



IBM Linux Technology Center

HPC Linux Systems

ScicomP 9 – Bologna, Italy

Ken Rozendal
Linux Technology Center
March 26nd, 2004

Special Notices

This presentation was produced in the United States. IBM may not offer the products, programs, services or features discussed herein in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products, programs, services, and features available in your area. Any reference to an IBM product, program, service or feature is not intended to state or imply that only IBM's product, program, service or feature may be used. Any functionally equivalent product, program, service or feature that does not infringe on any of IBM's intellectual property rights may be used instead of the IBM product, program, service or feature.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. Sources for non-IBM list prices and performance numbers are taken from publicly available information including D.H. Brown, vendor announcements, vendor www Home Pages, SPEC Home Page, GPC (Graphics Processing Council) Home Page and TPC (Transaction Processing Performance Council) Home Page. IBM has not tested these 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.

IBM may have patents or pending patent applications covering subject matter in this presentation. The furnishing of this presentation does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of a specific Statement of General Direction.

The information contained in this presentation has not been submitted to any formal IBM test and is distributed "AS IS". While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. The use of this information or the implementation of any techniques described herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. Customers attempting to adapt these techniques to their own environments do so at their own risk.

IBM is not responsible for printing errors in this presentation that result in pricing or information inaccuracies.

The information contained in this presentation represents the current views of IBM on the issues discussed as of the date of publication.

IBM cannot guarantee the accuracy of any information presented after the date of publication.

All prices shown are IBM's suggested list prices; dealer prices may vary.

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

Information about non-IBM products was obtained from suppliers of those products. IBM makes no representations or warranties regarding these products. Non-IBM products are offered and warranted by third-parties, not IBM.

Special Notices

Information provided in this presentation and information contained on IBM's past and present Year 2000 Internet Web site pages regarding products and services offered by IBM and its subsidiaries are "Year 2000 Readiness Disclosures" under the Year 2000 Information and Readiness Disclosure Act of 1998, a U.S. statute enacted on October 19, 1998. Internet Web site pages have been and will continue to be our primary mechanism for communicating year 2000 information. Please see the "legal" icon on IBM's Year 2000 Web site (www.ibm.com/year2000) for further information regarding this statute and its applicability to IBM.

Any performance data contained in this presentation was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements quoted in this presentation may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this presentation should verify the applicable data for their specific environment. The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIXwindows, AS/400, C Set++, CICS, CICS/6000, DataHub, DataJoiner, DB2, DEEP BLUE, DYNIX, DYNIX/ptx, e(logo), ESCON, IBM, IBM(logo), Information Warehouse, Intellistation, IQ-Link, LANStreamer, LoadLeveler, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NUMA-Q, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, POWERparallel, PowerPC, PowerPC(logo), ptx/ADMIN, RISC System/6000, RS/6000, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, SP2, System/390, The Engines of e-business, ThinkPad, Tivoli(logo), TURBOWAYS, VisualAge, WebSphere. The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: AIX/L, AIX/L(logo), AIX PVMe, Application Region Manager, AS/400e, Blue Gene, Chipkill, ClusterProven, DB2 OLAP Server, DB2 Universal Database, e-business (logo), ^, GigaProcessor, HACMP/6000, Intelligent Miner, iSeries, Network Station, NUMACenter, PowerPC Architecture, PowerPC 604, POWER2 Architecture, pSeries, Sequent (logo), SequentLINK, Service Director, Shark, SmoothStart, SP, Tivoli Enterprise, TME 10, Videocharger, Visualization Data Explorer, xSeries, zSeries. A full list of U.S. trademarks owned by IBM may be found at <http://iplswww.nas.ibm.com/wpts/trademarks/trademar.htm>. Lotus and Lotus Notes are registered trademarks and Domino and Notes are trademarks of Lotus Development Corporation in the United States and/or other countries.

NetView, Tivoli and TME are registered trademarks and TME Enterprise is a trademark of Tivoli Systems, Inc. in the United States and/or other countries. Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States and/or other countries. UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group. LINUX is a registered trademark of Linus Torvalds. Intel and Pentium are registered trademarks and MMX, Itanium, Pentium II Xeon and Pentium III Xeon are trademarks of Intel Corporation in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries. Other company, product and service names may be trademarks or service marks of others.

Agenda

- HPC Cluster Roadmap
- IBM Linux Strategy
- LTC Focus Areas
- Linux on POWER
 - ▶ Motivation
 - ▶ High Level Objectives
 - ▶ Detailed Plan
- LTC Detailed Focus Areas



IBM Linux Technology Center

HPC Cluster Roadmap

High Performance Computing

- A very important segment for IBM
- Goal:
 - Continue to develop complete, balanced, well integrated solutions
 - Continue to deliver highest uptime in industry, supporting on-demand
 - Continue to develop a comprehensive software environment to enable programmers to get maximum delivered TF on actual user problems
- Continue to invest in leading edge HPC technologies
 - ▶ Hardware
 - Processors (Power4+, Power 5, Power 5+, ...)
 - Also support IA-32 and Opteron based clusters
 - CECs (Regatta, Squadrons, ...)
 - Interconnects (eg. Infiniband)
 - ▶ Software
 - Linux operating System
 - Cluster File Systems & Parallel File System (GPFS)
 - Programming Models (MPI, LAPI) & Runtime environments (POE)
 - Resource Management & Scheduling (LL)
 - Application development tools (PE, ESSL, PESSL)
 - Systems Management (CSM)

pSeries Linux HPC Overview

- 4Q03
 - SLES8 on pSeries
 - CSM, GPFS over GigE, ESSL, Parallel ESSL over Myrinet
- 2Q04
 - RedHat3 on pSeries
 - SLES8 on pSeries and JS20
 - CSM, ESSL, Parallel ESSL
 - GPFS and LoadLeveler (SLES8 only)
- 3Q04
 - RedHat3 on JS20 and Power5
- 4Q04
 - SLES9 on Power5, JS20, JS20+
- 1Q05
 - RedHat4 on Power5, JS20, JS20+

pSeries Linux HPC Announcements and Plans

	4Q03	2Q04
User SW	Customer/ISV Parallel Apps.	Customer/ISV Parallel Apps.
HPC SW	<u>GM2.0.8/MPICH-GM 1.2.5.10</u>	<u>LoadLeveler 3.2 *</u>
	<u>ESSL 4.1</u>	GM/MPICH-GM
	<u>Par. ESSL 3.1</u>	<u>ESSL 4.1.1</u>
	<u>GPFS 2.2</u>	<u>Par. ESSL 3.1.1</u>
OS	SLES8 and <u>CSM 1.3.2</u>	GPFS 2.2*
SYSTEMS	p655 P650 P630/615	SLES8, <u>RHAS3</u> , and <u>CSM 1.3.3</u> <small>(Squadrons Support on RHAS3(ESSL/PESSL/CSM</small>
FABRIC	GBEnet Myrinet	p655 P650, <u>ML8 (Squadrons on RH Only)</u> P630/615, <u>SF2, L4 (Squadrons on RH Only)</u> <u>JS20 (SLES8 Only)</u>
	GBEnet Myrinet	GBEnet Myrinet

* GPFS and LL on SLES8 only initially

pSeries Linux HPC Plans

3Q04

4Q04

User SW	Customer/ISV Parallel Apps.		Customer/ISV Parallel Apps.	
HPC SW	LoadLeveler 3.2		LoadLeveler 3.2	
	GM/MPICH-GM		GM/MPICH-GM	
	ESSL 4.1.1		<u>ESSL 4.2</u>	
		Par. ESSL 3.1.1		<u>Par. ESSL 3.2</u>
OS	GPFS 2.2		<u>GPFS 2.3</u>	
	SLES8, RHAS3 and CSM 1.3.3 (Squadrons Support on RHAS3(ESSL/PESSL/CSM))		<u>SLES9, RH AS3, CSM 1.4</u>	
SYSTEMS	p655 P650, ML8 P630/615, SF2, L4 (Squadrons on RH Only) JS20 (SLES8 <u>and RH</u>)		p655 P650, ML8 P630/615, SF2, L4 JS20, <u>JS20+</u>	
FABRIC	GBEnet	Myrinet	GBEnet	Myrinet

4Q03 pSeries Linux Cluster Announcements

Customer/ISV
Parallel Apps.

[GM2.0.8/MPICH-GM 1.2.5.10](#)

[ESSL 4.1](#)

[Par. ESSL 3.1](#)

[GPFS 2.2](#)

SLES8 and [CSM 1.3.2](#)

p655
P650
P630/615

GBEnet

Myrinet

- ESSL 4.1 – same function set as AIX version
- Parallel ESSL 3.1 over Myrinet using MPICH GM

- GPFS 2.2 Restrictions on pLinux
 - Over GigE
 - NO Sanergy Support
 - NO DMAPI
 - NO pLinux Interoperability (i.e. Interoperability Support is for AIX and xLinux ONLY, initially)

- CSM 1.3.2 – same function set as AIX version

2Q04 pSeries Linux Cluster Plans

Customer/ISV Parallel Apps.	
LoadLeveler 3.2 *	
GM/MPICH-GM	
ESSL 4.1.1	
	Par. ESSL 3.1.1
GPFS 2.2*	
SLES8, RHAS3 , and CSM 1.3.3 <small>(Squadrons Support on RHAS3(ESSL/PESSL/CSM)</small>	
p655 P650, ML8 (Squadrons on RH Only) P630/615, SF2, L4 (Squadrons on RH Only) JS20 (SLES8 Only)	
GBEnet	Myrinet

- Adds LoadLeveler 3.2 for basic Job Scheduling over Myrinet to pLinux with restrictions:
 - NO: C/R, Gang/Preemption, Consumable Resources through WLM, DCE/CtSec, Kernel Extension, IBM Switch Adapter Support, POE Jobs
 - SLES8 Only

- ESSL 4.1.1 and Parallel ESSL 3.1.1 on MPICH over Myrinet using MPICH-GM
 - Power based Blades: JS20 on SLES8
 - Other pServers on RH

- CSM 1.3.3 – same function set as AIX version
 - Power based Blades: JS20 on SLES8
 - Other pServers on RH

* GPFS and LL on SLES8 only initially

4Q04 pSeries Linux Cluster Plans

Customer/ISV Parallel Apps.	
LoadLeveler 3.2	
GM/MPICH-GM	
<u>ESSL 4.2</u>	
	<u>Par. ESSL 3.2</u>
<u>GPFS 2.3</u>	
<u>SLES9</u> , RH AS3, <u>CSM 1.4</u>	
p655 P650, ML8 P630/615, SF2, L4 JS20, <u>JS20+</u>	
GBEnet	Myrinet

- Initial pLinux Support on RHEL EL3 for LoadLeveler added, SLES9
- ESSL 4.2 Squadrons/JS20+ Tuning, SLES9
- Parallel ESSL 3.2 -SLES9
- GPFS 2.3 over GigE for all Platforms listed, SLES9
- CSM 1.4 – same function set as AIX version, SLES9

xSeries Linux HPC Announcements and Plans


	4Q03	2Q04	3Q04	4Q04	1H05
IA32 xSeries (including BladeCenter) (Myrinet/ Ind. Std.)	<ul style="list-style-type: none"> ▪SLES 8, RHEL 3, RH 9 ▪CSM 1.3.2 ▪GPFS 2.2 	<ul style="list-style-type: none"> ▪SLES 8, RHEL 3 ▪CSM 1.3.3 ▪LoadLeveler 3.2 (RHEL3 only) ▪GPFS 2.2 	<ul style="list-style-type: none"> ▪SLES 8, RHEL 3 ▪CSM 1.3.3 ▪LoadLeveler 3.2 (RHEL3 only) ▪GPFS 2.2 	<ul style="list-style-type: none"> ▪X336, x346 ▪SLES 9, RHEL 3 ▪CSM 1.4 ▪LoadLeveler 3.2 ▪GPFS 2.3 	<ul style="list-style-type: none"> ▪SLES 9, RHEL 4 ▪CSM 1.4.1 ▪LoadLeveler 3.3 ▪GPFS 2.3
Opteron xSeries (Myrinet/ Ind. Std.)	<ul style="list-style-type: none"> ▪SLES 8, RHEL 3 ▪CSM 1.3.2 	<ul style="list-style-type: none"> ▪SLES 8, RHEL 3 ▪CSM 1.3.3 ▪GPFS 2.2 	<ul style="list-style-type: none"> ▪SLES 8 RHEL 3 ▪CSM 1.3.3 ▪GPFS 2.2 	<ul style="list-style-type: none"> ▪SLES 9, RHEL 3 ▪CSM 1.4 ▪LoadLeveler 3.2 ▪GPFS 2.3 	<ul style="list-style-type: none"> ▪SLES 9, RHEL 4 ▪CSM 1.4.1 ▪LoadLeveler 3.3 ▪GPFS 2.3



IBM Linux Technology Center

IBM's Linux Strategy

Software Group



IBM Global Services

Middleware Enablement

- DB2
- WAS
- MQSeries
- ISV's
- QuickStart

Technical Support

- Support Line - 24/7
- Account advocate
- Advanced Support
- All @server including clusters
- Support for major Linux distributions
- Presales technical support teams

Web Hosting

- xSeries
- zSeries Linux
- e-Sourcing

Training / Learning Services

- 5 languages; 20 countries
- all @server
- Web-based and classroom
- Cluster workshops - New!
- Redbooks
- LPI Certification

Clusters

- Support Line
- x1300 Installations
- OEM Procurement
- Hardware Setup
- Software Installation

Consulting

- Open Source SW Consulting
- Application Migration
- Consult Line service options
- e-business Enablement

Workload Consolidation

- File/Print, Webserving
- Bynari, Sendmail
- Linux Solution for e-Business
- Application Porting
- Solaris to Linux Migration

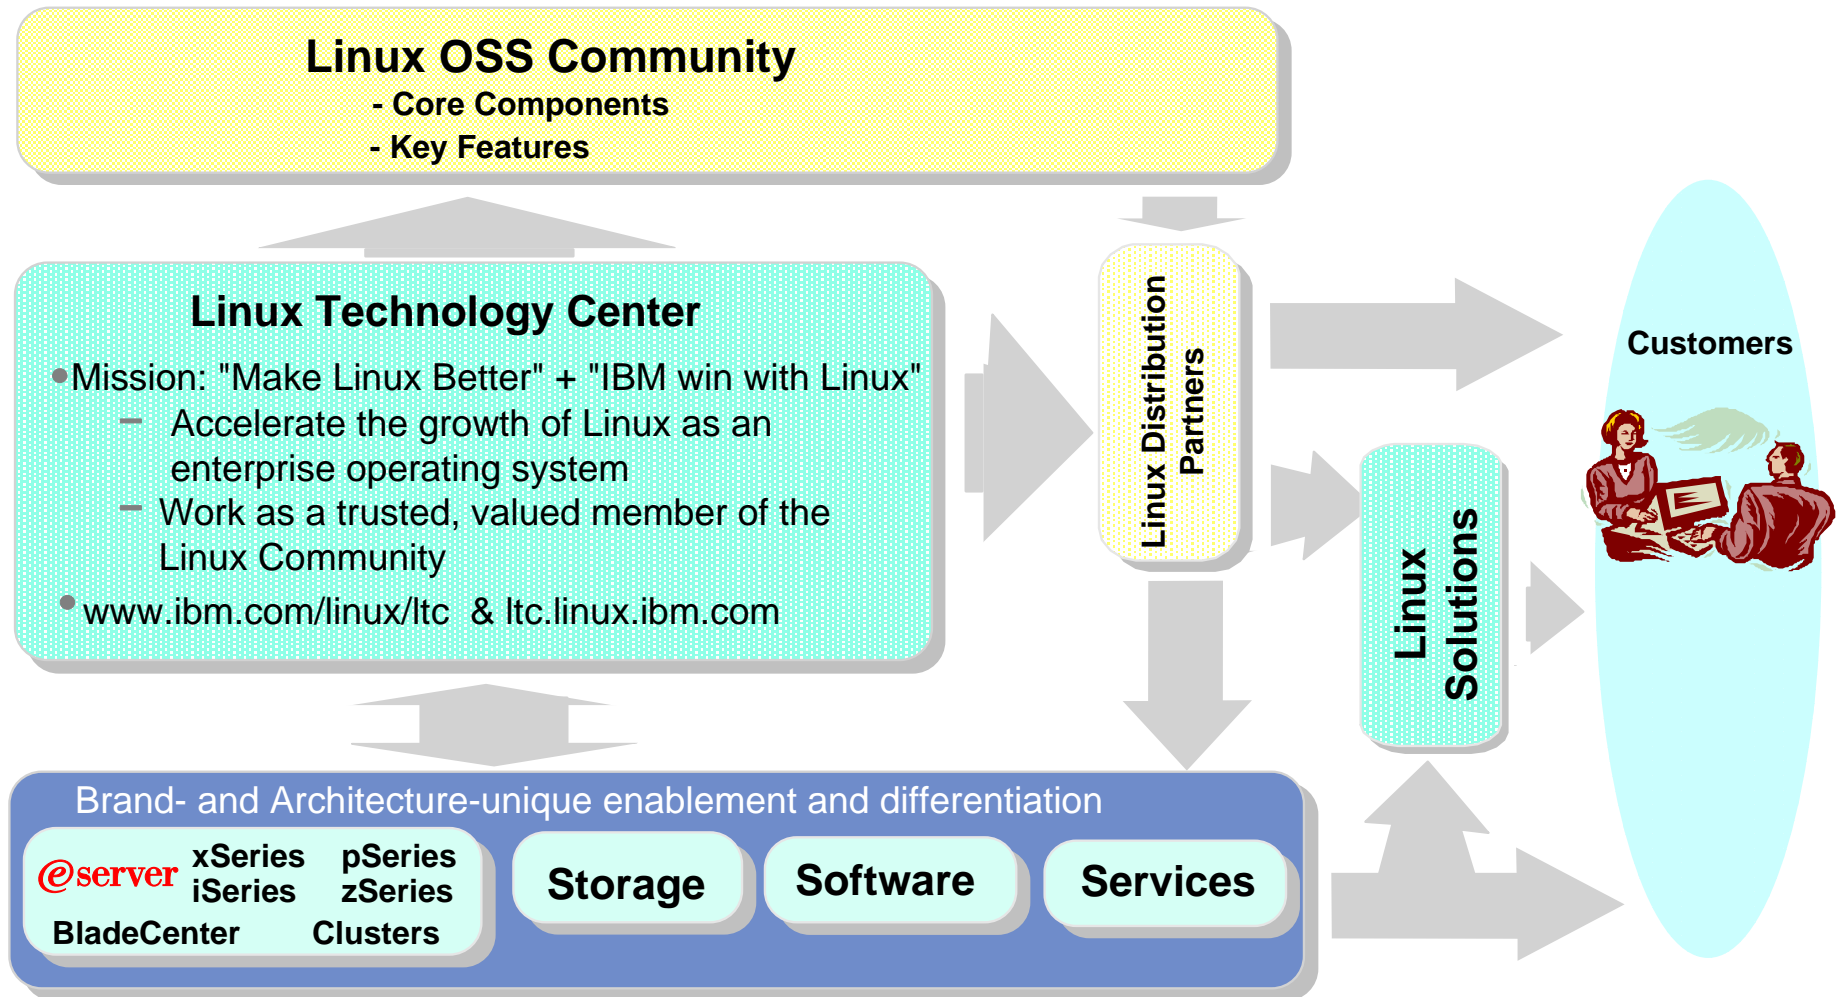




IBM Linux Technology Center

Linux Technology Center Focus Areas

Linux Technology Center



LTC Community Participation

- 450+ engineers directly contributing to open source components/projects
- The LTC is a key participant in community organizations:
 - ▶ Open Source Development Lab (OSDL) corporate sponsor.
 - OSDL provides hardware access and expertise to community open source projects.
 - ▶ Free Standards Group (FSG) driving Linux standards
- LTC provides support to other community groups / conferences:
 - ▶ OSDN (Open Source Developer Network)
 - ▶ 2.5 Kernel Developers' Summit
 - ▶ Linux World Expo
 - ▶ Linux Summit
 - ▶ O'Reilly open source conference
 - ▶ USENIX
 - ▶ GNOME
 - ▶ KDE
 - ▶ Free Software Foundation
 - ▶ Linux International

LTC Focus Areas

- Scalability
- Networking
- Security
- Storage I/O
- Networking I/O
- Volume Management
- JFS
- Serviceability
- File/Print
- Globalization
- NUMA
- Directory
- HPC Clusters & HA
- Systems Management
- Standards
- Test
- Performance
- Level 3 Support
- Accessibility
- Documentation
- Compilers
- Libraries
- Virtualization

Key Enhancements in 2.6 Kernel

- Scalability
 - ▶ 8 & 16-way SMP
 - ▶ per-CPU scheduler
 - ▶ per-CPU timers, counters, statistics
 - ▶ Improved resource locking
- VMM Enhancements
 - ▶ Support for 32 GB of memory on IA32
 - ▶ Large page (4 MB/2 MB) VM support
 - ▶ RMAP - reverse physical to virtual address mapping
 - ▶ Large page in-memory filesystem support
- Efficient support for large number of processes/threads
- I/O Enhancements
 - ▶ Large Block Raw I/O
 - ▶ Asynchronous I/O
 - ▶ Vectored block/raw I/O
- Hot Plug CPU, I/O
- Read-Copy-Update locking technology
 - ▶ IP route cache + RCU
 - ▶ IPC locking using RCU
 - ▶ dcache locking using RCU
- Futexes
- Logical Volume Management:
 - ▶ Device Mapper and EVMS
- NW Protocols: IPv6, IPsec, SCTP
- NUMA topology infrastructure & perf enhancements

Kernel Work - 2004

- Dynamic memory add/remove
- SCSI Multi-Path I/O
- Enhanced event logging
- Online diagnostics
- Hardening device drivers
- Resource Scalability (# tasks, IPC, I/O Capacity, etc.)
- Support for > 4000 I/O spindles
- Linux Kernel Locking and Cache Awareness
- Kernel Exported User Level APIs
- NUMA API
- NUMA Aware Scheduler and Multipath I/O
- IPv6 and NFSv4 hardening



IBM Linux Technology Center

Power Linux Strategy and Technical Roadmap

POWER Linux Motivation – Why POWER?

- Efficient virtual environment:
 - ▶ Dynamic LPAR
 - ▶ Sub-processor partitioning
 - ▶ Virtual I/O
- Extremely scalable architecture
 - ▶ Very good SMP hardware scalability
 - ▶ Very good I/O bandwidth
- Effective RAS capabilities
 - ▶ CPU and memory predictive failure
 - ▶ Capacity Upgrade on Demand
 - ▶ Effective platform and I/O error discovery, handling, and reporting

POWER Linux Motivation – Why Linux?

- Non-proprietary operating system
 - ▶ No “lock-in” to single OS vendor
 - ▶ Service and support available from multiple sources
- Runs on variety of hardware
 - ▶ Can choose hardware independently of OS
 - ▶ Porting applications is trivial (just recompile)
 - Just deal with 32-bit/64-bit and endian-ness
- Common programming environment
 - ▶ All tools freely available
 - ▶ Programming interfaces are well established

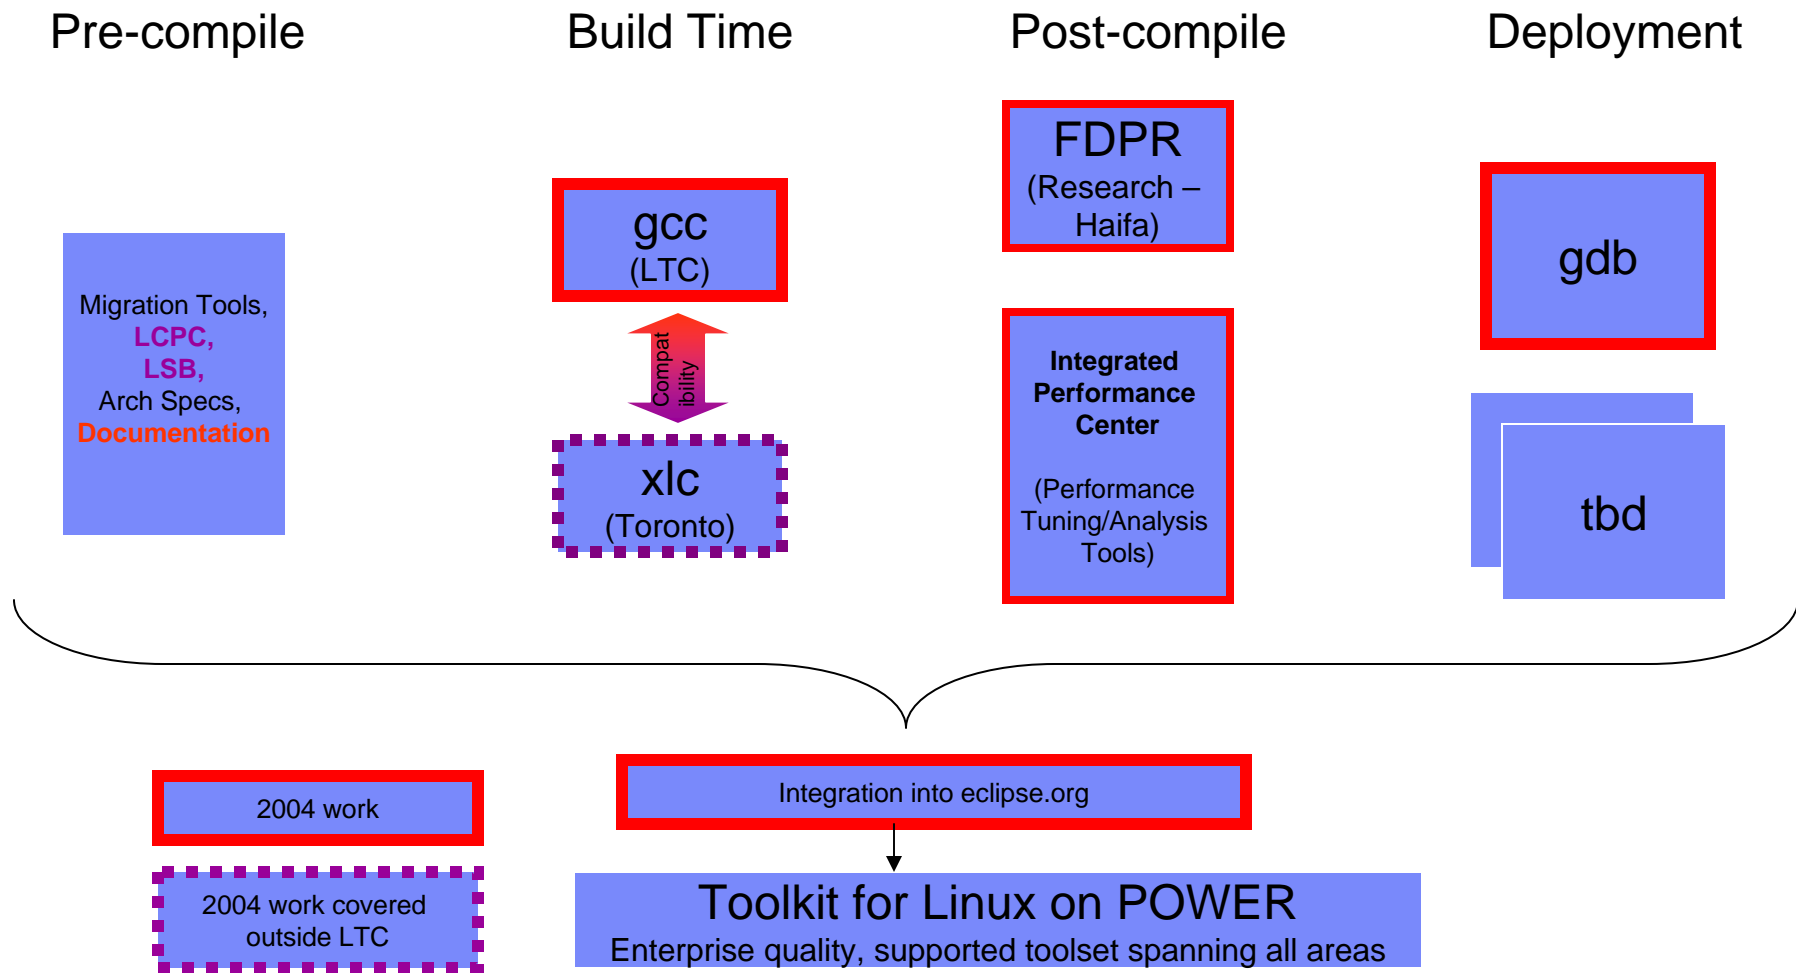
High Level Objectives

- Support all hardware systems
 - ▶ HV, blades, Squadrons
- Support broad set of I/O adapters
 - ▶ Initially, fully support a limited set of each type
 - ▶ Later, support a larger set of each type
- Fully support POWER virtualization capabilities
- Provide best possible price/performance at low end
- Provide best possible performance at high end
- Fully support POWER RAS capabilities
- Provide full hardware and software serviceability

Linux on POWER Roadmap – HW/SW Function

	2004		2005		2006
	1H	2H	1H	2H	
Distros					
2.4 kernel	•RHEL AS 3 U2 (TBD)		•RHEL AS 3 U4 (TBD)		
	•RHEL AS 3 U3 (TBD)				
2.6 kernel	•UL 2.0				
	•SLES 9	•SLES 9 SP1	•SLES 9 SP2	•SLES 9 SP3	•SLES 10
	•RHEL AS 4 (TBD)		•RHEL AS 4 U1 (TBD)		
Linux on Power	<p>Hardware</p> <ul style="list-style-type: none"> p655+, p690+ JS20 Blades GA1 Dual GPUL (PPC 970) <p>Features</p> <ul style="list-style-type: none"> EAL3 certification GR Processor Support SMT support 16 way SMP tuning Squadrons HW support DLPAR (CPU & IO) Virtual SCSI Virtual Ethernet Hot Plug (PCI and Disk) Subprocessor DLPAR Dynamic CUoD CPU Guard CPU Sparing Converged i/p LDP packaging 64-bit Altivec VMX Squadrons benchmarks JS20 blade benchmarks <p>Serviceability Enhancements</p> <ul style="list-style-type: none"> I/O adapter microcode update Service Agent JS20 blade firmware update Online Concurrent Diagnostics <p>I/O</p> <ul style="list-style-type: none"> IDE, Broadcom Ethernet, Qlogic Fiber channel, USB, Myrinet Fiber channel boot 	<p>Hardware</p> <ul style="list-style-type: none"> JS20 970+ GA2 Squadrons SF2, L4, SF4 Squadrons ML8, ML12, ML16, H HV1 970+, HV2(Power 5), HV4 (Power 5) <p>Features</p> <ul style="list-style-type: none"> Integrated performance center Squadrons H benchmarks IBM Director agent HW Instrumentation Toronto Visual Age Compiler update (C, C++, Fortran) support for JS20 GCC compiler optimizations Library optimizations FDPR-Pro port Squadrons benchmarks JS20 benchmarks HV benchmarks <p>Service Enhancements</p> <ul style="list-style-type: none"> Fast Dump Support EEH Recovery Device driver support of error logging framework Device driver support of diagnostics CIM based framework Linux standalone diagnostics <p>I/O</p> <ul style="list-style-type: none"> LSI SCSI expansion (JS20 blade side car) 2-D graphics adapter 	<p>Hardware</p> <ul style="list-style-type: none"> EAL4 certification 32-way SMP tuning IBM Director port to Linux on POWER Large page size support Dynamic memory add GCC C optimized for POWER architecture GCC C++ optimized for POWER architecture Libraries tuned for POWER architecture GNU Fortran compiler tuning for Power architecture NUMA tuning Squadrons benchmarks JS20 benchmarks HV benchmarks <p>I/O</p> <ul style="list-style-type: none"> Infiniband technology preview Vendor qualified multipath I/O 	<p>Hardware</p> <ul style="list-style-type: none"> Virtual I/O hosting partition (TBD) Multipath I/O Ethernet device scalability Storage device scalability NUMA APIs Squadrons benchmarks JS20 benchmarks HV benchmarks <p>I/O</p> <ul style="list-style-type: none"> 10Gb Ethernet 10Gb fiber channel 	<p>Hardware</p> <ul style="list-style-type: none"> 64-way SMP tuning Dynamic add/remove memory Compiler optimized for future architectures multiple page sizes large segment sizes JS20 benchmarks HV benchmarks <p>Service Enhancements</p> <ul style="list-style-type: none"> Memory Guard <p>I/O</p> <ul style="list-style-type: none"> 12x Infiniband

POWER Strategic Tooling Scope



2004 Toolchain work items

Line Item	Description
Gcc - Scheduler mods	Current scheduler designed for in-order execution, not GP-type cores
Gcc - Register alloc	Improve register allocator use of large PowerPC register set relative to few IA-32 registers.
Gcc - Loop optimization use of PPC CTR branch on count reg	SuSE rewrote GCC loop optimization with x86 requirements (for AMD) causing POWER branch on count reg optimization regression.
Gcc - C++ abstraction penalty	Stepanov benchmark shows perfect scaling for x86 but penalty for POWER.
Gcc - Compiler optimization heuristics	Ensure heuristics thresholds optimal for POWER (most developers have x86 or work performed under contract for various chip manufacturers).
Gcc - Autovectorization	Utilize AltiVec/VMX and advantage over MMX/SSE (requires early loop opt pass in Tree-SSA)
Gcc - Software pipelining	Optimize loops to pipeline through floating point unit with register alloc constraints.
Gcc - Address corner case failures for AltiVec/VMS	The AltiVec/VMX ABI for POWER is somewhat baroque and the current GCC functionality fails on some corner cases due to limitations in GCC's infrastructure for passing arguments to functions
Gcc release manager support	IBM financial support of Mark Mitchell, gcc release manager. Works at Code Sourcery
FDPR-Pro	Bring FDPR to POWER linux based environment, enhance global code reordering, aggressive function in lining followed by appropriate register allocation for eliminating spills and redundant code and by code reordering, static data reordering, other optimizations
Integrated Performance Center	Tool similar to Intel's VTune and Apple's CHUD.. Based on eclipse tool environment. Provide perf analysis and tuning perspective to Eclipse, HW analysis, SW app analysis, integrate other existing auxilliary tools

2004 Toolchain work items (cont'd)

Line Item	Description
Libm and compiler builtins optimization	Analyze existing math, memset, dyn mem alloc functions and replace them where needed with code tuned to platform
Gdb and other tools	TBD – investigation and plan for specific enhancements required, full staffing staged post 1Q04
GNU Fortran compiler	Join community effort to create skill base internal to LTC/IBM Definition of vendor-specific Fortran ABI
Performance Baselineing and Measurement	Setting baseline between x86 and ppc64, setting performance target using internal and gcc community benchmarks to measure progress towards goal
Documentation	Tech writer to provide doc support to improve gcc internals manual , doc on building cross compilers, docs/spec on ppc arch, etc
Test	Regression testing, compatibility testing between gcc / xlc New test case development Inclusion of some xlc test cases into gcc source base to detect regressions
C++ ABI test suite	License Code Sourcery's C++ ABI test suite. Finance code sourcery to port and maintain for ppc32/64
Project manager	Coordinating SOWs and external resource activity with team deliverables, working with Toronto on availability of overall power toolchain issues (gcc & IBM compilers), coordinating with IMD/Apple compiler work as well



IBM Linux Technology Center

Linux Technology Center Detailed Focus Areas

Ken Rozendal
February 3, 2004

© 2003 IBM Corporation

LTC's Overall Enterprise Focus Areas

- **Kernel Scalability**
 - ▶ POSIX asynchronous I/O
 - ▶ Network asynchronous I/O
 - ▶ Large block I/O support
 - ▶ direct I/O, raw I/O
 - ▶ locking granularity improvements
 - ▶ multi-queue scheduler
 - ▶ NUMA API and process binding
 - ▶ cache line, counter, stat optimizations
 - ▶ read/copy update lock primitive
- POSIX Threading
- I/O & Storage improvements
- SCSI, RAID, Fibre Channel driver hardening
- multi-path I/O
- persistent device naming
- large number of devices
- PCI Hot Plug
- CPU Hot Plug
- physical memory add/remove

LTC's Overall Enterprise Focus Areas (cont'd)

- **Enterprise Volume Mgmt System**
 - ▶ interoperability w/ other volume formats
 - ▶ drive linking, bad block relocation, snapshot (writeable), boot time volume discovery, 64-bit
- **Journalled File System**
 - ▶ in 2.4 and 2.6 kernels
- **Networking**
 - ▶ IPv4
 - ▶ IPv6
 - ▶ MobileIP
 - ▶ SCTP
 - ▶ TCP/IP Performance
- **Networking (cont'd)**
 - ▶ IPSec
 - ▶ Etherchannel
 - ▶ Channel bonding
 - ▶ GB, 10/100 Ethernet driver hardening

LTC Overall Enterprise Focus Areas (cont.)

- **Serviceability**
 - ▶ dynamic probes
 - ▶ POSIX event logging
 - ▶ crash dump facility
 - ▶ kernel debugger
 - ▶ online/concurrent diagnostics
- **File/Print enhancements**
 - ▶ Samba (Kerberos ticket enablement, AD printer publishing)
 - ▶ native CIFS client
 - ▶ Domain Controller (new)
 - ▶ OMNI Printer driver
 - ▶ NFS over TCP
- **Directory**
 - ▶ OpenLDAP evolution
 - ▶ Internationalization
 - ▶ Performance
 - ▶ Additional protocol support

LTC Overall Enterprise Focus Areas (cont.)

- **Security**
 - ▶ Kerberos
 - ▶ PKI
 - ▶ OpenSSL
 - ▶ Loadable Security Modules
 - ▶ XCrypto device support
 - ▶ PKCS-11 API
- **HPC Clusters**
 - ▶ OSCAR
 - ▶ System Installation Suite
- **Systems Management**
 - ▶ CIM infrastructure (WBEM)
 - ▶ Pegasus (C++) and SNIA (Java) CIMOMs
 - ▶ Linux schema implementation (SBLIM)
- **High Availability**
 - ▶ Linux-HA
 - ▶ Heartbeat

LTC Overall Enterprise Focus Areas (cont.)

- **Standards**
 - ▶ Linux Standard Base
 - ▶ Linux Internationalization (Li18nux)
 - ▶ IETF
 - Networking, IO, Directory, etc.
 - ▶ The Open Group
- **Linux Test Project**
 - ▶ testing each kernel release candidate for maintainer
 - ▶ test suite and test harness
- **Performance Benchmarking/Analysis**
 - ▶ TPC-H
 - ▶ TPC-W
 - ▶ VolanoMark
 - ▶ SpecWeb99
 - ▶ Netperf
 - ▶ dbench
 - ▶ Imbench
 - ▶ SPEC SDET

LTC Enterprise Focus Areas (cont.)

- Performance Instrumentation/Tools
 - ▶ lockmeter
 - ▶ kernprof
 - ▶ sstat
 - ▶ SGI co-pilot analysis tool
- Level 3 Support team
- Linux Documentation Project
- Accessibility
- Internationalization
 - ▶ Unicode
 - ▶ enablement of xterm, awk, grep, sed, regexp, iconv converters, bash and text utils
 - ▶ Li18nux Specification and Certification Test suite
- Graphics Workstation support
 - ▶ DRI, XFree86, Mesa (OpenGL)
 - ▶ 3D Graphics Adapter enablement

Key Maintainers/Core Members

- Free Standards Group (LSB, Li18nux)
- Samba
- OpenLDAP
- IPv6
- SCTP
- Various device drivers
- EVMS
- JFS
- SBLIM, Pegasus
- LSM, Bastille, xCrypto
- PCI hot plug
- USB
- APM
- OMNI Print
- PPC32, PPC64
- Linux-HA, Heartbeat
- Linux Test Project
-and growing