



Technical Forum & Executive Briefing

17 al 21
Octubre
2011

Imagine **PODER** Imagine **CAPACIDAD**


DB2 pureScale on POWER Workload Optimized Systems




Workloads Are Different

Business Analytics

OLTP




- Complex queries against a data warehouse
- Read only
- Star schema design often used to speed up queries
- Different modes of operation




- Transactions against operational data
- Reads and writes
- Normalized schema design to eliminate redundancy
- Multiple user throughput operation

Modes Of Operation Have Different Performance Objectives


Business Analytics		OLTP
Complex analytic queries	Operational reports with high concurrency	Multiple users running transactions concurrently



Complete queries in the fastest



Achieve maximum throughput



Achieve maximum throughput

Competitor's One-Size-Fits-All Database Machine

versus


IBM Workload Optimized Systems

IBM workload optimized systems provide reliable transaction integrity, leading performance, and system flexibility.


Optimized for Workloads that ...



***Ensure every
business transaction
completes
successfully and
valuable data is
protected***



***Provide top speed
and 24x365 availability
to serve the globally
integrated enterprise***



***Allow configurations
that meet today's needs,
with rapid and easy
growth to keep pace
with business***

...and built for the way your business really works

Our systems – optimized for transactions – have a proven track record

1950s...1960s

TPF: Airline Reservation System

S/360: 1ST Binary Compatible computer family

IMS: Transaction & Database System



1990s

IMS, CICS, and DB2 Parallel Sysplex:

High-scale Application and Data Serving



1970s...1980s

System 38 and AS/400:

Integrated Application and Data Serving

S/370: Online Transaction Processing

DB2: Relational Database System



2000s

DB2 pureScale on PowerHA: High-Scale Database Management

WebSphere Edge Server: High-scale Web Application Serving

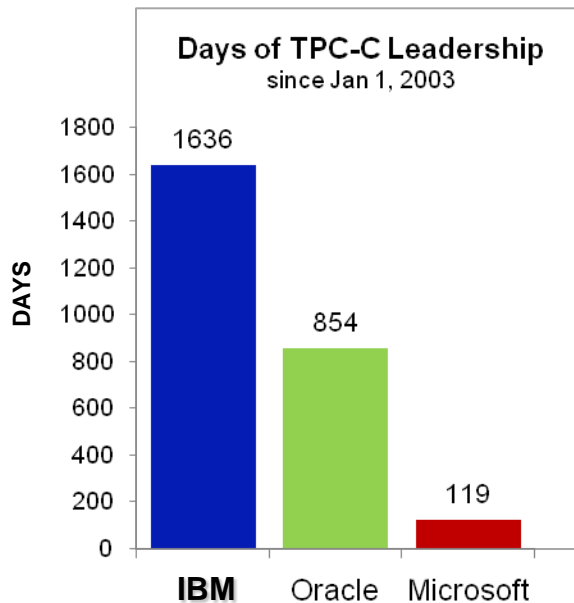
DB2 9: Industry's 1st Hybrid Relational & XML Database System

Datapower: XML & Web services appliance


Power 7

Trusted performance is measured by the speed to handle the world's most demanding business transactions.


TPC-C performance leadership



SAP SD performance leadership



SPECjEnterprise2010 Performance leadership




Source: IBM-maintained records of performance benchmark leadership. TPC-C and SAP 3-Tier SD leadership days are up to and including 22 Feb 2010. SPECjEnterprise2010 up to 05 Mar 2010.

IBM offers the highest levels of reliability for critical business systems.

Database Reliability*

Normalized Monthly Outages (all)



Data Availability**

Node Failure

% of Data Available

Time (~seconds)

DB2 pureScale

Oracle RAC

* Source: Solitaire whitepaper: DB2 Performance on IBM System p® and System x®11g

** Source: Internal IBM testing for DB2 pureScale availability combined with several public sources for Oracle RAC availability.

Do You Know -

What the Cost of a One Hour Outage is?




Do You Know -

What the Cost of a One Hour Outage is?

Industry	Average Cost Downtime Per Hour
Brokerage Services	\$6.5 million
Energy	\$2.8 million
Credit Card	\$2.6 million
Telecomm	\$2.0 million
Manufacturing	\$1.6 million
Financial	\$1.5 million
Retail	\$1.1 million
Pharmaceutical	\$1.0 million
Industry Average	\$2.4 million

Source: Network Computing, The Meta Group, and Contingency Planning Research

DB2 pureScale



Unlimited Capacity

- Buy only what you need, add more without service interruption

Application Transparency

- Avoid the risk and cost of application changes

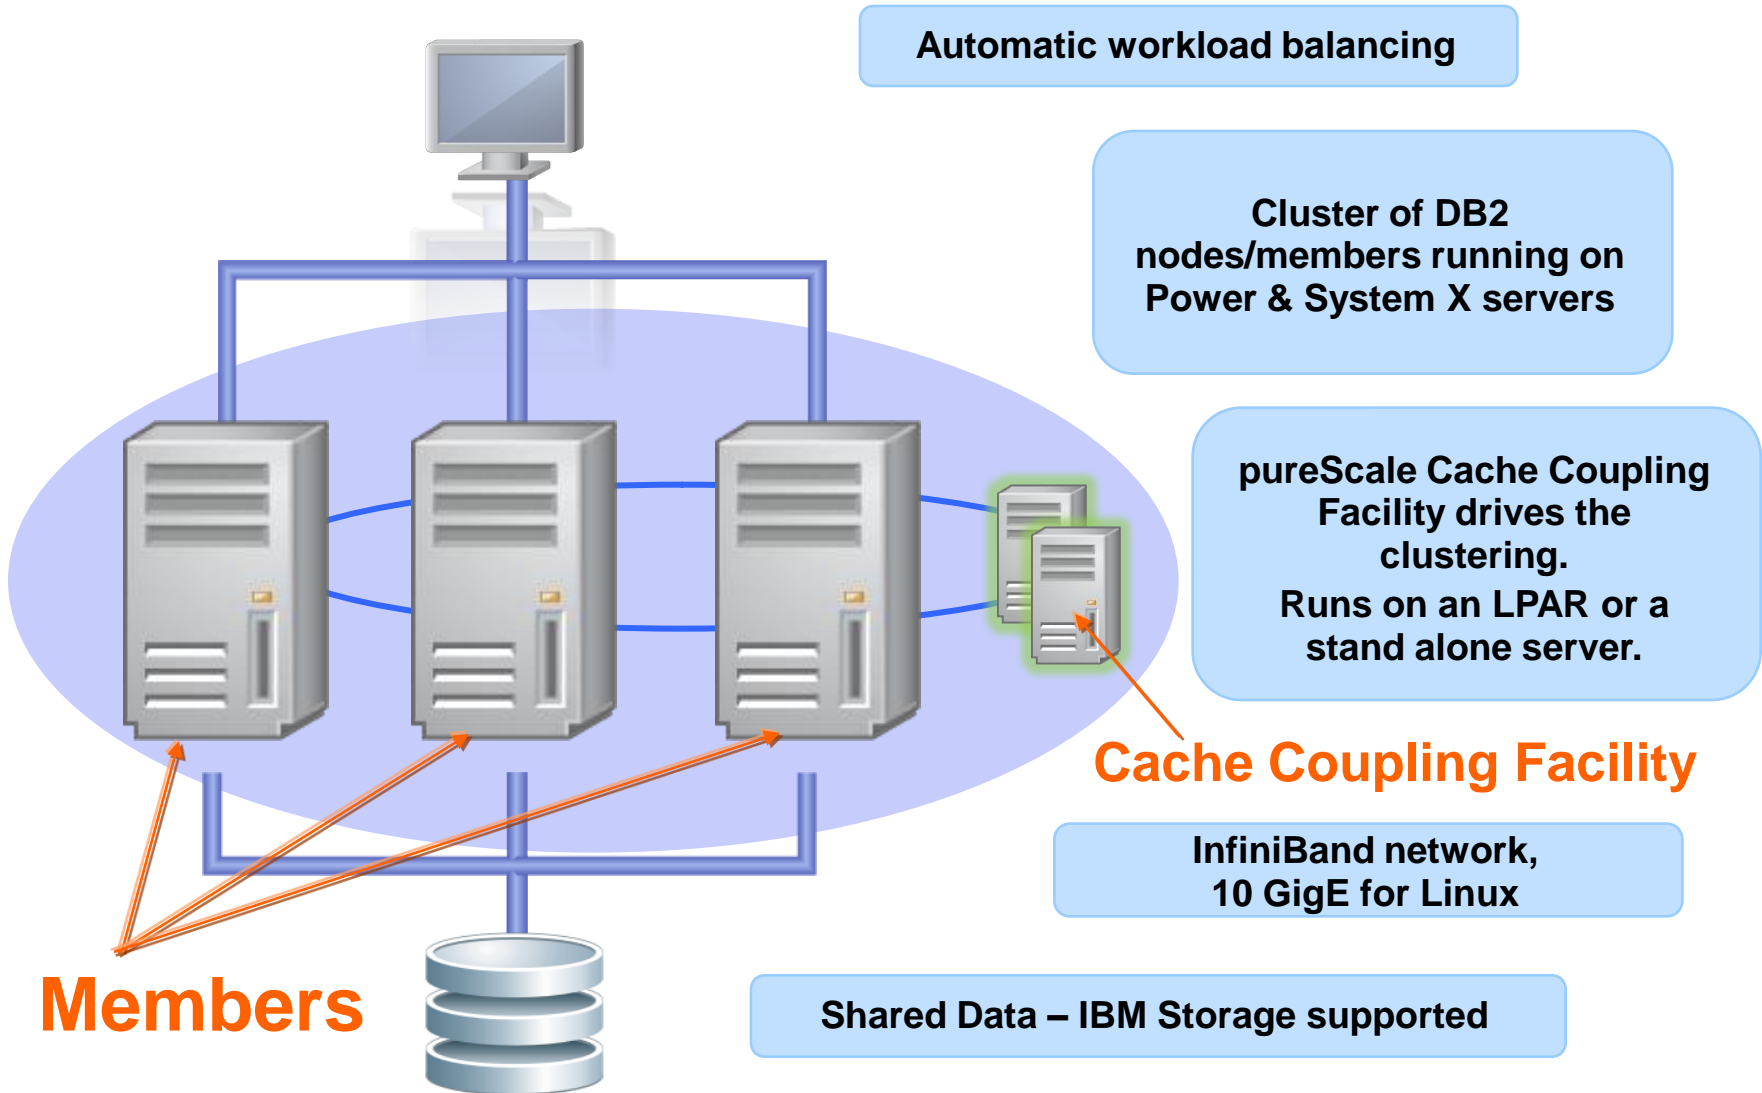
Continuous Availability

- Deliver uninterrupted access to your data with consistent performance

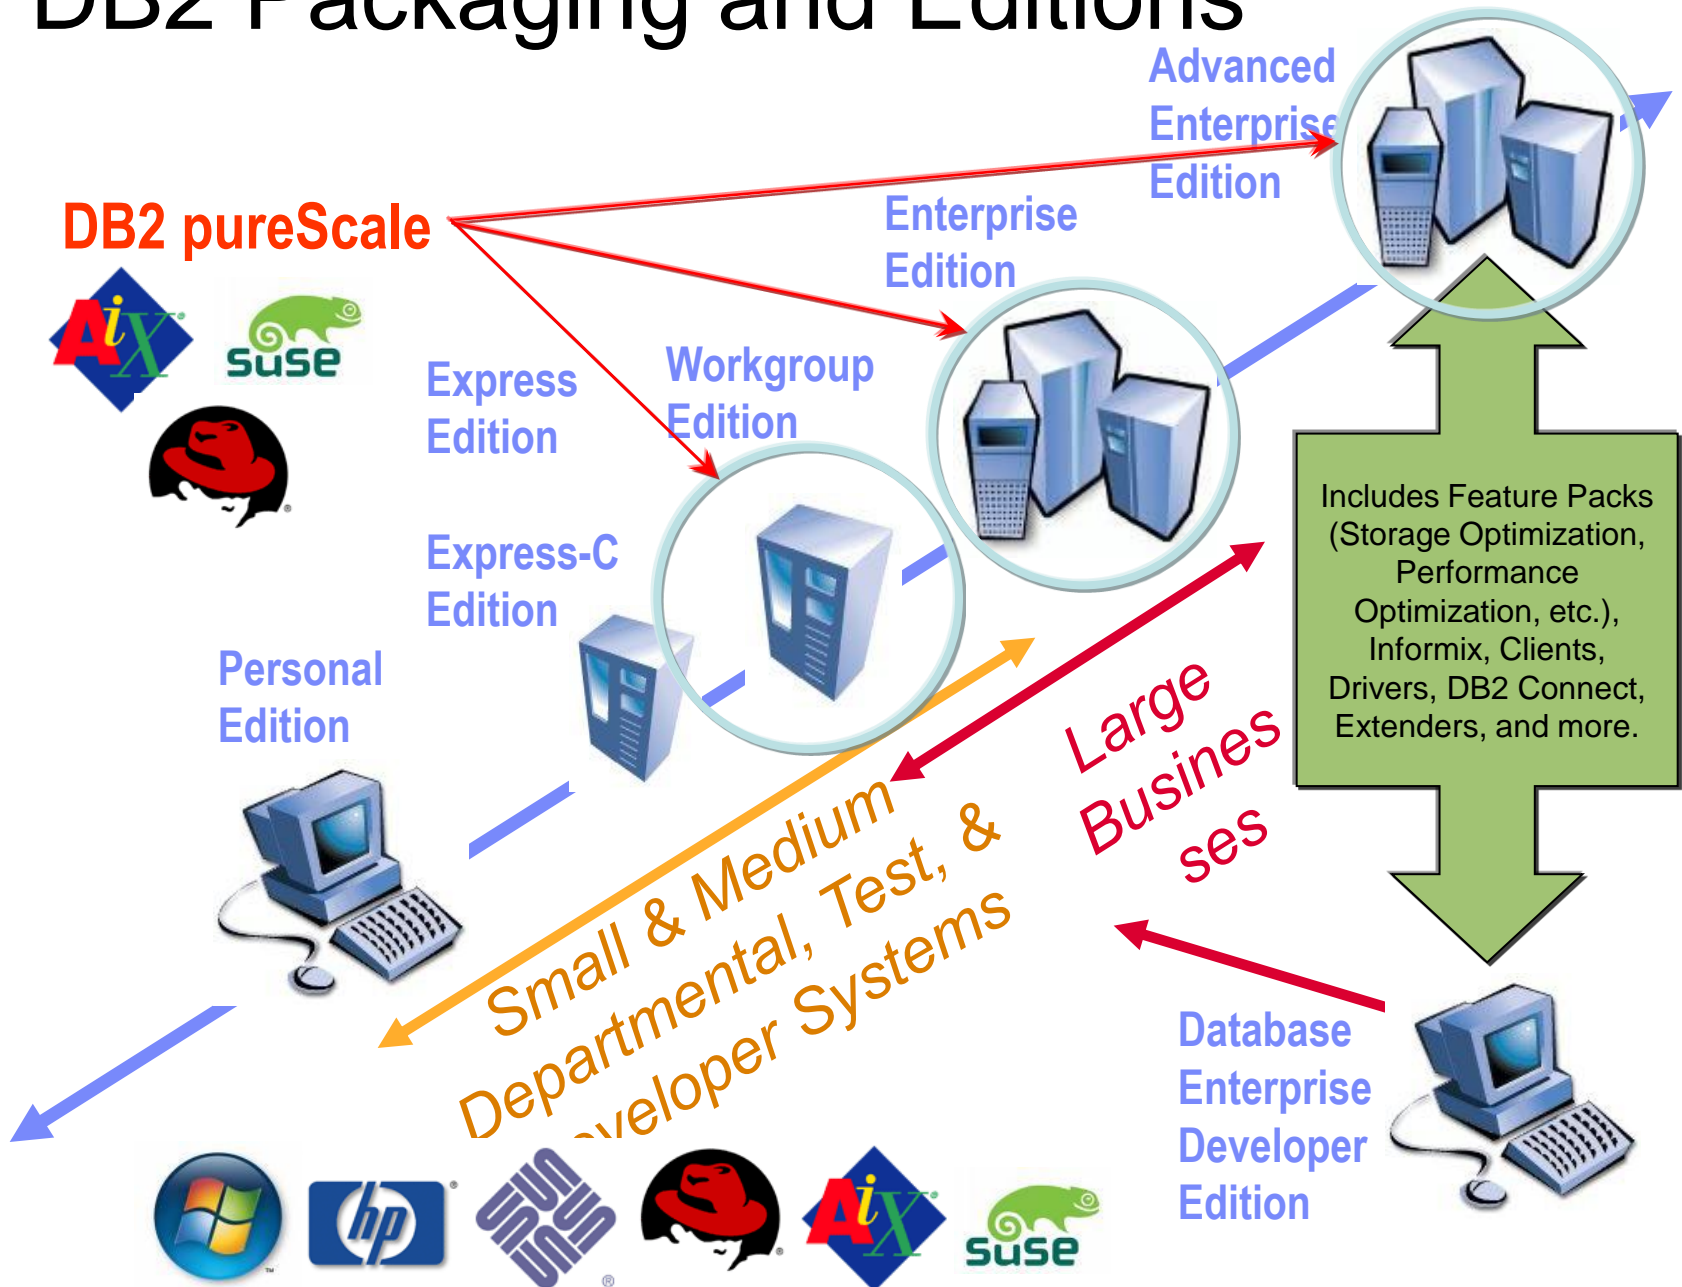
DB2's new OLTP continuous availability feature provides unlimited capacity, transparent to applications

**Leverages the architecture of z/OS:
the Gold Standard of reliability and scalability**

DB2 pureScale Architecture



DB2 Packaging and Editions




The Key to Scalability and High Availability

- Efficient Centralized Locking and Caching
 - As the cluster grows, DB2 maintains one place to go for locking information and shared pages
 - Optimized for very high speed access
 - DB2 pureScale uses Remote Direct Memory Access (RDMA) to communicate with the powerHA pureScale server
 - No IP socket calls, no interrupts, no context switching

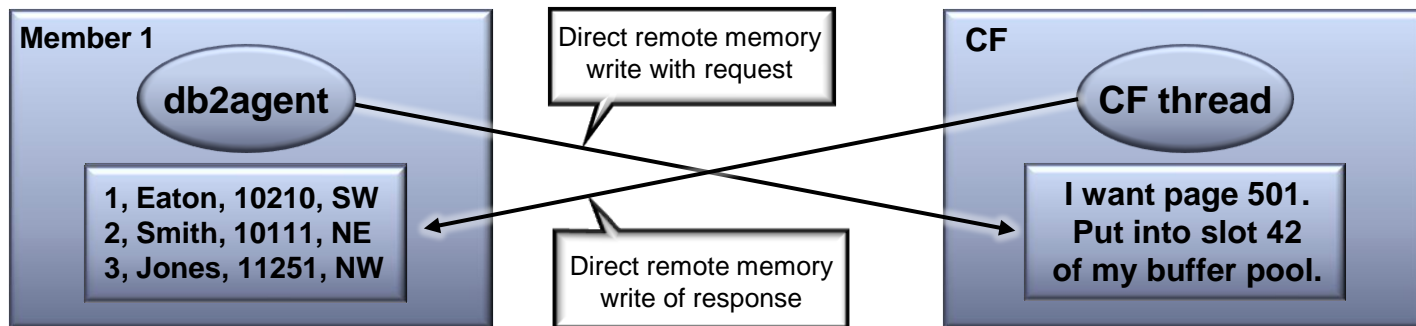
- Results

- Near Linear Scalability to large numbers of servers
- Constant awareness of what each member is doing
 - If one member fails, no need to block I/O from other members
 - Recovery runs at memory speeds



The Advantage of DB2 Read and Register with RDMA

1. DB2 agent on Member 1 writes directly into CF memory with:
 - Page number it wants to read
 - Buffer pool slot that it wants the page to go into
 2. CF either responds by writing directly into memory on Member 1:
 - That it does not have the page **or**
 - With the requested page of data
- Total end to end time for RAR is measured in microseconds
 - Calls are very fast, the agent may even stay on the CPU for the response




Much more scalable, does not require locality of data

Proof of DB2 pureScale Architecture Scalability

- How far will it scale?
- Take a web commerce type workload
 - Read mostly but **not read only**
- Don't make the application cluster aware
 - **No routing of transactions to members**
 - Demonstrate transparent application scaling
- Scale out to the 128 member limit and measure scalability

The Result




Unlimited Capacity

- DB2 pureScale has been designed to grow to whatever capacity your business requires
- Flexible licensing designed for minimizing costs of peak times
- Only pay for additional capacity when you use it even if for only a single day

Solution:

Use DB2 pureScale and add another server for those two days, and only pay sw license fees for the days you use it.




Over 100+ node architecture validation has been run by IBM

Application Transparency

Take advantage of extra capacity instantly

- No need to modify your application code
- No need to tune your database infrastructure



Your DBAs can add capacity without re-tuning or re-testing

Your developers don't even need to know more nodes are being added

Continuous Availability

- Protect from infrastructure outages
 - Architected for no single point of failure


Automatic workload balancing

Duplexed secondary global lock and memory manager

Tivoli System Automation automatically handles all component failures


DB2 pureScale stays up even with multiple node failures

Shared disk failure handled using disk replication technology

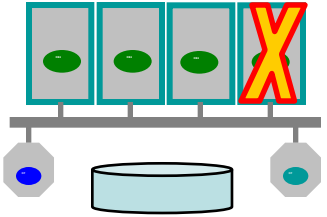


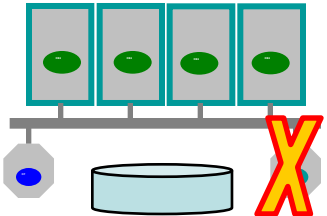


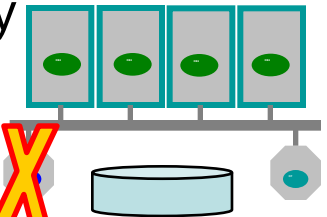




Recover Instantaneously From Node Failure

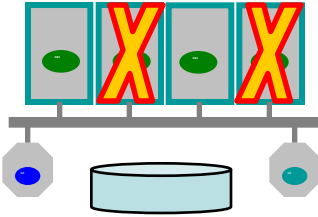


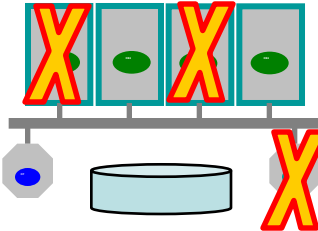


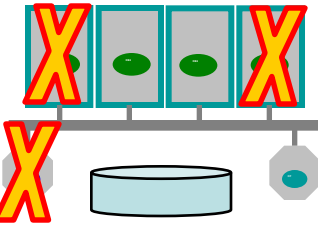


- Protect from infrastructure related outages
 - Redistribute workload to surviving nodes immediately
 - Completely redundant architecture
 - Recover in-flight transactions on failing node in as little as 15 seconds including detection of the problem



Summary (Single Failures)


Failure Mode	Other Members Remain Online ?	Automatic & Transparent ?	Comments
Member 		 Connections to failed member transparently move to another member	Only data that was inflight on failed member remains locked temporarily.
Primary <i>PowerHA pureScale</i> 			Momentary “blip” in CCF <i>pureScale</i> service. Transparent to members (In-flight CCF <i>pureScale</i> requests just take a few more seconds before completing normally.)
Secondary <i>PowerHA pureScale</i> 			Momentary “blip” in <i>PowerHA pureScale</i> service. Transparent to members (In-flight <i>PowerHA pureScale</i> requests just take a few more seconds before completing normally.)

Simultaneous Failures

Failure Mode	Other Members Remain Online ?	Automatic & Transparent ?	Comments
		 Connections to failed member transparently move to another member	Only data that was inflight on failed members remains locked temporarily. Recovers done in parallel.
		 Connections to failed member transparently move to another member	Same as member failure. Momentary, transparent, "blip" in <i>PowerHA pureScale</i> service.
		 Connections to failed member transparently move to another member	Same as member failure. Momentary, transparent, "blip" in <i>PowerHA pureScale</i> service.

Moving Your Applications to DB2 is Easy

Break Free From High Database Costs



Proven
Results

- Easily move your applications from Oracle database to DB2
- Leverage existing skills and people without re-training
- Applications moved to DB2 run quickly with full native execution
- Customers are no longer locked into Oracle RAC
- Integrated, cross-platform tools supporting Oracle database as well
- Customers and partners have moved in only days

Summary – What can DB2 pureScale Do For You?

- Deliver higher levels of scalability and superior availability
- Better concurrency during regular operations
- Better concurrency during member failure
- Result in less application design and rework for scalability
- Improved SLA attainment
- Lower overall costs for applications that require high transactional performance and ultra high availability

Questions & Answers





ibm.com/software/data/infosphere/smart-analytics-system

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by © are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.