 Connect
  - implements the DRDA Application Requester
  - is available on LUW (incl. Linux for zSeries and p/iSeries)

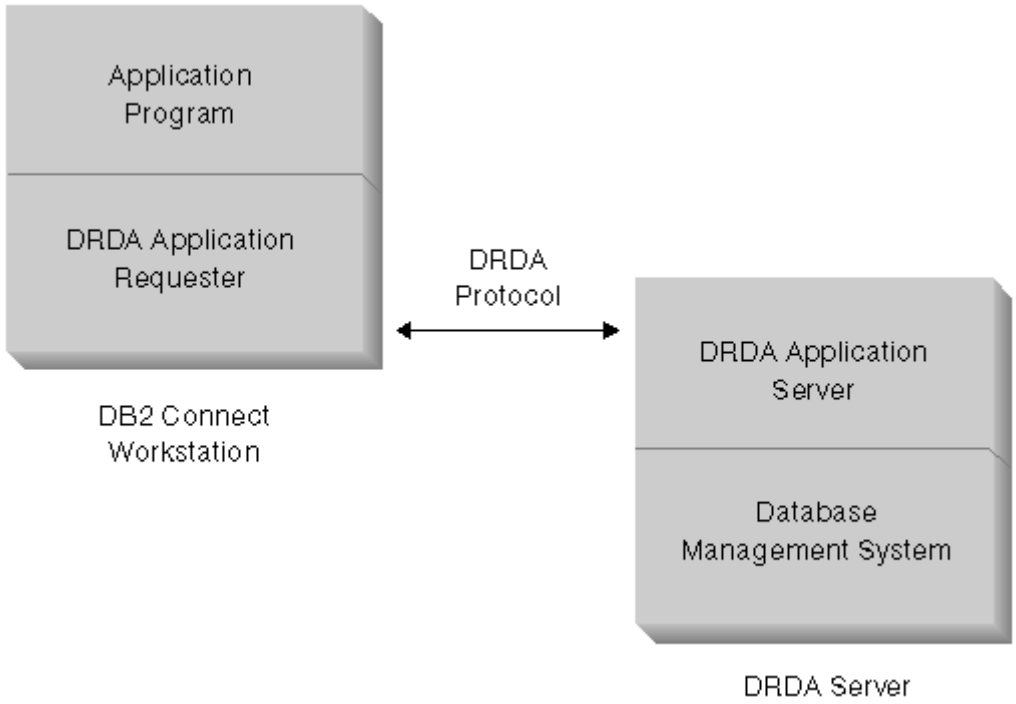

DB2 CONNECT Solutions before V8

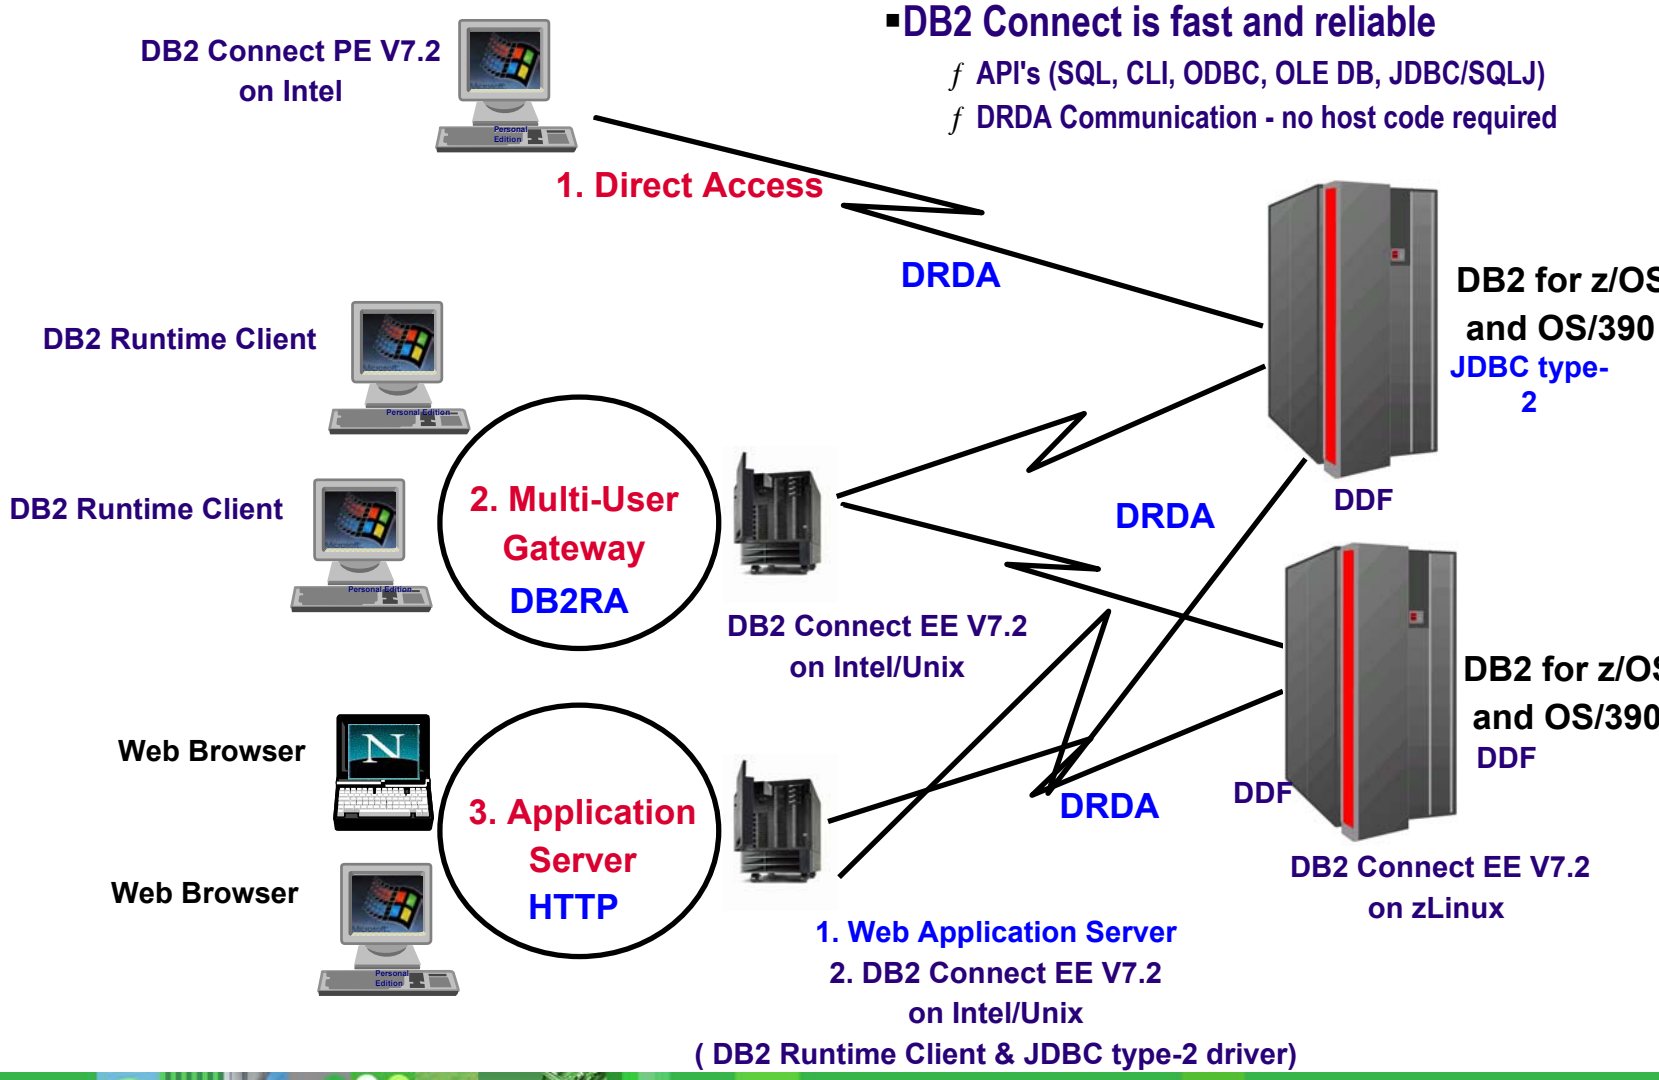

DB2 CONNECT & DB2 COMMON Client Solutions

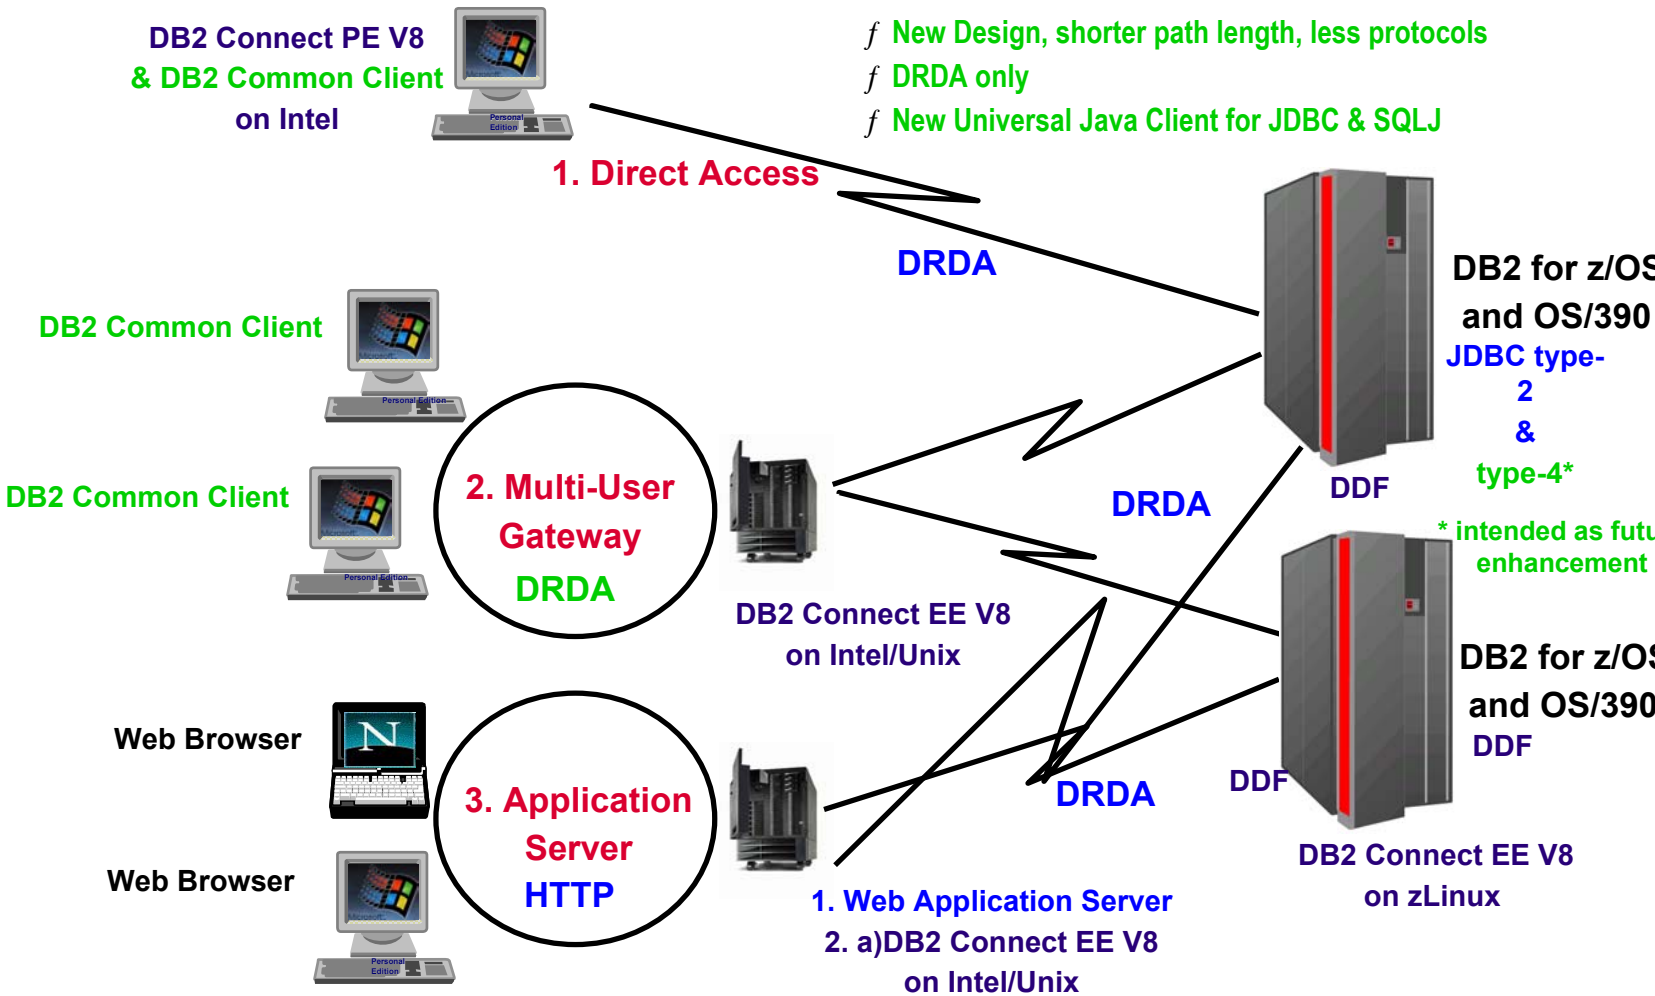

- **DB2 Common Client**
  - New Design, shorter path length, less protocols
  - DRDA only
  - New Universal Java Client for JDBC & SQLJ

Connectivity zSeries & Linux

- ✓ centralized hardware
- ✓ compatible environment
- ✓ high speed connection to z/OS data

DB2 V8.1 ESE Soft-/Hardwarevoraussetzungen

- Hardware
  - For 31-bit DB2 products running on Linux for S/390 and IBM eServer zSeries, a 9672 G5, or higher, Multiprise 3000
  - For 64-bit DB2 products for Linux on zSeries, any zSeries system that supports running Linux
- Software
  - Red Hat Enterprise Linux (RHEL) 2.1, Update 4
  - Red Hat Enterprise Linux (RHEL) 3, Update 2
  - SUSE LINUX Enterprise Server (SLES) 8 (SP3)
  - SUSE LINUX Enterprise Server (SLES) 9

Detailierte Informationen (Kernel, glibc Level) finden Sie unter:

http://www-3.ibm.com/software/data/db2/linux/validate

Überblick Linux on zSeries

DB2. Data Management Software

e business software

Was ist Linux for S/390 and zSeries (1)

- pure Linux, ASCII Environment
- 31 und 64 Bit sind verfügbar
- kein Ersatz von z/OS or other S/390

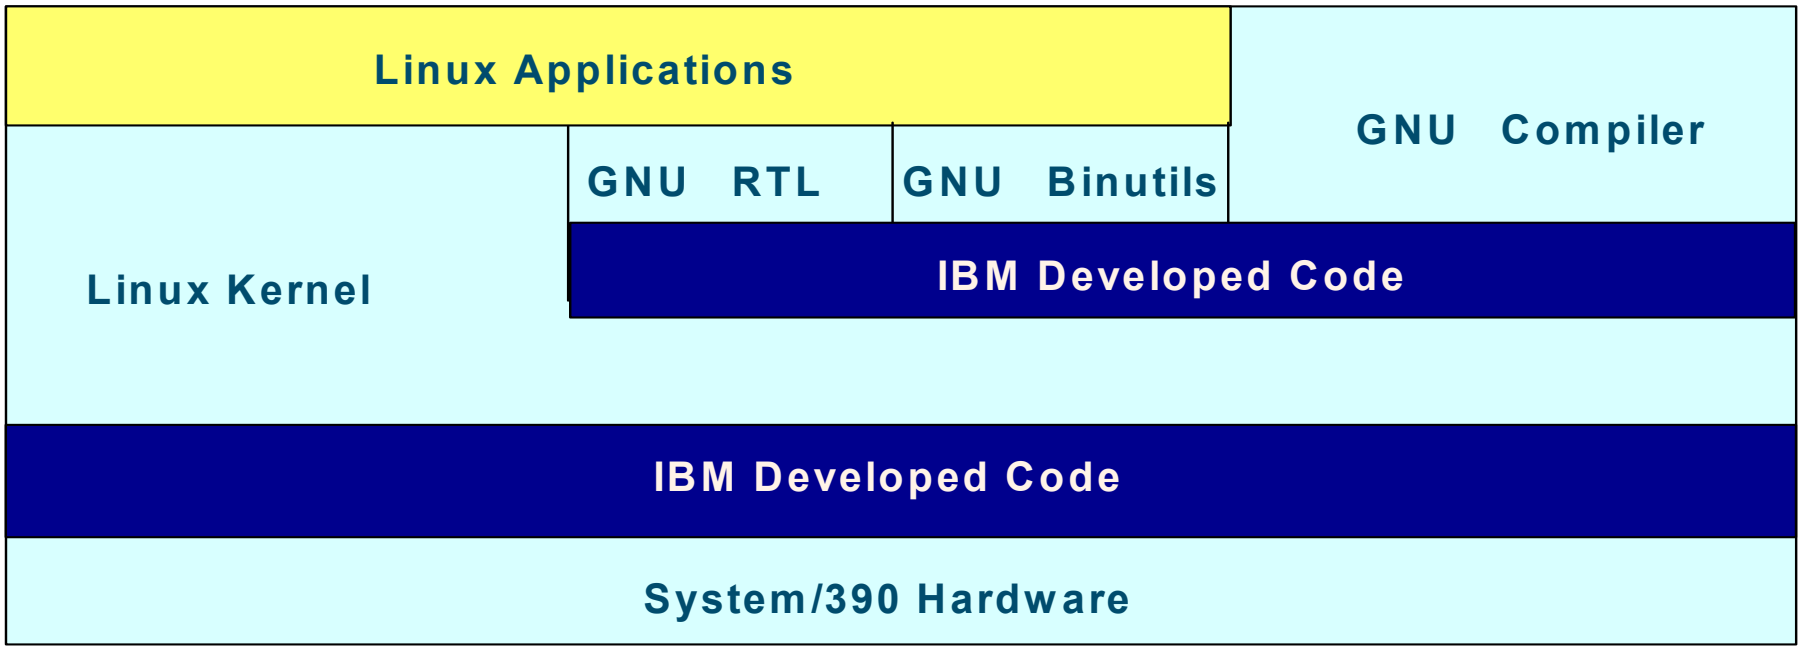

Was ist Linux for S/390 and zSeries (2)

- läuft native, in einer LPAR, oder unter z/VM
- unter z/VM sind mehrere hundert Linux Images möglich

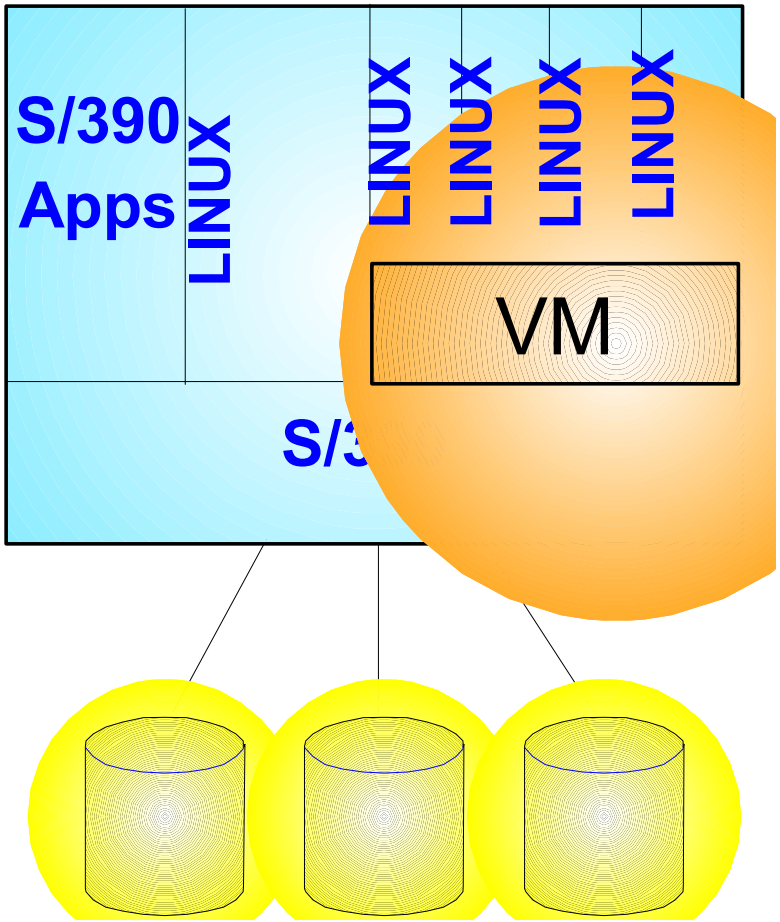

IFL

ntegrated Facility for Linux

IFL CPUs können „nur“ für Linux genutzt werden
reduziert Lizenzkosten, da Anwendungen auf Linux unter IFL nur für die IFL CPUs lizensiert werden müssen

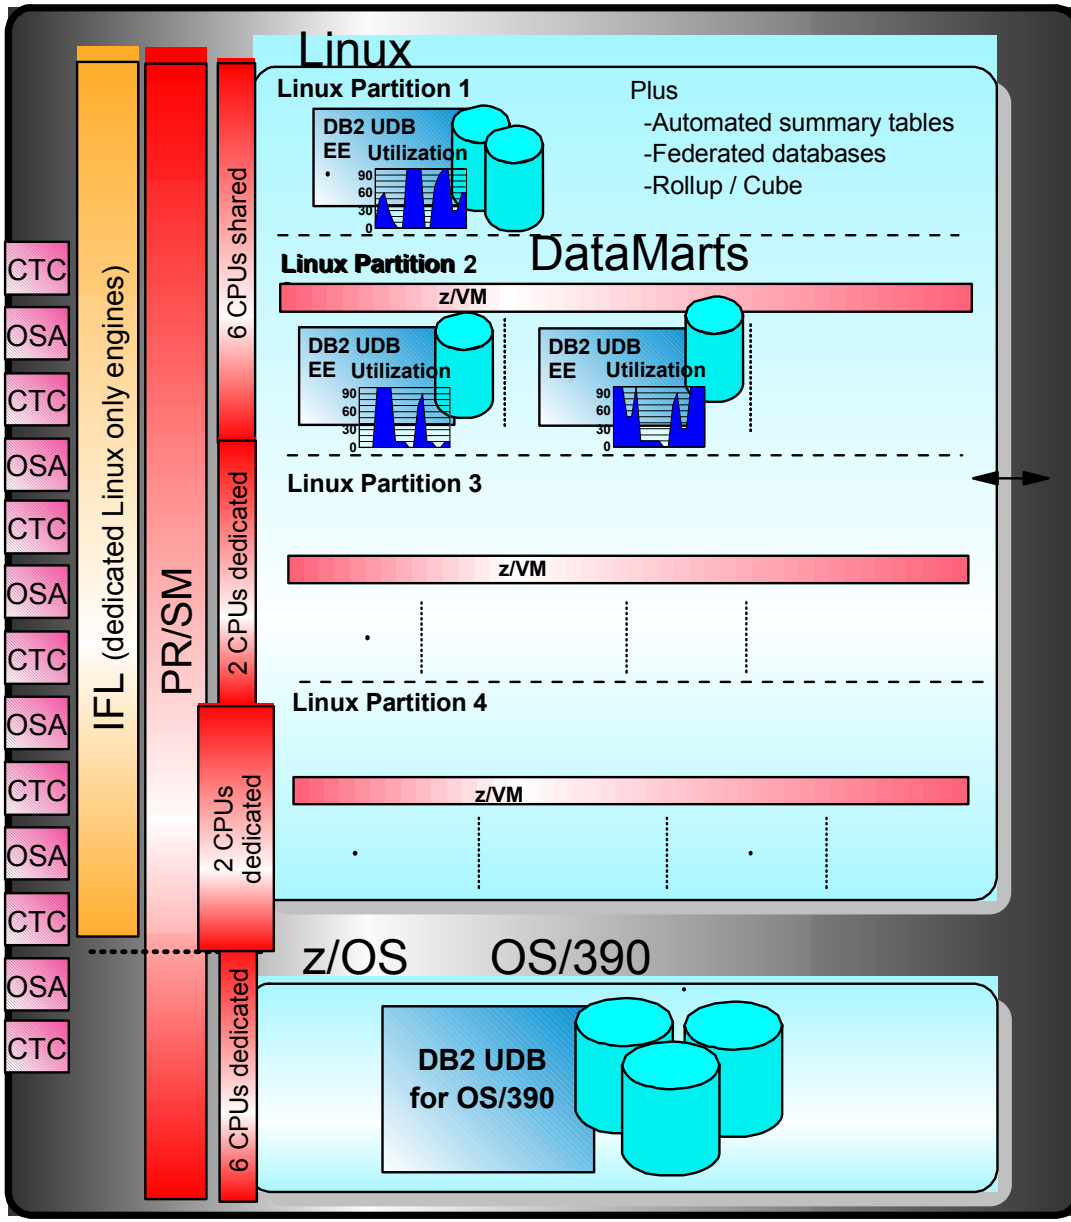

IBM Software Group

DB2 V8.2 aka. „Stinger“ on Linux for zSeries

e business on demand software

DB2 WW
Pre-sales
Support

DB2 V8.2 aka. “Stinger” aka. DB2 V8.1 + FP7

- **More than 200 new Features**
- **High Availability Disaster Recovery (HADR) and client reroute** features enable 24x7 information availability and resilience required by on demand enterprises
- **Design Advisor** (…assists DBAs in making optimal and comprehensive database design decisions
- **Autonomic Object Maintenance** feature automatically performs policy-based administration and maintenance functions, such as table reorganization, statistics collection, and database backup.
- DB2's **Linux leadership** is extended by supporting the distributions based on the new 2.6 kernel. Exploitation of new I/O and memory management features found in the latest Linux distributions is provided. 64-bit Linux versions of DB2 servers are provided for Linux on Intel EM64T, Linux on POWER (pSeries and iSeries), and Linux on zSeries hardware architectures
- **Planned availability dates:**
  - September 17, 2004 (electronic software delivery).
  - October 29, 2004 (media and documentation
- **Announcement**: http://w3-3.ibm.com/sales/ssi/rep_ia/0/877/ENUSZP04-0370/index.html

64 and 32/31 Bit Everywhere!!

- New support with Stinger
  - Native 64 Bit DB2 support for:
    - Power (IBM eServer iSeries and pSeries systems)
    - zSeries (IBM eServer zSeries systems)
    - Intel EM64T
- Existng support (here for completeness)
  - X86 (32-bit edition for Intel Pentium, Xeon and AMD Athlon based systems)
  - AMD64 (64-bit edition for AMD Opteron and Athlon64 based systems)
  - IA64 (64-bit edition for Intel Itanium based systems)
  - POWER (32-bit edition for IBM eServer iSeries and pSeries systems)
  - S/390 (31-bit edition for IBM eServer zSeries systems)

Supported with DB2 V8.1 FP6

- Linux(TM) (S/390(R) and zSeries(R) 31-bit)
  - Product list:
    - DB2 Administration Client, Version 8
    - DB2 Connect Application Server Edition, Version 8
    - DB2 Connect Enterprise Edition, Version 8
    - DB2 Connect Unlimited Edition, Version 8
    - DB2 HTML Documentation, Version 8
    - DB2 PDF Documentation, Version 8
    - DB2 Run-Time Client, Version 8
    - DB2 Spatial Extender, Version 8
    - DB2 UDB Enterprise Server Edition, Version 8

- Linux(TM) (S/390(R) and zSeries(R) 64-bit)
  - Product list:
    - DB2 Connect Application Server Edition, Version 8
    - DB2 Connect Enterprise Edition, Version 8
    - DB2 Connect Unlimited Edition, Version 8

Supported with DB2 V8.1 FP7 = V8.2 = Stinger

- Linux(TM) (S/390(R) and zSeries(R) 64-bit)
  - Product list:
    - * DB2 Administration Client, Version 8
    - * DB2 Application Development Client, Version 8
    - * DB2 Connect Application Server Edition, Version 8
    - * DB2 Connect Enterprise Edition, Version 8
    - * DB2 Connect Unlimited Edition, Version 8
    - * DB2 Run-Time Client, Version 8
    - * DB2 UDB Enterprise Server Edition, Version 8

DIRECT IO for Linux

- Currently, the file system caching policy of all I/O operations in DB2 are performed in buffered mode by default
- While this caching policy is extremely effective when the cache hit ratio is high, it has an overhead of making an extra copy of the buffer from the disk to file cache (in case of read) or from file cache to disk (in case of write).
  - Since the buffer is already cached in DB2's bufferpool layer, this dual level of caching proves to be unnecessary in situations where the cache hit ratio is low and a lot of I/O is performed.
- Direct I/O (DIO) is an alternate caching policy that reduces CPU utilization for reads and writes by eliminating the copy from file cache to user buffer.
- This feature extends DIO support on Linux distribution which are based on a 2.6 kernal (this was not ported back to previous kernel levels)
- Lab tests indicate a 12% performance improvement (OLTP workload)
- Specified by using the NO FILE SYSTEM CACHING parameter on the CREATE/ALTER TABLESPACE or CREATE DATABASE statement

Vector I/O

- Non-Vector I/O (default) requires the prefetchers to
  - read into the bufferpool cache
  - Move pages in bufferpool cache into bufferpool
- Vector I/O reads data directly into the bufferpool
- Vector I/O is supplied with Linux Kernel 2.6 and 2.4
  - Supported with:
    - Red Hat Advanced Server 2.1
    - Red Hat Enterprise Linux 3
    - SuSE Linux Enterprise 8
- Enabled by DB2SET DB2_SCATTERED_IO=ON
- Lab tests indicate a 15% performance improvement

Asynchronous I/O on Linux

- Allows DB2 to take advantage of the kernel-based asynchronous I/O capabilities of the 2.6 Linux kernel
- AIO is not enabled by default in DB2 stinger
- use db2set and the DB2NOLIOAIO variable to enable AIO
  - If false, AIO will be enabled in DB2
  - If true, AIO will be disabled.
- DB2 uses the Linux AIO interface for it's page cleaners
- AIO can noticeably improve performance on Red Hat AS 2.1. SuSE SLES 8 (2.4 kernel) also supports it, so customers can manually enable and run AIO with DB2, even with existing code.
- Lab tests have shown a 14% performance improvement when running OLTP workload on a file system. 5% on Raw Devices

Tivoli System Automation for Linux

- Included with DB2
  - No Charge IF:
    - 2 Node Cluster
    - DB2 products only (no failover for other products on the cluster)
- Detects Outages
- Recovery Policy
  - Entire system failover
  - Partial system failure
- For more information on Tivoli System Automation for Linux
  - http://www.ibm.com/software/tivoli/products/sys-auto-linux/

64-bit DB2 for zLinux

There is no inherent performance advantage from 64-bit DB2 beyond memory addressability of large RAM !

64-bit DB2 for zLinux

- zSeries offers highest levels of availability
- Best workload management
- Ideal for consolidation of low-utilization servers

"A major advantage of Linux on zSeries is its ability to consolidate many lightly to moderately loaded distributed servers that do not peak concurrently on a single zSeries server. "

-Chris Panetta

DB2 V8.2 (Optional)

DB2 V8.2 - Technology Strategy Priorities

- Optimize for Multiple Workload Environments
- Deliver Optimized, Transparent Access to All Forms of Digitized Information
- Ease Application Development
- Improved DBA Productivity – New Wizards
- Deliver High Quality Database Services
- Be the Best ISV Partner

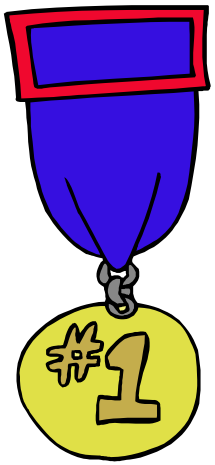

Internal SQL PL

- Internal SQL PL eliminates the current dependency on a C compiler for the creation of SQL procedures
  - Eliminating this dependency has been the most consistent PSM-related requirement
- New SQL PL generates a plan and bytecode to execute in the engine
  - Performance should be comparable to existing implementation
- Bytecode stored in the catalog and loaded at execution time

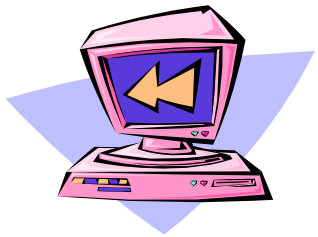

Larger SQL Statements

- The 64KB limit on statement size currently limits the total size of statements (e.g. 'CREATE PROCEDURE' or 'CREATE TRIGGER') which therefore limits the size of the object.
- This limit has been increased in order to allow partners who have the majority of their application logic in stored procedures or triggers, or need to use other large statements.
  - Currently partners have to break up the procedures and make them more complex in DB2 than in other RDBMS
  - In some cases, it is not possible to migrate a trigger or psm statement from other RDBMS to DB2 with DB2's 64K maximum statement size
- The limit will be increased to 2MB.

Data Encryption

- Encrypts user data during client/server communication over the network
- Two new authentication types
  - SQL_AUTHENTICATION_DATAENC
    - connections must use the data encryption
  - SQL_AUTHENTICATION_DATAENC_CMP
    - allows for a compatibility mode with down-level products that do not support the new authentication type
    - they will be allowed to connect with SERVER_ENCRYPT and not encrypt user data
- Encrypts all sensitive data

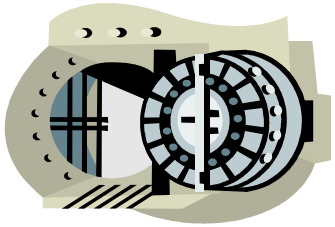

DB2 Security Exit

- DB2 does not have it's own mechanism for maintaining userids and passwords or userid group memberships
  - All authentication in DB2 is managed either through the underlying operating system or an external security system such as Kerberos
- This feature will allow the customers to create their own authentication mechanisms to handle:
  - group membership
  - authentication on the client side
  - authentication on the server

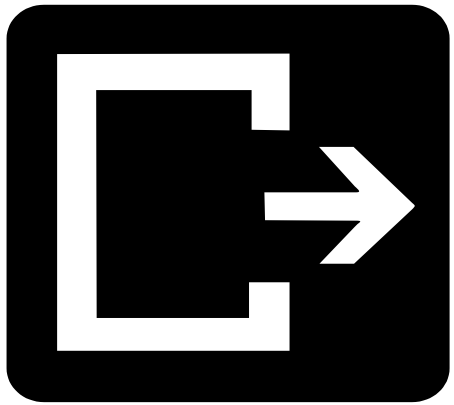

Client Reroute

- Whenever a server crashes, each client that is connected to that server gets a communication error which terminates the connection leading to an application error.
  - The DB2 UDB client code attempts to re-establish the connection to either the original server or to a new server.
  - When the connection is re-established, the application will receive an error that informs it of the transaction failure, but the application can continue with the next transaction.
- The Automatic Client Reroute feature could be used in following configuration environments:
  - Enterprise Server Edition (ESE) / Data Partitioning Feature (DPF)
  - Dpropr-style Replication
  - High Availability Cluster Multiprocessor (HACMP)
  - High Availability Disaster Recovery (HADR) environment
- Also available with LDAP Clients

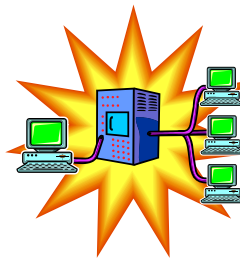

High Availability Disaster Recovery (HADR)

- Challenge
  - 24 x 7 Availability
- Solution
  - Offsite Disaster Recovery:
    - Failover to disaster site without losing work
  - Onsite Standby:
    - Apply a security patch without taking down the database

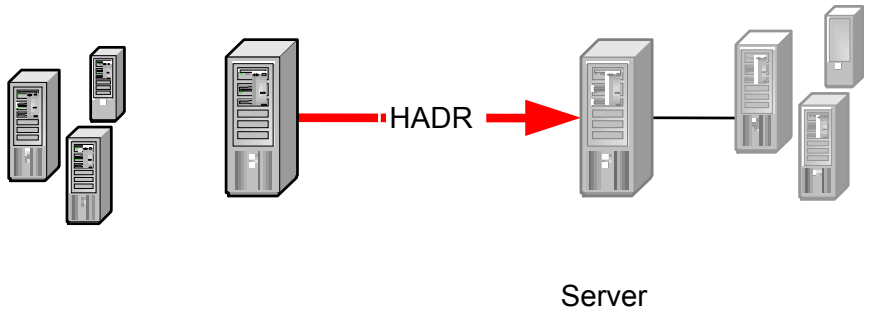

Offsite Disaster Recovery

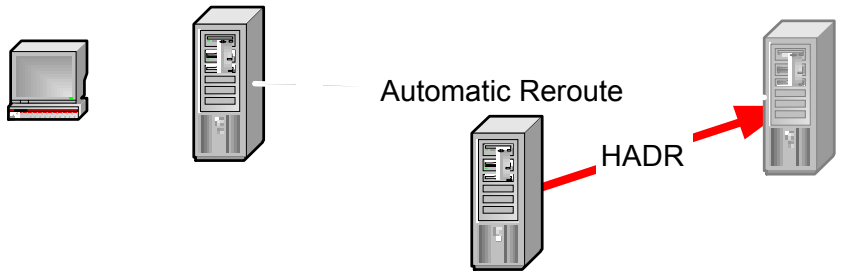

Onsite Standby

High Availability Disaster Recovery

- DB2 High Availability Disaster Recovery (HADR) is an easy to use data replication feature that provides a high availability (HA) solution for both partial and complete site failures
  - HADR is intended to be the primary HA solution for complete site failures as well as the choice for applications demanding ultra-fast failover for partial site failures.
- HADR replicates data changes from a source database (called the primary) to a target database (called the standby)
  - Using synchronous mode, HADR can guarantee that any transaction committed on the primary is also committed on the standby
  - HADR allows failover and failback between the two systems
  - HADR requires the same hardware, OS and DB2 software on the two systems (except some minor differences during rolling upgrade).
- HADR propagates data changes by shipping database log records from the primary to the standby
  - It is tightly coupled with DB2 logging and recovery.

Basic Principles of HADR

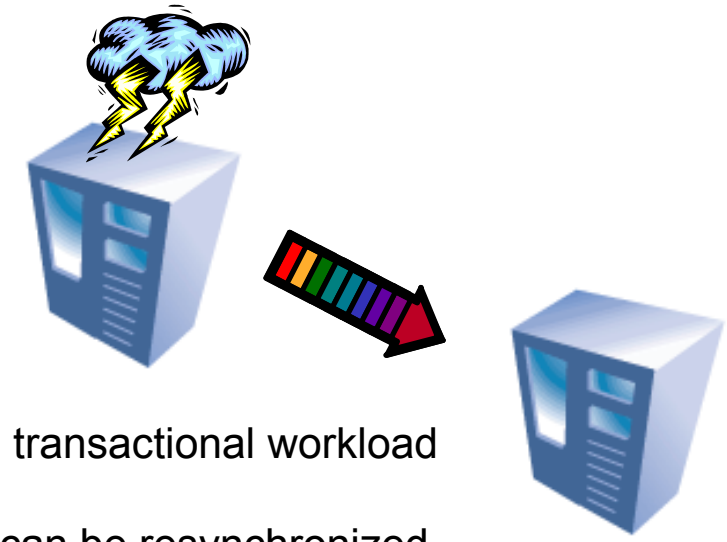

- Two active machines
  - Primary
    - Processes transactions
    - Ships log entries to the other machine
  - Standby
    - Cloned from the primary
    - Receives and stores log entries from the primary
    - Re-applies the transactions
- If the primary fails, the standby can take over the transactional workload
  - The standby becomes the new primary
- If the failed machine becomes available again, it can be resynchronized
  - The old primary becomes the new standby

Synchronization modes

- Synchronous (zero data loss)
  - The log data is flushed to stable storage at the standby before committing on the primary.
  - The log flush and TCP send() – receive() will be serialized.
  - The primary cannot proceed on to the next log flush until it receives the ACK notifying it that the standby has written the log.
  - A commit succeeds when the log data is on disk at the primary and at the standby.
- Near synchronous
  - The log data has successfully been sent to the standby site but it might not have been flushed to stable storage when the commit on the primary succeeds.
  - The log write at the primary and the send to the standby are performed in parallel at the primary.
  - A commit succeeds when the log data is on disk at the primary and it has been received by the standby
- Asynchronous
  - The log data has been given to TCP/IP and the socket send call returned successfully.
  - This is not an acknowledgment of its successful receipt by the standby, so if the connection breaks, the standby may not have received everything the primary sent.
  - However, TCP/IP sockets do guarantee delivery order, so while the socket stays alive there will never be missing or out-of-order packets as seen by the standby.
  - The log write at the primary and the send to the standby are performed in parallel at the primary.
  - A commit succeeds when the log data is on disk at the primary and it has been sent to the standby

DB2 JDBC/SQLJ drivers (open standards)

- Type 2 (app driver – db2java.zip)
  - Used for WebSphere Applications
  - High performance – all IBM benchmarks
- Type 2 (Universal JDBC/SQLJ driver – db2jcc.jar)
  - New in DB2 v8 FP2
- Type 3 (net driver – db2java.zip)
  - Native Java driver
  - Requires DB2 JDBC applet server running
- Type 4 (Universal JDBC/SQLJ driver – db2jcc.jar)
  - Native Java driver – uses DRDA
  - New in DB2 v8
- DRDA – Distributed Relational Database Architecture
  - http://www.opengroup.org/dbiop/

Java Enhancements V8.2
New Drivers and Improved Help

- Enhanced Driver Options
  - Type 4 JDBC driver
    - J2EE 1.4 / JDBC 3.0 compliance-ready
    - JTA - two-phase commit
  - Light Weight Java Runtime Client
  - Royalty Free Java Type 4 Driver Distribution
- Integrated DB2 Help based on Eclipse Help System

Who's Using DB2 on Linux Today?

Merrill Lynch

Canon

BLUESCOPE STEEL

BBDO WORLDWIDE

GROHE
WATER TECHNOLOGY

satellite

HABIB'S

HSBC Trinkaus & Burkhardt
Bank seit 1785

Ventaglio

NYSE. AMERICAN STOCK EXCHANGE

E*TRADE FINANCIAL

MERCANTIL
BANCO UNIVERSAL

TOMMY HILFIGER

CORNERCAP
INVESTMENT COUNSEL

Redbooks/Redpieces

- **Up and Running with DB2 for Linux -** SG24-6899-00
- **DB2 Universal Database Enterprise Server Edition on a Linux Cluster System**
  - on **ibm.com**/redbooks

Diskussion / Fragen

Kontakt

DB2 : Günter Weber, IBM Germany

weberg@de.ibm.com

Links & additional Information

- Die DB2 Web-Site: www-3.ibm.com/software/data/db2/linux
- Software für Linux for S/390:
  - Linux.s390.org
- IBM Linux Seite
  - www.ibm.com/linux
- DB2 Perl Database Interface
  - www.ibm.com/software/data/db2/perl
- DB2 for Linux validate-Site:
  - www-3.ibm.com/software/data/db2/linux/validate