

May 2003



**DB2** Information Management Software

# **DB2 Integrated Cluster Environment (ICE) for Linux**

## **A Solution Blueprint**

*by,  
Boris Bialek  
Rav Ahuja*

---

**Contents**

---

<b><i>Introduction</i></b>	<b>2</b>
<b><i>Ingredients for Success</i></b>	<b>3</b>
<b><i>Benefits</i></b>	<b>4</b>
<b><i>Customer Story</i></b>	<b>8</b>
<b><i>A Real Blueprint</i></b>	<b>10</b>
<b><i>Blueprint Architects</i></b>	<b>12</b>
<b><i>DB2</i></b>	<b>13</b>
<b><i>eServer xSeries</i></b>	<b>15</b>
<b><i>Mellanox</i></b>	<b>19</b>
<b><i>SAP</i></b>	<b>22</b>
<b><i>SuSE</i></b>	<b>25</b>
<b><i>Tivoli</i></b>	<b>29</b>
<b><i>Voltaire</i></b>	<b>31</b>

---

**Introduction**

---

***The Challenge***

Organizations today face the challenging job of integrating technologies to pull together a comprehensive solution that fits their business needs and IT budgets.

A complete data center solution consists of enterprise class applications, databases, management tools, network infrastructure and server hardware. Until recently, such solutions were based on proprietary systems and involved significant investment in capital, time and resources to assemble, integrate, test, manage and support. This scenario has changed dramatically with Linux clusters and best-of-breed software, hardware and services from IBM and its partners.

***The Answer***

The DB2® Integrated Cluster Environment (ICE) for Linux provides customers with a high performance data management system that can scale from 2 to 1,000 nodes in a clustered Linux environment. It is based on DB2 Universal Database™ for Linux from IBM®, the worldwide market share leader in relational database systems. The other components in a DB2 ICE solution can include:

- IBM eServer™ xSeries® platform
- IBM WebSphere® Application Server
- IBM Tivoli® systems management software
- SuSE Enterprise Linux Server
- Enterprise applications from SAP
- InfiniBand products from Mellanox and Voltaire

DB2 ICE for Linux enables customers to build on their existing business assets, lowering the total cost of computing and generating a fast return on investment. This solution is targeted at the financial, retail, manufacturing and public sector industries and offers customers the flexibility of Linux and open standards.

This paper outlines a blueprint for a DB2 ICE based solution and illustrates how even the strictest requirements for a complete data management solution can easily be satisfied with the right combination of products.

---

## Ingredients for Success

---

### **Autonomic database**

The benefits of Linux clusters, including superior performance at a low cost, are well known. When you add to the mix the built-in clustering capabilities of IBM DB2 Universal Database (UDB) for Linux, Enterprise Server Edition (ESE), you get a rock-solid foundation for all your enterprise e-business data processing needs.

### **Powerful hardware**

IBM has extensive and proven experience with clustered UNIX<sup>®</sup> computers. IBM xSeries<sup>®</sup> has applied that knowledge to produce servers that are armor-plated for Linux, optimized for database work loads, and deliver tremendous power at a fraction of the price.

### **Demanding applications**

The combination of DB2 UDB and Linux on IBM eServer xSeries, with the latest Intel<sup>®</sup> processors is powerful enough to run highly demanding business applications, including SAP R/3, mySAP Customer Relationship Management (CRM), mySAP Business Intelligence (BI) and IBM WebSphere Application Server.

### **Reliable management**

Managing a clustered environment is easy using autonomic capabilities built into DB2 software and xSeries hardware along with management tools from IBM<sup>®</sup> Tivoli<sup>®</sup> software. The availability of the solution is further enhanced using Tivoli System Automation for Linux.

### **High-speed network**

Network products from Voltaire, based on InfiniBand architecture provide the network fabric to meet the high-speed connectivity requirements of database servers, storage, and serve as gateways to the conventional Ethernet infrastructure. The InfiniBand fabric features Mellanox chips.

### **Platform support and services**

You don't have to go at it alone. IBM Global Services Linux professionals help you install, migrate, configure, tune and service your database solution. IBM's strong relationships with its Linux partners, supplemented by expertise from SuSE Linux, United Linux, RedHat, Mandrakesoft and others give you the confidence and support to deploy even the most critical solution.

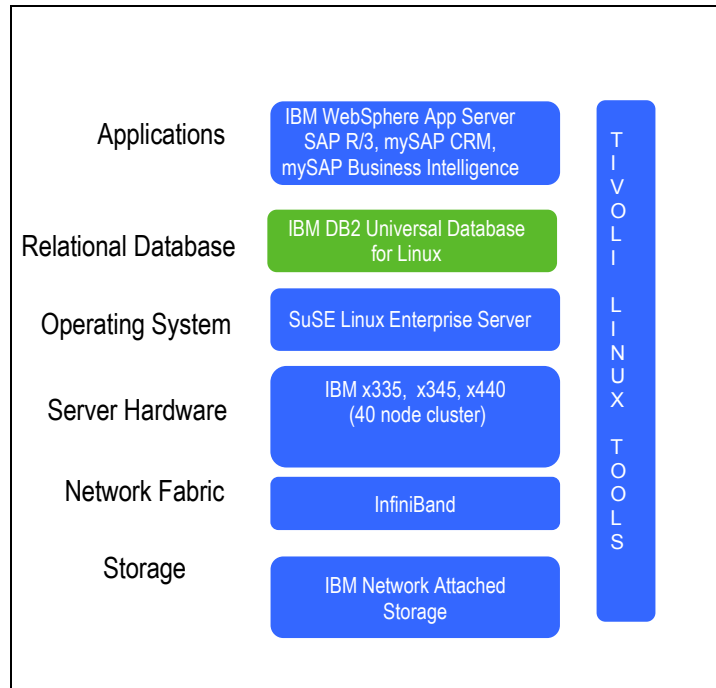


Figure 1: Ingredients of a DB2 ICE based solution

---

## Benefits

---

Managers and IT professionals tasked with selecting and implementing a data management solution that is robust enough for mission-critical enterprise needs, yet flexible enough to deliver exceptional value to small businesses, often evaluate technologies based on the following criteria:

- ✓ Scalability
- ✓ Price/Performance
- ✓ Availability
- ✓ Manageability
- ✓ Interoperability

### **Scalability pays great returns**

Many IT professionals are concerned about whether their relational database system on Linux will be able to scale as workloads increase. DB2 software and xSeries servers for Linux help you avoid large upfront costs as well as migration and training costs later. They can be deployed based on your initial requirements and then scale as large as your business demands.

The possibilities are indeed empowering. You can build powerful clusters using IBM xSeries servers running Linux. IBM DB2 Universal Database Enterprise Server Edition provides a single database that can be distributed across such a cluster. You can easily scale up by adding a new machine to the cluster, and trigger DB2 to automatically redistribute data to the new partition.

### **Proven performance at a low cost**

Operating DB2 Universal Database on IBM xSeries servers for the Linux platform, the customer can expect the high performance your business needs to support its growth. The combination of IBM DB2 Universal Database and IBM xSeries servers is a proven leader in several performance benchmarks on Linux.

Numerous application and industry standard benchmarks have delivered performance not only in lab conditions but also at countless customer installations around the world. A key reason for this success is the close cooperation with IBM business partners such as Mellanox, SuSE and Voltaire who participated in these efforts to deliver a superior. Host channel adaptors and switches based on Mellanox silicon featured in components from Voltaire can meet the most demanding needs for fast, low-latency communication between database partitions in the cluster. This solution combines leading technologies into a powerful supercomputer at a fraction of the costs of a special purpose machine with equivalent capabilities.

### High levels of availability

DB2 Universal Database is setting the standard for quality and reliability in the database industry. As more mission-critical applications are implemented on Linux, IBM's ability to bring mainframe-level reliability to this environment has become a major factor for customers choosing the DB2 family.

IBM @server offers a reliable foundation for leading Linux-based computing. With the IBM X-Architecture™ advantage, you can feel confident about the availability of the hardware running your core business applications.

The Tivoli System Automation (TSA) for Linux increases the availability of DB2 operating environments on xSeries hardware by effectively monitoring system and application health. In the event of an outage, TSA allows applications to automatically fail over to other servers in the cluster providing continuous access to DB2 and restored data. TSA is based on the autonomic research inside IBM, the longstanding experience with IBM cluster technology and specifically designed for the benefits of enterprise class solution stacks. The tight integration with DB2 for Linux offers smooth and seamless production operation with highest levels of availability.

### Easily managed

DB2 greatly reduces the complexity of data management by eliminating, simplifying, and automating many tasks traditionally associated with maintaining an enterprise class database. These advances are the first implementation of the Self-Managing and Resource Tuning (SMART) project and the first step towards making autonomic computing a reality for database implementations.

Tivoli management products simplify the management of distributed systems. Tivoli Monitoring provides monitoring for essential system resources, to detect bottlenecks and potential problems, and to automatically recover from critical situations. Tivoli NetView® extends traditional network management to ensure the availability of critical business systems and to provide rapid resolution of problems.

Tivoli Storage Manager (TSM) protects critical business information by simplifying, centralizing, and automating backup and recovery operations. Tivoli TSM support for DB2 provides online backup and quick, granular-level recovery, helping to reduce downtime and administrative costs.

### **Seamless interoperability**

We recognize that today's business environment is heterogeneous and there is a great need for different technologies to integrate seamlessly. Starting with a vision of making Linux ready for business, coupled with strong relationships, the technology vendors featured here have created a highly capable enterprise class solution. Furthermore, DB2 and xSeries have introduced validation programs designed to set standards in interoperability on Linux.

The right products working in harmony enhance productivity, reduce costs of deployment and simplify management of critical e-business applications. In other words, they empower you to focus on your business rather than the technology on which it runs.

### **Conclusion**

This solution is a proof point for implementing total enterprise solutions on Linux using scalable data management software, running on a powerful server cluster and an extremely fast network fabric, complemented with tools that provide extreme reliability and intelligent management.

*“We didn’t even consider any major relational database vendors other than IBM for the BISP project, because no other products offer the scalability and robust parallel processing that DB2 Universal Database Enterprise-Extended Edition does.”*

*–Wolfgang Schlippes-Thiede BISP  
Product Manager and  
Operational Manager of Hosting  
Center, Triaton*

---

## Customer Story

---

According to market analysts, business intelligence (BI) has become the most crucial and sought after realm of information technology for the enterprise. In today’s economic climate, mid-size and large businesses can no longer survive without cost effective and rapid access to business information for their key decision makers. Yet many businesses continue to encounter budgetary, staffing or other constraints that prevent them from introducing in-house BI systems.

### **e-business—accelerating the pace of business**

Multinational e-business consulting firm Triaton GmbH (Triaton) has developed a hosted e-business solution that allows its clients to pinpoint customer needs, improve customer relations, forecast market trends, streamline production planning and enhance their competitiveness. Leveraging BI and data management software from IBM, Triaton offers a comprehensive, customized, integrated and fully serviced BI solution—without the huge investment in hardware, software and staff that an enterprise BI infrastructure often requires. Known as Business Intelligence Service Providing (BISP), the hosted solution extracts product, customer, sales and marketing information from the client’s enterprise, hosts it in a data warehouse on Triaton’s own servers and performs the sophisticated analysis that is needed in order to glean knowledge from the data. BISP has enabled typical clients to achieve full payback within six to nine months of implementation and reduce the total cost of their IT operations by up to 30 percent.

Wolfgang Schlippes-Thiede, product manager for BISP and operational manager of the hosting center at Triaton, says BISP was born of a growing frustration on the part of many Triaton customers: “We found that many mid-size and large enterprises were hesitant to introduce BI solutions because they lacked the necessary in-house expertise or could not estimate the associated cost. The hardware, software, maintenance and operation required by such systems often amounts to more than 50 percent of companies’ IT budgets, and many of our customers were simply unable to afford the expense.”

### **Scalable database, scalable skills**

Seeking a data management platform that could power its e-business solution, Triaton evaluated DB2 Universal Database against NCR Teradata, and chose DB2 Universal Database.

Schlippes-Thiede explains why. “From the standpoint of cost performance, we found DB2 to be far superior to Teradata. We didn’t



even consider any major relational database vendors other than IBM for the BISP project, because no other products offer the scalability and robust parallel processing that DB2 Universal Database Enterprise-Extended Edition does. Its flexibility enables us to reap maximum benefit from the core DB2 competencies of our IT staff and avoid the cost of hiring or training specialists in other systems.”

*“DB2 for Linux lets us maximize our benefit from the economies and performance of the Linux operating system, providing a powerful, cost-efficient data management foundation for our BI cluster.”*

*–Wolfgang Schlippe-Thiede,  
BISP Product Manager and  
Operational Manager of Hosting  
Center, Triaton*



**Data integration for immediate intelligence**

BISP accesses and integrates data from various sources at the customer’s enterprise, including ERP, CRM and e-procurement systems, XML sources, relational databases as well as mainframe and other legacy systems. The information is extracted and transferred over the Internet or leased lines to Triaton’s data center. At the data center, the data is and is then loaded into the data warehouse, providing the foundation for data analysis.

Triaton’s three-tier solution is based on IBM xSeries servers running IBM DB2 Enterprise-Extended Edition, Version 7.2 on SuSE Linux. “DB2 for Linux lets us maximize our benefit from the economies and performance of the Linux operating system, providing a powerful, cost-efficient data management foundation for our BI cluster,” Schlippe-Thiede notes.

**Update**

Triaton has incorporated DB2 Version 8 since the original implementation of the solution. Triaton has also successfully completed a Proof-of-concept with DB2 on Intel’s 64-bit architecture systems.

---

## A Real Blueprint – 40-node cluster

---

There are numerous ways of implementing a solution based on DB2 ICE for Linux. The blueprint outlined below is one such configuration that has been tested and demonstrated at numerous industry events.

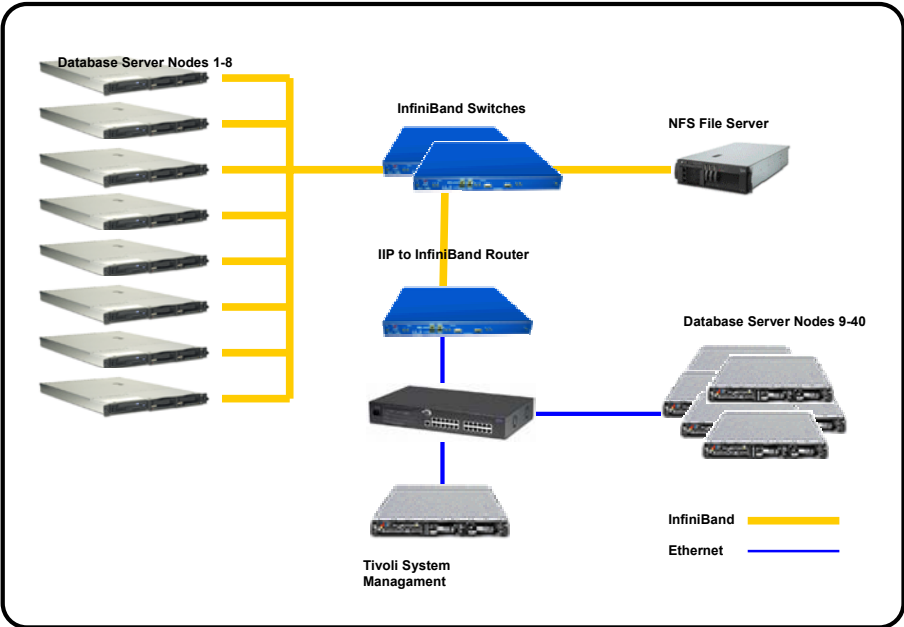
At the heart of the solution blueprint is a 40-node cluster powered by IBM DB2 Universal Database for Linux V8.1 (Enterprise Server Edition with Data Partitioning Feature). Each of the 40 nodes contains a single database partition and runs on an IBM eServer xSeries x335, x345 or x440 server, for X-architecture™ performance, reliability and scalability.

DB2 uses a shared-nothing architecture that makes full use of the cluster's computing power. With this approach, each node works on a subset of the entire database and is fully independent with its own memory, CPU(s) and disk(s). Therefore there are no resource contentions between the database partitions, ensuring extremely fast performance. This approach also makes the cluster highly scalable. If 40-nodes aren't enough, just add more as needed!

Although no special hardware is needed for the servers of a distributed database, communication needs between database nodes warrants the use of high-speed network interconnect with low-latency and high bandwidth. InfiniBand Data Center Fabric (IB-DCAF) enables high-speed communication between the nodes of the distributed database, the application servers and the storage infrastructure. Although all of the nodes in the cluster could have been connected via InfiniBand, we also showcase Gigabit Ethernet to demonstrate interoperability with existing networks.

The first three nodes in the cluster act as application server nodes for the three SAP applications. The first node in the cluster is the SAP R/3 4.6d central instance and its the first database partition. Similarly the second node is the SAP CRM central instance and its first database partition. The third acts as the application server for the SAP business information warehouse (BW) and contains its first database partition. Each of the application server nodes hold only the data for the actual application while the other 37 servers are equally distributed with all three databases.

The "FACT" tables for SAP BW are spread across the nodes while the SAP data dictionary resides on the first node.



Not all components included in this DB2 ICE blueprint are necessary for a minimum configuration but this provides the overview of the options available and discussed in the earlier sections.

---

## Blueprint Architects

---



### **DB2** Information Management Software

IBM's DB2 database software is the worldwide market share leader in the industry. The integrated strength of IBM's data management products, based on DB2, powers the industry's most demanding solutions.

<http://ibm.com/db2>

### @server xSeries

The xSeries is IBM's line of scalable industry-standard, Intel-based servers that enable customers to confidently run their e-businesses. The xSeries leverages IBM's X-Architecture leadership to provide highly scalable servers.

<http://ibm.com/eserver>

### **Tivoli** software

Tivoli software from IBM enables an IT organization to reduce the total cost of ownership and improve service levels by helping manage security, storage, performance and availability, and configuration and operations.

<http://ibm.com/tivoli>

### **WebSphere** software

IBM WebSphere is the market-leading Internet infrastructure software, or middleware, for creating, running and integrating e-business applications across a variety of computing platforms. Built on open standards such as J2EE™, XML and the new Web services standards

<http://ibm.com/websphere>



Mellanox Technologies is a leading supplier of InfiniBand semiconductors; providing switch, host and target channel adapter (HCA/TCA) devices to the server, communications and storage markets.

<http://www.mellanox.com>



Founded in 1972, SAP is the recognized leader in providing collaborative e-business solutions for all types of industries and for every major market.

<http://www.sap.com>



SuSE Linux is the international technology leader and solutions provider in Open Source operating system software. SuSE Linux is a privately held company focused entirely on supporting the Linux community.

<http://www.suse.com>



Voltaire provides intelligent connectivity solutions that enable data center managers to take full advantage of InfiniBand, the next generation I/O standard, when connecting to existing networks.

<http://www.voltaire.com>

In the following pages we provide details about the various components included in this blueprint.

**DB2** Information Management  
Software

### **IBM DB2 Universal Database**

In today's data-driven business world, you want to leverage your information assets in new and innovative ways for competitive gain. But you can't afford to spend too much time or devote too many resources to managing your database. With IBM DB2 Universal Database, Version 8.1, you don't have to.

IBM DB2 is a self-managing and self-tuning relational database management system, strong enough to meet the demands of large corporations and flexible enough to serve medium-sized and small e-businesses. IBM DB2 provides database technology leadership, easing the complexity and cost of managing e-business applications, and laying the foundation for integrating information across the enterprise.

DB2 is easier to use and administer than ever, with new self-management capabilities that automate many of the tasks traditionally associated with maintaining an enterprise-class database. It continues to boast the price/performance that has made firm believers out of 400,000 companies and more than 16,000 IBM Business Partners. No wonder IBM has been named the worldwide leader in database management system software.

### **Central to your most critical business applications**

As the nucleus of many e-businesses, DB2 is optimized for key applications. It is ideal foundation for information integration, business intelligence, transaction processing, Web-based solutions as well as a backend for packaged solutions like ERP, CRM, and SCM. DB2 has the muscle to manage the information contained in the largest online catalogs. And with its support for the latest Web services standards and Internet protocols, you can easily share your data with partners, customers and suppliers.

To analyze customer behavior and explore marketing trends, you'll have plenty of high-powered business intelligence support with DB2 data management software from IBM. IBM DB2 Warehouse Edition provides a framework for building, managing and accessing largest of data warehouses. With IBM DB2 Intelligent Miner™ data mining offerings, you can uncover patterns and trends—even in realtime—that can enhance decision making, helping you strengthen customer relationships and increase your competitive advantage.

Your business information comes in various formats from multiple sources and platforms—especially as your enterprise expands. IBM is at

the forefront of integration technology, and allows you to combine XML with powerful data integration tools across relational databases, flat files, spreadsheets, Web services and more.

### **Easy to maintain with minimal resources**

IBM has developed the latest version of DB2 with self-managing and resource tuning technology. This SMART technology brings self-administration, self-tuning and self-healing properties to DB2, minimizing DBA intervention and maintenance needs. The DB2 Control Center adds another element of management convenience. From one terminal, DB2 Control Center provides a common, integrated toolset for managing local and remote databases across different software and client hardware platforms.

You can also take advantage of two new features to monitor the well-being of your DB2 system. The Health Monitor and Health Center alert you to potential system issues, so you can address them before they become harmful problems. The Health Monitor is a server-side tool that provides constant monitoring without user interaction. The Health Center is the graphical interface to the Health Monitor, and you can use it to configure the Monitor.

### **DB2 and Linux**

DB2 has long been known for its technology leadership. Therefore it was not surprising when IBM took the lead in bringing DB2's proven performance, scalability and ease of use features to clustered Linux environments. The reasons why major companies and governments around the globe choose to deploy DB2 for Linux in the enterprise setting are quite simple. DB2's rich set of features have been running on Linux longer than other commercial databases on Linux, and during this time while the Linux kernel matured through the efforts of thousands of programmers and volunteers, the IBM teams were busy further hardening the kernel and DB2 on Linux for enterprise workloads.

### **A solid, all-around value**

While redefining information management through its innovative technologies, IBM remains committed to delivering value. DB2 offers one of the market's lowest-TCO data management solutions. From its scalability to its fast, reliable performance and easy development environment, DB2 can rock your world and help you win the e-business game.



The combination of IBM eServer xSeries and IBM DB2 Universal Database offers an elegant foundation for Linux based database solutions. xSeries and DB2 UDB will help customers realize new levels of productivity while lowering total cost of ownership resulting in faster return on investment. With xSeries and DB2 UDB customers will enjoy a foundation that delivers industry-leading availability, scalability, and performance for mission-critical database solutions.

### **Advanced server technology for DB2 UDB**

The xSeries server design leverages decades of experience IBM has gained in servers - building servers that are among the leaders in performance, scalability, reliability, and manageability, just right for demanding DB2 Universal Database solution. xSeries is setting the bar for within the industry-standard space delivering new levels of availability and performance through unique mainframe-inspired innovations which are designed to help minimize solution deployment risk while attaining new levels of productivity for database implementations.

IBM X-Architecture was introduced in 1998 as a technology blueprint to bring down mainframe inspired technology into the industry standard space. This offers database solutions enhanced availability, scalability, and systems management – all on industry-standard hardware. Over the past several years, IBM xSeries has continuously delivered against this blueprint introducing technologies such as Active PCI, C2T Interconnect™ cabling, Chipkill™ memory, Predictive Failure Analysis®, Light Path Diagnostics™, and IBM Director Software Rejuvenation to name a few.

Continuing to build on the IBM xSeries commitment of bringing mainframe inspired technology into the industry-standard space, IBM announced in 2001 the next wave of innovation for the Intel-based server platform. Called Enterprise X-Architecture™ (EXA), it combines industry-standard features with IBM mainframe-inspired capabilities to produce revolutionary advances in the I/O, memory, and performance of IBM xSeries servers. Some of the unique technologies enabled through EXA are: XpandOnDemand™ Scalability, Remote I/O, Real Time Diagnostics, Xcel4™ Server Accelerator Cache, Active Memory™, and Memory ProteXion™. The results of EXA helps bring more powerful and easier-to-manage servers with significantly lower total cost of ownership helping delivering concrete, measurable ROI and operational efficiencies.

### **Broad Range of Servers**

The xSeries servers with mainframe-inspired innovations at their core, help customers do more with their IT dollars, delivering against the critical areas of availability, reliability and performance, setting a new standard for Intel-based servers powering database solutions. In addition, the this powerful line up of servers introduces new levels of flexibility helping you to establish a server architecture that will properly handle your growing database needs as you add new users for your database solution.

IBM xSeries offers a broad innovative product line ranging from affordable, edge-of-network servers to enterprise-class servers. Below is a summary of the product categories and servers offered by IBM xSeries.

**High Performance Servers** – A range of high-performance Intel based servers built on revolutionizing Enterprise X-Architecture technology with scalability up to 16-way SMP. Servers include: x360, x440, and x450.

**Rack-Optimized** – A line of highly rack-optimized servers ranging from Uni to 2-way SMP for customers who want more compute power in less space. Servers include x305, x335, and x345.

**Universal Servers** – Uni to 4-way Universal servers optimized for internal storage for small businesses, departmental, and distributed workloads. Servers include x205, x225, x235, and x255.

**BladeCenter** – An innovative technology that introduces servers designed as cards that easily plug into a 7U chassis capable of holding fourteen 2-way Intel Xeon Processor DP servers, a capability especially attractive for datacenters deploying scale-out solutions.

The IBM eServer xSeries servers running DB2 UDB solutions offers an elegant foundation for businesses of many sizes and complexities.

### **OnForever™ – Driving Increased Availability and Reliability**

xSeries has set the bar for the industry on providing the most reliable servers and the best tools to manage these servers. The IBM OnForever initiative is designed to provide customers with the ultimate in availability and reliability – ultimately driving towards continuous availability. IBM xSeries servers are designed from the ground up for high Reliability and Availability, along with ease of Serviceability (RAS).



The primary goal of the IBM OnForever initiative is to bring xSeries servers the level of availability that is traditionally associated with IBM mainframes. The focus of OnForever is on the system-level solution, designed to provide even higher availability and greater manageability for the entire system, including hardware, operating systems, middleware and applications. Not only do these mainframe-based innovations help xSeries servers prevent most outages—and recover rapidly from the few that do occur—they also make xSeries servers simpler to service, and easier to manage, while using industry-standard hardware and operating systems. These high-availability features help keep customers' servers running continuously—a requirement of databases supporting mission-critical solution environments that run around the clock and is critical to the competitiveness and results of a business.

### **Self Managing Capabilities**

A business' success can be dependent on how well their database application runs, therefore it is critical that the server it's running on supports proactive tools to manage this IT environment. Without these tools there will be devastating impacts on the bottom line. xSeries systems management tools simplify—and in many cases, automate—performance planning, preventive maintenance, diagnostic monitoring and event responses to help maintain consistent, high levels of server productivity. As a result, xSeries servers protect customers' businesses from the devastating effects of downtime, and help keep overall maintenance costs incredibly low – an important factor for almost all database solution environments.

IBM xSeries offers systems management tools in a combination of hardware instrumentation and advanced software – helping to mitigate risks and costs while ensuring rapid deployment and consistent proactive and predictive management of your database solution. With IBM Director as the core systems management tool, xSeries servers are easy to administer and manage in a mission-critical database solution environment. Key IBM Director extensions such as Capacity Manager, Software Rejuvenation, Real Time Diagnostics, Process Control brings new levels of manageability to solution environments helping to reduce total cost of ownership with improved return on investment through increased uptime.

Another key component within the systems management picture is the IBM Remote Supervisor Adapter. This adapter is standard on select servers and provides capabilities to manage the server remotely even if the OS is not functioning. Many people like to think of this adapter as a computer within a computer – dedicated to monitoring the health of the

server. It's a key element in helping reduce the total cost of ownership for the server, which in turn lowers the total cost of ownership for the entire database solution.

**Strong support for next generation e-business on xSeries & Linux**

IBM is a recognized leader in the Linux community and the premier provider of end-to-end Linux solutions: Linux-ready hardware, software, and support. With IBM's leadership position, IBM xSeries servers are a safe, smart choice to run Linux DB2 UDB database solutions. In addition, xSeries leverages IBM's unique solution development network which includes Linux Technology Centers, Linux Development Centers, Linux Competency Centers, and Solution partnership Centers delivering expertise for porting, configuration testing, development and ISV solution validation.



### IBM DB2 and InfiniBand – A Winning Combination


The clustering capabilities of IBM DB2 Universal Database enable exceptional database performance with the added benefit of superb reliability. Of course the clustering interconnect needs to deliver these same features and InfiniBand<sup>SM</sup> delivers both performance and RAS (reliability, availability, and serviceability).

### Delivering Usable 10 Gbps Performance

While other technologies boast of 10Gbps link speeds only InfiniBand delivers usable 10Gbps bandwidth with minimal CPU overhead. Mellanox HCAs offer hardware transport and kernel bypass to make server CPU power available to deliver the maximum DB2 performance, without wasting precious cycles on delivering reliable communications between servers. Furthermore highly integrated switches with low cost copper cabling deliver the most cost effective 10Gbps technology available.

### Ease of Integration

PCI /PCI-X based HCA boards allow InfiniBand connections to be simply and quickly integrated into existing high performance server platforms. Furthermore with full support for all TCP/IP and UDP based processes, InfiniBand integrates simply and quickly into existing data center environments. When adopting new technologies such as InfiniBand it is always reassuring to know that the existing software infrastructure just works, without any changes.



InfiniBand offers more than backwards compatibility and simple integration. The real benefits of InfiniBand are recognized by taking advantage of advanced features of the technology such as kernel bypass

and RDMA (Remote Direct Memory Access). Again InfiniBand technology and DB2 complement each other as the highly scalable DB2 Universal Database Architecture is designed to be able to take advantage of these high performance InfiniBand capabilities.

### **Advanced Management and RAS**

InfiniBand technology offers Management and RAS (Reliability Availability, and Serviceability) capabilities at every level of the architecture from hot swappable modules, embedded intelligent health monitoring agents, to failover capabilities. The InfiniBand Architecture achieves new level of manageability by incorporating intelligent management agents in each channel adapter and switch element in the cluster network. These mandatory management agents make detailed transaction level information available to application management software. Providing this information at unprecedented levels enables management packages to become proactive - recognizing and fixing problems before they occur. Furthermore, with the new level of information available root cause analysis can be performed at the application and transaction level insuring even higher levels of transactional reliability. InfiniBand includes advanced failover mechanisms such as APM (Automatic Path Migration) that enable a self-healing InfiniBand fabric. The DB2 Universal Database can take advantage of these powerful management and RAS mechanisms and integrate this with state of the art management tools to provide the highest levels of robustness and uptime.

### **Advanced Platforms for DB2**

Today InfiniBand and DB2 can be delivered on standard server platforms using traditional PCI/PCIX HCA add-in cards. Even greater performance, density, and Total Cost of Ownership (TCO) benefits are just around the corner with new modular InfiniBand compute platforms from the major server vendors. A vision of how these new platforms might look can be gleaned from the Mellanox Nitro II server blades reference design.



The Nitro II incorporates three powerful components into a highly integrated modular chassis. These three components are: Up to twelve 4U high server blades, dual 16-Port 10Gb/s switch blades, and integrated chassis with passive backplane supporting 480 Gb/sec of aggregate bandwidth. Commercial platforms assembled from industry standard components similar to this will support the DB2 Universal Database and will be able to deliver new levels of scalability, performance, and reliability. Diskless server blades are made possible by InfiniBand's low-latency hardware transport, thereby overcoming the latency and bandwidth penalties of LAN-based remote storage. Eliminating the local disk improves reliability, lowers cost and enables more power for improved CPU and memory performance. Such highly integrated InfiniBand server blade platforms eliminate cabling, reduce power and cooling overhead, and deliver better total cost of ownership and will once again showcase the synergy between InfiniBand technology and the IBM DB2 Universal Database.



The various SAP applications rely on a database to store the actual user data. While the SAP application code, better known under the name of its 4GL programming language ABAP is actually stored in the database and is limited in its size -- the large user data (like part numbers), customer data, orders and more are the critical component where the database to application server relation as well as the actual scalability of the database plays a critical role.

The introduction of newer applications such as SAP BW business information warehouse, customer relationship management and supply chain functionalities with SAP's Advanced Planner and Optimizer put the database to the test even more.

Database users and developers have always sought to break the barriers imposed by traditional technologies in disk and I/O hardware. For the past decades, I/O throughput was the number one limitation for response times of database queries; these are, of course, a core component of every solution accessing enterprise-wide data resources.

Even with larger cache and memory available to database servers, the deployment of enterprise-wide OLTP applications, like SAP R/3, is ultimately constrained by the capabilities of storage devices.

The industry leader in database technology, IBM DB2 UDB is poised to be the leading database technology for OLTP and OLAP. With OLAP applications, such as SAP's Business Information Warehouse, additional challenges are created for database vendors.

SAP NetWeaver (formerly known as mySAP Technology) is the technical foundation for all SAP applications as well as partners' products and solutions. One of the core components of SAP NetWeaver is SAP BW.

SAP BW delivers enterprise-wide data warehousing, a business intelligence platform and a suite of business intelligence tools. It was designed from the ground up to provide superior “modelling-the-business” tools and methodologies. Following the main functional areas:

*Data warehousing* forming the application-neutral foundation for SAP Business Intelligence (SAP BI). SAP BW supports the complete data warehousing process, from data integration, data transformation, consolidation and cleansing to data provision for analysis.

*Business intelligence platform* serving as the technological infrastructure to support information access and comprehensive analytics.

*Business intelligence suite* transforming data into insightful information and feed a wide variety of users for decision-making.

As part of the SAP NetWeaver architecture, SAP BW draws from and utilizes the capabilities of the other components for business intelligence. The sum of the functionality of SAP BW and the contribution of other components of SAP NetWeaver form a BI platform that represents the next major step in the evolution of business intelligence.

SAP Business Intelligence (SAP BI) is an integrated end-to-end solution that transforms data into information. The SAP business content contains predefined templates for reports and analyses giving you faster access to your business information. Some of these objects are industry specific and some are domain specific for areas such as customer relationship management and supply chain management. The predefined business content speeds up the implementation of your data warehouse solution.

SAP BW supports collaborative BI scenarios with knowledge management technologies such as collaboration and content management. This means that you and your team can acquire the insights and knowledge you need to make the right business decisions. In addition SAP BI acts as the platform to develop, extend, and deploy analytical applications delivered through mySAP business suite. SAP analytical applications run on data stored in the data warehouse, and leverage the tools and technologies of SAP BI. In addition to integrating business processes, they provide predefined “closed-loop” business

scenarios and predefined metrics that measure the effectiveness of business operations.

SAP analytical applications improve business processes across a company's entire value chain, like Customer Relationship Management (CRM), Supply Chain Management (SCM), Financials and Human Resources (HR).

Clustering with DB2 UDB is currently the only database solution with parallel capabilities supported by SAP BW. This technology enables scalable performance on large databases by breaking the processing into separate execution components that can be run concurrently on multiple processors. Processing the individual queries with database utility operations in parallel can dramatically reduce elapsed times for queries and database operations. The demand for mySAP.com BW for this show-case implementation was generated by wide-spread deployment of this solution with DB2 and its extremely scalable and high-performing database power. The generation of data cubes with 10,000 data sets that spread in several dimensions demands a scalable architecture like DB2 UDB. Nonetheless, the same implementation can be done using a transactional mySAP.com system and any other application server based solution.

The data warehouse with its FACT tables was the choice to start with for an InfiniBand-based implementation. The flawless and fast deployment allowed the testing of mySAP.com – the well-known ERP package from SAP – as an additional product. In contrast to the OLAP based data warehouse, mySAP.com 4.6c is an OLTP application with different characteristics. The SAP R/3 uses the same layout for the data dictionary and table spaces but has a load on table spaces such as STAB, BTAB and POOL that are not used for SAP BW. A simple distribution of these table spaces showed good scalability of the SAP overall installation.

Beyond the SAP BW solution stack a new feature of SAP's technology set is used called MCOD. MCOD allows loading multiple SAP instances into one single database instance. For the showcase we add to the existing SAP BW configuration a SAP R/3 and a CRM system to show the ease of implementing SAP in a clustered database environment.





### **Ready for enterprise business SuSE Linux Enterprise Server**

Proven security, scalability, reliability, standards compliance, maintenance and support - these features define the SuSE Linux Enterprise Server. Enterprises of all sizes can benefit from an incomparable offering in the market: The SuSE Linux Enterprise Server combines maximum reliability and performance with unique scalability, being supported across the whole range of IBM eServer hardware, plus all of the major hardware manufacturers.

Founded 11 years ago in 1992, SuSE has built an enviable reputation of delivering some of the most technically advanced Linux solutions for business of all sizes. With the ability to support DB2 on the largest range of hardware, from the smallest server to the largest 64-bit mainframe, SuSE offers true flexibility when choosing the right solution.

SuSE Linux Enterprise Server and DB2 are a combination boasting excellent performance and optimized for high availability. SuSE has chosen DB2 for Linux to enhance their own internal CRM solution for responding to customer support requests. Using DB2, SuSE has achieved the optimum level of reliability and scalability. Measurably better customer service and therefore, higher customer satisfaction, are the key business benefits of the SuSE Linux Enterprise Server / DB2 implementation. Along with higher customer satisfaction comes more effective communications and higher profitability.

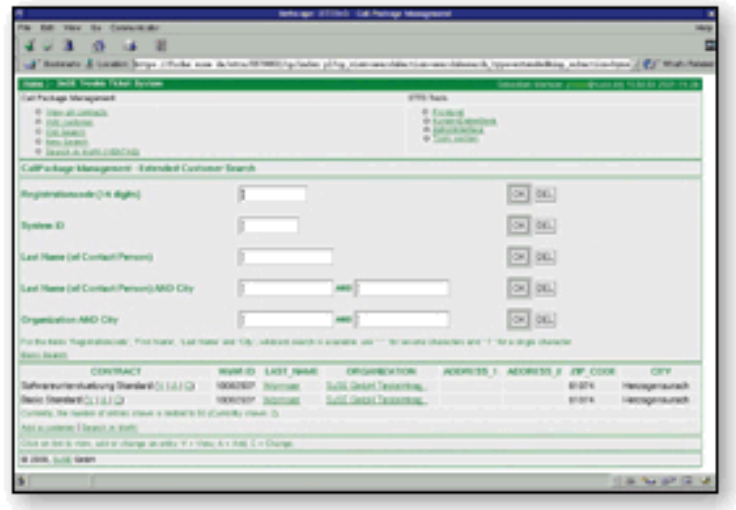


### **Groundbreaking Enterprise Linux Solution**

By certifying on enterprise platforms of leading ISVs and IHVs and by implementing and maintaining the entire major industry standards, the SuSE Linux Enterprise Server is a groundbreaking enterprise Linux solution. The planning for SuSE Linux Enterprise Server implementations is made easier by the one-year release cycle and comprehensive maintenance services. SuSE provides a tailor made maintenance service for the supported hardware and software platforms, delivering only relevant, certified patches and fixes that may be required. Quality assurance is provided by comprehensive testing in all server environments and close development ties with major industry partners. Combined with SuSE's extensive know-how in support, consulting and training services, the SuSE Linux Enterprise Server is the ideal choice for Linux deployment in all enterprises.

### SuSE sets standards in the quality of Open Source solutions

The strength of Open Source is the Linux community's productivity and responsiveness. With the largest, most experienced Open Source development team worldwide, SuSE is one of the major contributors to the Linux community and drives innovation in Linux. SuSE continues to develop new features such as memory management, high availability, performance features, large file support and journaling file systems: all necessary to run databases efficiently.



*With DB2 Universal Database the SuSE Trouble Ticket System has been trouble-free and available 24 hours a day, 7 days a week to help customers solve their problems.*

In order to consolidate the broad range of Linux features and functions into an easy to install and manage package, SuSE has created a development and quality assurance process named "AutoBuild", which is totally unique in the Linux market. "AutoBuild" guarantees a certifiable production process for new products, resulting in outstanding quality, rapid development times and the flexibility to deliver support for varied hardware platforms using the same source code. Using the same source code enables SuSE to provide support for multiple hardware types, while maintaining the same software interfaces. Hence migrating applications from one hardware platform to another is often simply a case of recompiling.

### **SuSE Linux Enterprise Server - Available across all relevant hardware platforms**

SuSE is the ONLY operating system provider worldwide, who delivers an Enterprise Server operating system across all of the relevant hardware architectures. In January 2002, IBM commissioned the Linux experts at SuSE to develop and maintain the SuSE Linux Enterprise Server operating system environment for the entire IBM eServer product line. In another agreement, announced in March 2002, IBM Global Services and SuSE Linux offer global Linux services and support for corporate customers, ranging from joint offerings for Linux solutions and training to worldwide support and maintenance services. By working closely with IBM development teams for both software and hardware, SuSE and IGS deliver world-class support.

SuSE also believes that Linux solutions should be easy to install, so, wherever possible and relevant, the SuSE Linux Enterprise Server ships proprietary hardware drivers in binary. This saves lost time searching for modules and recompiling for specific hardware, and has the extra benefit that the total package has been assured by SuSE.

### **Investment-security and Future-proofing**

In contrast to conventional Linux products, the SuSE Linux Enterprise Server provides an actively maintained server operating system optimized for business deployment. SuSE makes all of the relevant enhancements available for SuSE Linux Enterprise Server customers with certified, quality assured and well-documented updates. For the first time, enterprise customers have long term planning and investment security through clearly defined release cycles and the SuSE system maintenance guarantee. As SuSE ensures that the configuration of system interfaces remains unchanged after updates, companies can operate their IT infrastructure knowing that applications are still certified and that high-level performance, extraordinary reliability and maximum security are enhanced by the updates.

### **Certified for business solutions of all major ISVs and IHVs**

Due to the high quality and cross-platform availability of SuSE Linux products, leading ISVs and IHVs, such as IBM, SAP certify their products for the SuSE Linux Enterprise Server. SuSE has worked with the industry leaders for some years, and SuSE Linux customers benefit from this close co-operation. For instance, SuSE is the only Linux distribution to offer a certified DB2 / Linux environment across all of the IBM eServer platforms.

In September 2002 SuSE became the first Linux partner to reach the status of SAP Global Technology Partner', another milestone in the delivery of customer-focused solutions.



As co-founder of the UnitedLinux initiative, SuSE has been actively involved in the project. The goal is to streamline Linux development and certification around a global, uniform business distribution of Linux. With UnitedLinux based products, ISVs and IHVs can concentrate their development and certification efforts on a major business Linux distribution. By reducing the number of separate Linux distributions, ISVs and IHVs can reduce the effort and expense of developing solutions for several different Linux environments. Business customers will benefit by having a superior, certified Linux distribution, delivering quality, scalability and continuity that is tuned to businesses most pressing needs.

Within the initiative, SuSE performs the role of the product integrator, being responsible for the development and the quality assurance of the final product releases that are based on the high quality of the SuSE Linux Enterprise Server. Through the company's role within the UnitedLinux initiative, SuSE is enabling the future development of Linux as the leading enterprise operating system.



The system and solution management is an important piece of the overall solution. The broad variety of Tivoli solution elements implemented in the cluster enriches the usability and reduces cost.

### **IBM Tivoli Storage Manger**

IBM Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. Scaling to protect hundreds of computers running a dozen OS ranging from laptops to mainframes and whether connected together via the internet, WAN or LAN, Storage Manager's centralized Web-based management, smart-data-move and store techniques and comprehensive policy-based automation all work together to provide companies with a powerful suite of automated tools for enterprise-wide data protection that helps to minimize storage administration costs and the impact to both computers and networks. Tivoli Storage Manager helps companies to save valuable resources by providing comprehensive, resilient data protection, simplifying and automating data protection, and providing integrated storage across disparate systems.

### **IBM Tivoli Monitoring**

IBM Tivoli Monitoring is a robust monitoring engine that provides monitoring for essential system resources, to detect bottlenecks and potential problems, and to automatically recover from these critical situations. Tivoli Monitoring helps companies analyze and predict system problems and can be configured to alert administrators to critical system events. Tivoli Monitoring also helps companies to not only monitor critical systems resources but also to detect and automatically respond to potential problems and bottlenecks. The proactive respond to the identified environmental situations resolves problematic events and reduces the overall cost of systems management

Tivoli Monitoring saves system administrators from manually scanning through extensive performance data before problems can be solved. Using industry best practices, Tivoli Monitoring can provide immediate value to the enterprise.

### **IBM Tivoli System Automation**

IBM Tivoli System Automation for Linux provides self-healing capabilities for business applications by detecting failing IT components and then repairing or working around those failures. It manages the availability of business applications according to customer-defined goals. This solution is able to run on a single Linux system or in an xSeries clustered environment.

The IBM Tivoli System Automation for Linux policy is based on resource information, groups of resources, and relationships. Scripts and procedures are not required, new resources or systems can be added without re-writing scripts. This helps ease application growth and application scaling. Tivoli Systems Automation provides additional value through fast detection of outages, sophisticated knowledge about application components and their relationships and quick and consistent recovery of failed resources and applications either in place or on another system of a Linux cluster

IBM Tivoli Systems Automation also performs activities such as shutting down the remainder of an application or even a whole failed system, moving a TCP/IP addresses, or recovering data from a log.

### **IBM Tivoli NetView**

IBM Tivoli NetView helps companies maintain system availability and streamline support by discovering and consolidating TCP/IP networks, displaying network topologies, and managing events and SNMP traps. It also helps monitor network health and gathers performance data from across the enterprise providing a single platform for automating problem diagnosis. It can quickly guide support personnel to an appropriate response or even take corrective action automatically. Tivoli NetView helps to reduce support and training cost for your IT infrastructure by enabling small teams to manager large environments.



Voltaire provides intelligent connectivity solutions that enable data center managers to seamlessly deploy InfiniBand fabrics and take full advantage of the powerful capabilities of the high-speed industry standard fabric. Voltaire's flagship product line is a family of modular InfiniBand Switch Routers based on Voltaire's unique TCP Termination Architecture, which enables corporate data centers to maximize the benefits of InfiniBand when connecting to IP networks. Voltaire's routers are located at the InfiniBand fabric edge and act as a transparent interface between existing network devices, specifically the TCP/IP infrastructure, and InfiniBand-enabled server and storage devices.

Voltaire's end-to-end InfiniBand solutions also include a powerful software stack. The simple to install stack contains a range of drivers and protocols - from basic drivers such as IP over InfiniBand (IPoIB) to powerful upper layer protocols such as Socket Direct Protocol (SDP). SDP is an industry-standard high-performance RDMA protocol and is an important part of the InfiniBand specification. Voltaire's SDP implementation provides a high-bandwidth, low-latency communication while reducing the networking overhead on the servers' CPU to minimal levels. This combination improves applications' performance in orders of magnitude when compared with Ethernet based fabrics. In addition SDP is based on the same Sockets application interface that TCP/IP uses and therefore does not require any changes in existing applications.

Voltaire switch routers process and fully terminate incoming TCP/IP traffic and then pass it to the servers or storage devices using native InfiniBand protocols, offering much better CPU utilization and greater bandwidth compared to other connectivity solutions that simply bridge or tunnel the IP traffic. Voltaire routers do not require the use of any proprietary drivers, they leverage on standard protocols such as TCP/IP and SDP.

The combination of Voltaire's modular switch routers and its powerful software stack has a strong and unique value proposition:


**High-Performance connectivity based on an industry standard with no changes required in applications!**

The Voltaire TCP Termination Architecture lends itself to applications such as load balancing, automatic server fail-over mechanisms, packet filtering, traffic management, traffic monitoring, server/network partitioning and provisioning. As a result, it enables the implementation of the concept of clustering as needed.

**InfiniBand based clustering**

Voltaire's InfiniBand database cluster designated products combined with entry-level servers provide IT managers with a powerful and

scalable architecture for high volume, high transaction applications such as distributed databases, reducing Total Cost of Ownership of database clusters by up to 80 percent!



In addition to their optimized IP to InfiniBand capabilities, Voltaire's switch routers offer

InfiniBand scalable switching capabilities. These two functionalities make Voltaire routers ideally suited for InfiniBand database clusters. The switching capability serves for internal communication between the database nodes while the routing capability connects the cluster to the rest of the datacenter infrastructure. The communication with the application server is dramatically improved through the use of Voltaire routers. Furthermore, Voltaire's TCP termination architecture allows database services such as backup and replication to use InfiniBand native protocols such as SDP significantly reducing the overhead that these services generate.

### **Application Server Clustering: Robust Simplicity**

During recent years the complexity of application servers has increased steadily. Although considered "stateless", significant amounts of temporary information (typically associated with user information, cookies, carts, etc) is scattered among the application servers and some of the networking devices such as the load balancers. This has created a reality in which traffic routing to the applications servers is complex and the temporary data is not protected from failures as it is randomly scattered among the servers.

Through the use of its RDMA capabilities InfiniBand offers new capabilities to Tier II, which simplify connectivity and protect the temporary data from being lost upon failures. In this solution application servers share and maintain the temporary data structures in complete



synchronization, while the rest of the application servers access this data on these servers whenever needed.

InfiniBand is the only technology that enables such a solution due to its low-latency capabilities. RDMA operations will serve both for the temporary data synchronization for high-availability purposes and for access by the rest of the application servers.




### **Enabling Utility Computing Solutions**

This model enables the application server to be completely “stateless”, setting the ground for powerful utility computing solutions to be implemented in Tier II. Such solutions will allow the dynamic provisioning of servers to different applications according to the data-center needs. For example the majority of application servers can be allocated during business hours to process consumer transactions while part of these servers will be re-allocated to run billing applications during night-time. This model dramatically increases the utilization of the data-center infrastructure.

### **Connecting it all to the existing infrastructure**

It is clear that when a new technology such as InfiniBand is being deployed in the data-center, the issue of integrating it with the existing infrastructure becomes critical. On one hand simplicity is required - in other words, seamless connectivity between the new InfiniBand fabric and the existing IP infrastructure. On the other hand InfiniBand fabrics must use InfiniBand native protocols such as SDP to fully utilize the benefits brought by the technology. The solution that answers both these needs is Voltaire’s InfiniBand Switch Router that fully terminates TCP/IP traffic and transforms it to SDP sessions going into the InfiniBand enabled servers.



© Copyright IBM Corporation 2003  
IBM Canada  
8200 Warden Avenue  
Markham, ON  
L6G 1C7  
Canada

Printed in United States of America  
5-03  
All Rights Reserved.

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

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.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of 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.

The information in this white paper is provided AS IS without warranty. Such information was obtained from publicly available sources, is current as of 01/30/2003, and is subject to change. Any performance data included in the paper was obtained in the specific operating environment and is provided as an illustration. Performance in other operating environments may vary. More specific information about the capabilities of products described should be obtained from the suppliers of those products.