



Linux and Grid Disruptive and Darwinian

Objective of this Listen and Talk

Linux

What is Linux

Why Linux

Grid / Open Services

What is GRID

Why GRID



What is Linux

- **UNIX-like operating system developed by Linus Torvalds**
 - Version 2.4 available
- **Developed / tested by the Open Source community**
 - Highly disciplined / structured
 - High quality
 - Secure
 - Stable
- **Packaged and shipped by distributors**
 - Red Hat
 - United Linux
 - SuSE
 - Caldera/SCO
 - Conectiva
 - Turbolinux
 - Debian
 - Other regional distributors (Red Flag, Mandrake, etc...)




Open Source Software Advantage


- **Open-Source software**
 - ▶ Software-- with source code -- available on the Internet for unrestricted distribution under an open-source license
 - ▶ Participatory community debugs, enhances, and maintains the source code
 - ▶ Strong technical leadership
 - ▶ High quality, secure code
- **Linux**
 - ▶ Linux distributions are an "open source" based operating system, distributed freely, with it's underlining kernel source code openly published enabling low cost of entry.
 - ▶ "Linux provides a common, stable scalable and inexpensive Unix [style] operating system across a multitude of platforms, hence lowering overall costs ... " (Giga, 1999)



- inner group who make decisions
- directly commit code
- team runs via consensus w/ veto
- shared vision required for progress/success
- conservative - quality first
- active participants
- design/coding
- respect for previous contributions
- sheer numbers have impact



Open Source Approach



Open Source Not Just Linux



Linux



Perl



GNU



Apache



OpenBSD



Python



Gimp



Mozilla



NetBSD



TCL



Ghostscript



PHP

Platforms

Languages

Tools

App Infra

Linux Verses Unix


“IDC interviewed 60 Oracle8i database enterprise and service provider customers around the world about the savings realized From deploying Oracle8i on both operating platforms. From the results, IDC was able to determine the business value return of deploying Oracle 8i on Linux and Unix.

The savings realized was about half the cost of a Unix deployment.”

IDC



We even use it!



IBM phase one email Migration.

We have moved a sample of our Notes servers to Linux with a annual estimated savings of 10 million dollars.

We use Linux as the OS for out email Anti-Virus checking for 1 TB of email monthly.

www.ibm.com/linux is on Linux servers with over 17 million hits per day.

Linux is the OS in our manufacturing equipment company control 2.5 billion dollar technology manufacturing Line.

Linux is not a 3rd world build it is Global

It would cost **over \$1 billion to develop this Linux distribution by conventional proprietary means** in the U.S. (in year 2000 U.S. dollars).

It includes **over 30 million physical source lines of code (SLOC)**.


It would have required about **8,000 person-years** of development time, as determined using the widely-used basic COCOMO model.

Red Hat Linux 7.1 represents over a **60% increase in size, effort, and traditional development costs** over Red Hat Linux 6.2 (which was released about one year earlier).

It is estimated that the value volunteer labor in R&D would exceed **2 Billion dollars** a year verse **80 to 100 million** for a proprietary OS.

Linux Carries Every Datacenter Workload¹

"Are you planning to increase your Linux usage in the next two years?"



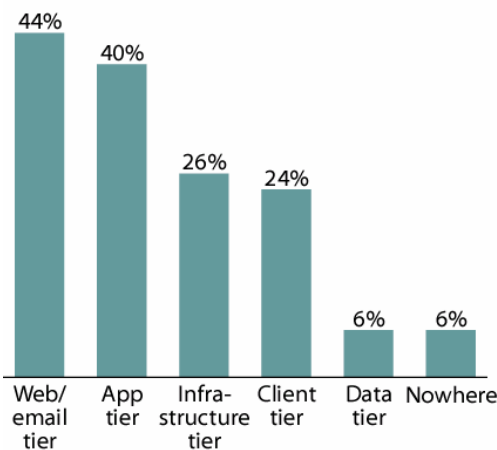
Base: 50 \$1 billion-plus companies

Source: OSDL



Linux Carries Every Datacenter Workload¹

"Where are you running Linux today?"



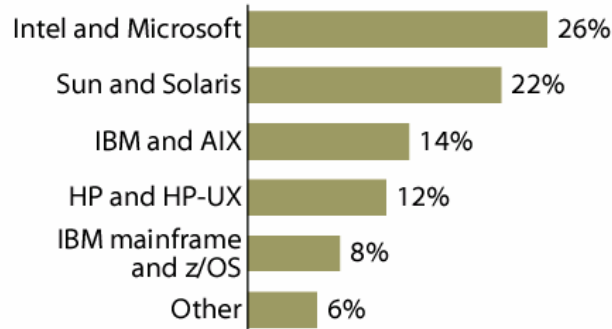
Base: 50 \$1 billion-plus companies (multiple responses accepted)

Source: OSDL



Every Operating System Is Losing Ground To Linux

“Which hardware and operating systems are you replacing with Linux?”

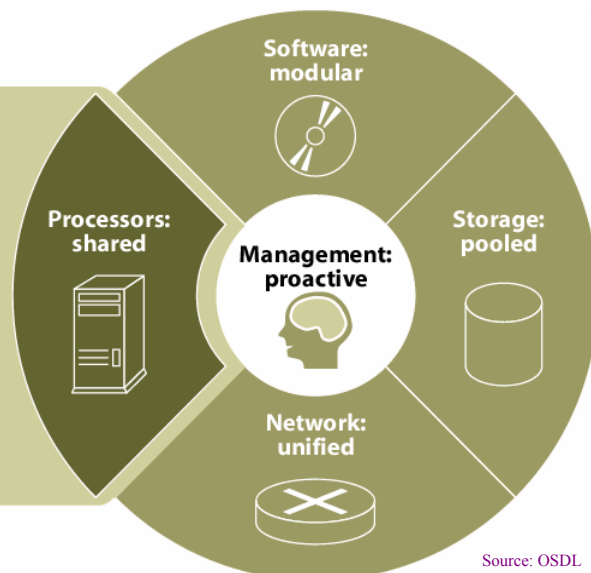


Base: 50 \$1 billion-plus companies
(multiple responses accepted)

Source: OSDL

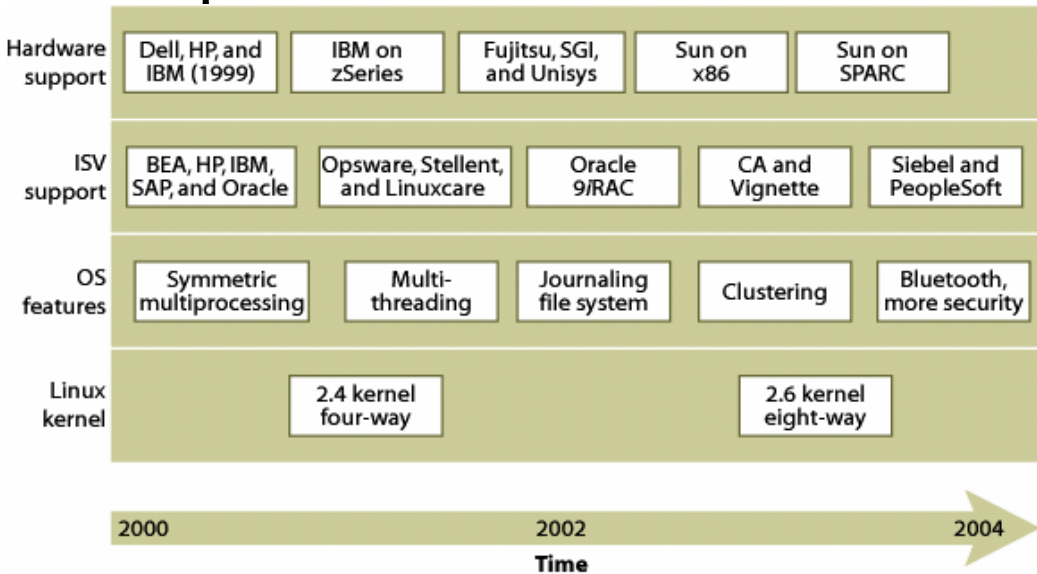
The Role Of Linux In Organic IT

- Linux is an abstraction that lets any application run on any processor.
- With managed deployment, firms can incrementally – and relentlessly – transform their servers into a shared pool of processors running Linux.
- An automated server pool is a solid step on the path to Organic IT.



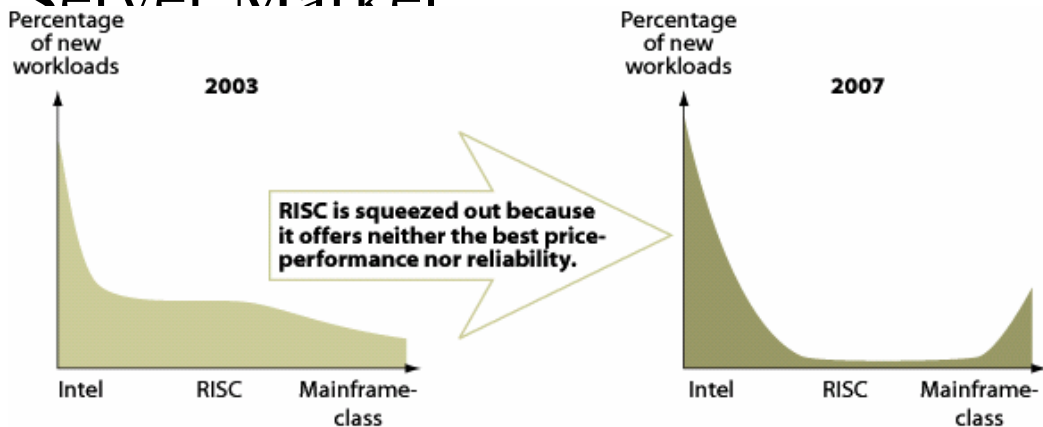
Source: OSDL

Linux Is Rapidly Gaining Enterprise Features



Source: OSDL

Linux Will Hollow Out The Server Market



Why do midrange RISC servers fade while both Intel and mainframe-class thrive? Because:

1. Linux makes it possible to run any application on any server -- servers will win based on their price and operating characteristics.
2. Intel's scale advantage guarantees that its processors will surpass RISC in price-performance.
3. Mainframe-class machines offer the highest grace under load and best business reliability.

Source: OSDL



Linux Distributors Combine their Power

UNITEDLINUX

UnitedLinux = Open Industry Consortium

■ Concept:

- ▶ Binary-compatible Linux distribution, branded "UnitedLinux"



■ Business Model: Add sw and services on top of UnitedLinux.

- ▶ Linux distributors maintain brand names
 - "UnitedLinux Inside"
- ▶ SuSE acts as UnitedLinux systems integrator



■ Implementation:

- ▶ Cross-platform support for:
 - all IBM @server platforms, key middleware, available YE 2001
- ▶ Worldwide language support
 - standards based (e.g. LSB 1.1, LI18NUX)



"IBM will continue to support Red Hat Linux across its key hardware, software and services offerings, we will also fully support UnitedLinux, which will make it easier than ever before to create a wide variety of Linux-based solutions for any size e-business,"

– Steve Mills, IBM senior vice president, SWG"

The Linux 80 20 rule

“In most datacenters Linux will support 80 percent of the systems. 20 percent will need to stay on current systems.”

**Michael Tiemann
CTO Redhat**



Consistency between Distributions LSB 1.3 (Linux Standards base)

Mission Statement

The goal of the LSB is to develop and promote a set of standards that will increase compatibility among Linux distributions and enable software applications to run on any compliant system. In addition, the LSB will help coordinate efforts to recruit software vendors to port and write products for Linux.

The word "OPEN" is written in a large, red, 3D, stylized font. The letters are thick and have a slight shadow, giving them a three-dimensional appearance. The font is bold and modern.



Binary Compatibility



- UNIX is a “source standard” requiring “porting” from/to AIX, HP-UX, and Solaris.
- LSB is a binary standard
 - Recompiling in only needed for differing architectures
 - Intel 32/64, PowerPC 32/64, S390 31/64
- <http://www.ibm.com/developerworks/linux/library/l-lsb.html>



Grid Computing

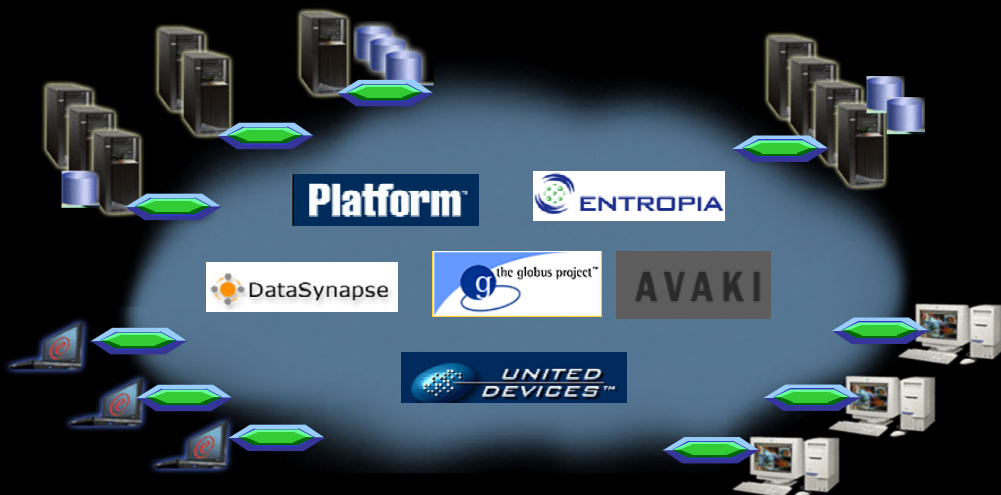
What is GRID?

Described by many as large scale distributed computing, Grid computing enables geographically dispersed computers or computing clusters to dynamically and virtually share applications, data, and computational resources.

More than just a vision, Grid Computing delivers measurable business value, particularly among organizations with compute or data-intensive demands. With Grid Computing, companies can balance the supply and demand of computing cycles and resources by providing users with a single, transparent, aggregated source of computing power. Ultimately, it provides the ability to lower the total cost of computing by providing on-demand, reliable and transparent access to available computer resources.



Grid Middleware Today



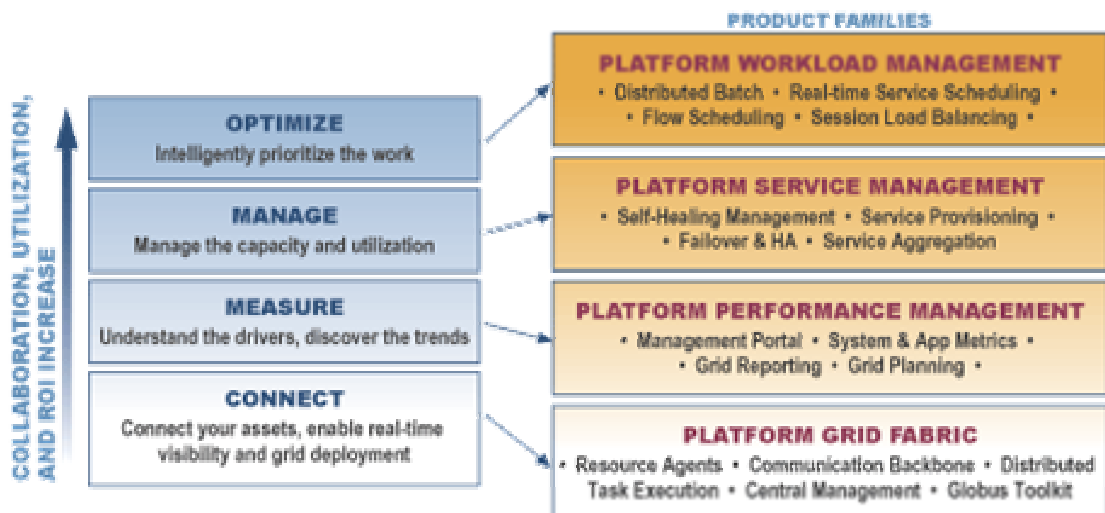
Grid Benefits

- ① Lower the cost of computing by optimization existing enterprise assets and computing resources.
- ② Improve efficiency and utilization, enabling users to solve complex computing projects such as "Grand Challenge" and "Big Science" problems.
- ③ Connect disparate, disconnected "processing silos" to improve collaboration for greater productivity.



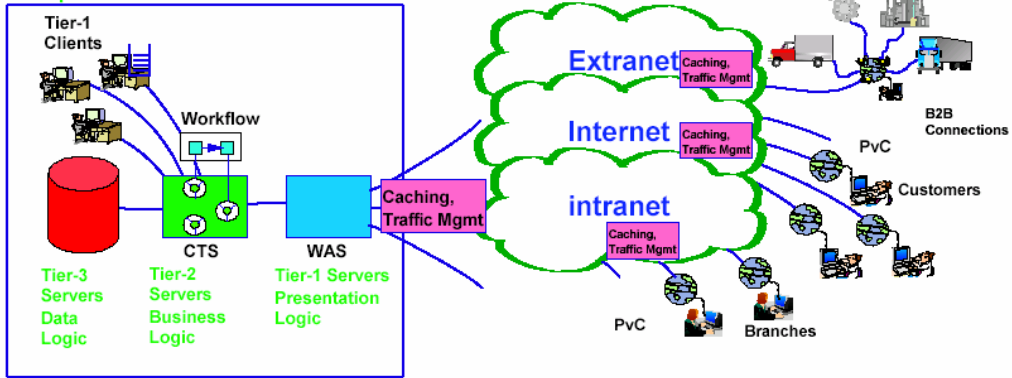
New Horizons

Enterprise Grid Computing solutions have been developed to address each of the four stages of Grid deployment, Platform enables organizations to maximize the potential of globally distributed computing resources.



De-composition of monolithic systems Re-integration of distributed compute resources

Enterprise Datacenter



This was originally a project to build a computer interface to fool a human that the interface was human.

The next generation of heterogeneous computing in a homogenous model (Grid) would require the intelligence of a Lizard.

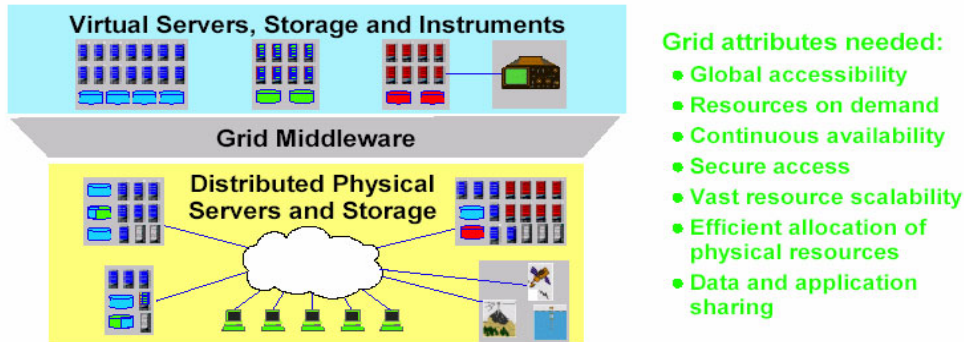


Trivia for today



Grid Computing

New open standards and software will enable resource sharing, collaboration and computing utilities on the Internet



Grid middleware will create virtual servers and storage from pools of servers, clients and storage dispersed throughout the Internet

Grid computing utilities will emerge and become like power utilities

Grid Suite Components

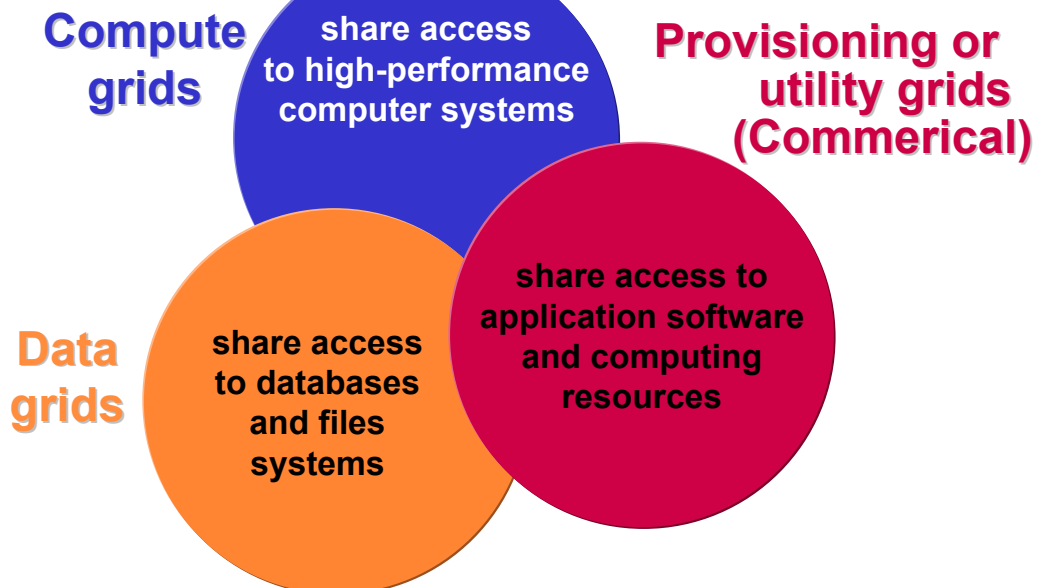
- **Integration of Platform technologies:**
 - Workload Management
 - Performance Management
 - Resource Management
- with best of breed technologies
- Integration of Globus services
- Leverages expertise and customer base
- Desktops to supercomputers

Drivers: Optimize Infrastructure

- **Resource optimization**
 - Maximize return on capital equipment by accessing spare cycles
- **Resource access**
 - Provide mechanism to share specialized resources across organizational boundaries
- **Cost sharing**
 - Allow multiple groups to contribute resources to a project while maintaining control of the resources
- **Improved management model**
 - Incorporate multiple systems in an organization under a single unified systems model



Grid types



Server / Storage Utilization

Peak-hour
Utilization

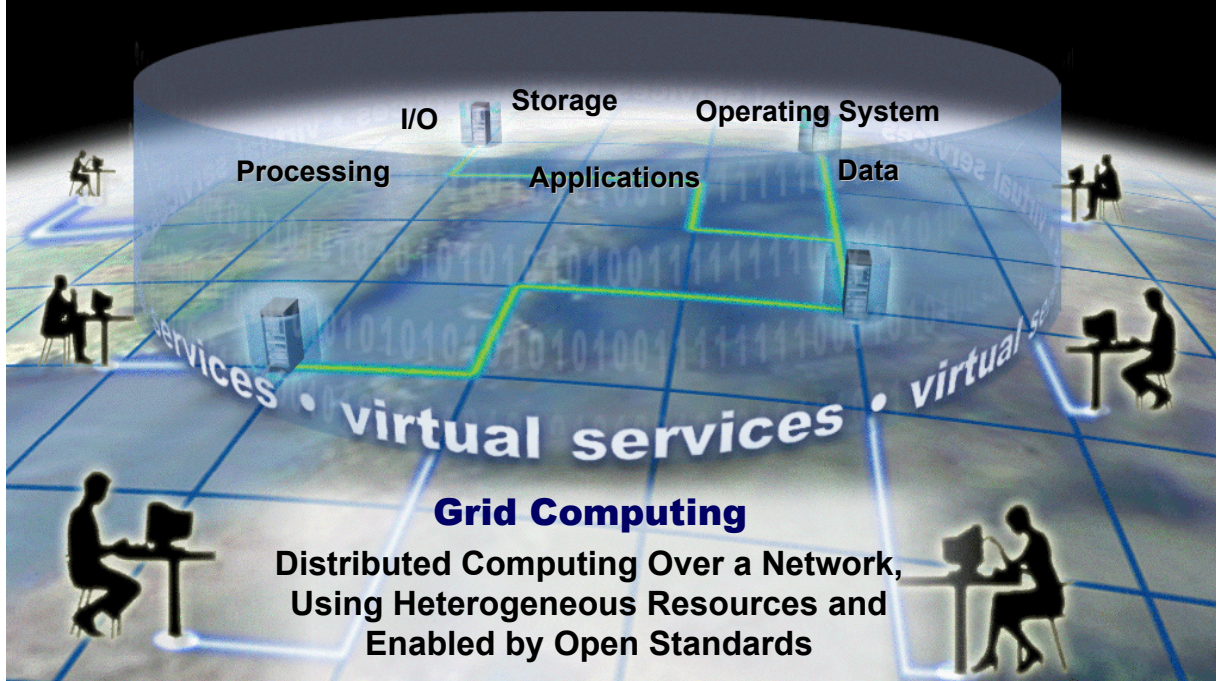
Prime-shift
Utilization

24-hour
Period
Utilization

	Peak-hour Utilization	Prime-shift Utilization	24-hour Period Utilization
Mainframes	85-100%	70%	60%
UNIX	50-70%	10-15%	<10%
Intel-based	30%	5-10%	2-5%
Storage	N/A	N/A	52%

Source: IBM Scorpion White Paper: Simplifying the Corporate IT Infrastructure, 2000

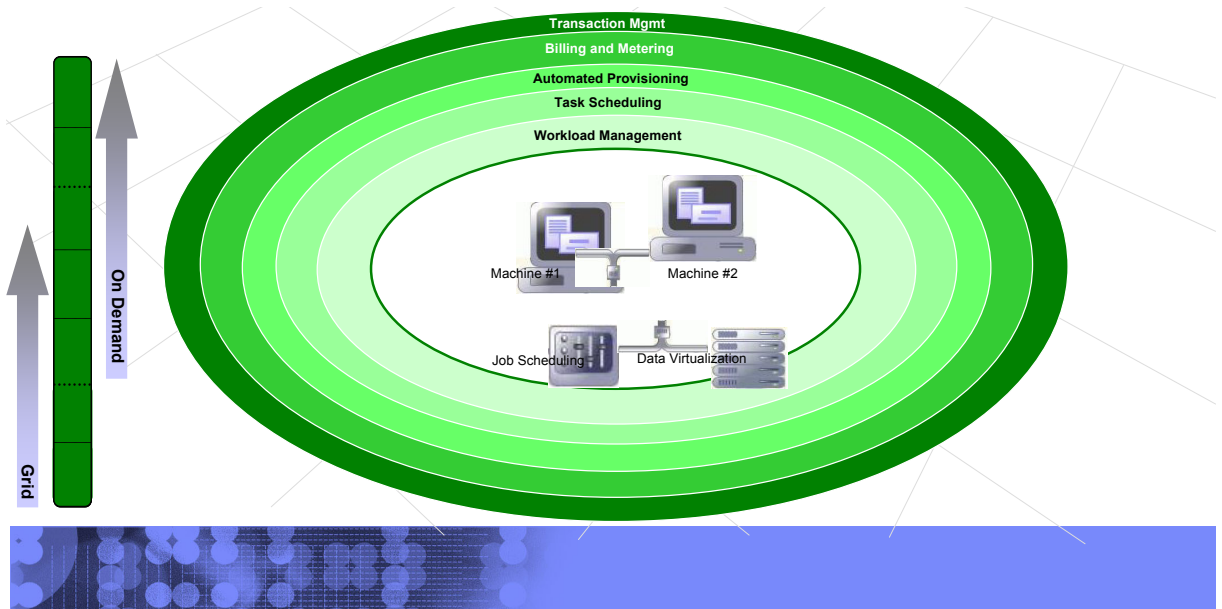
Virtualized



What Makes Up a Grid

Transaction Management:

- Manage the execution of e-business transactions across distributed resources
- Enable dynamic allocation of resources for transactional and parallel application models



Grid Addresses These Needs

- **Infrastructure Optimization**
 - Workload Management and Consolidation
 - Reduced Cycle Times
- **Increased Access to Data and Collaboration**
 - Federation of Data
 - Global Distribution
- **Resilient / Highly Available Infrastructure**
 - Business Continuity
 - Recovery and Failover



Charles Schwab

Business Analytics

- Reduced the processing time on an existing wealth management application, from more than four minutes to fifteen seconds
- Will allow Charles Schwab to increase customer satisfaction by responding to inquiries in real time...while the customer is on the phone
- Schwab is now planning to leverage Grid computing into other key business areas

"We believe that Grid computing ... has the potential to greatly improve our quality of service and be a truly disruptive technology."

Oren Leiman, Managing Director, Charles Schwab



charles SCHWAB

Analysts & Media Coverage: Jan. 27 Announcement

"Enterprises should see this announcement as an evolution of grid computing offerings, which are moving away from just custom engagements, to possible higher-volume packaged offerings." - **Gartner**

"IBM's recent grid announcement is a prime example of how to tightly telegraph business-process value propositions to individual vertical sectors." - **Summit Strategies**

"IBM, more so than any of its primary competitors, really has its act together in grid computing — and vendors such as Dell, Hewlett-Packard, and Sun should now be greatly concerned." - **Bloor Research**

"It's becoming more clear that IBM had something both specific and ambitious in mind when it talks about its vision of Grid Computing" - **Illuminata**




Smallpox Research Grid: Feb. 5

Announcement Summary

- Launched with IBM, United Devices, Accelrys and the U.S. Department of Defense.
- Focused on the development of new drugs that could potentially combat the smallpox virus post infection.
- 2 Million users PC's will provide 1.1 Petaflops or 1.1 Quadrillion operations per second.
- Over 1600 IBMers participating – 500 years of CPU time donated
- Powered by an IBM infrastructure, which includes IBM eserver™ p690™ systems and IBM's Shark Enterprise Storage Server running DB2® database software using AIX and Linux.
- Contains elements from recently announced the IBM Grid Offering for Analytics Acceleration in Life Sciences.



"While grids are currently popular in academic and research settings, IBM and others believe they will be useful in businesses as well. IBM last month released 10 grid products."



"Smallpox Researchers Seek Help From Millions of Computer Users"

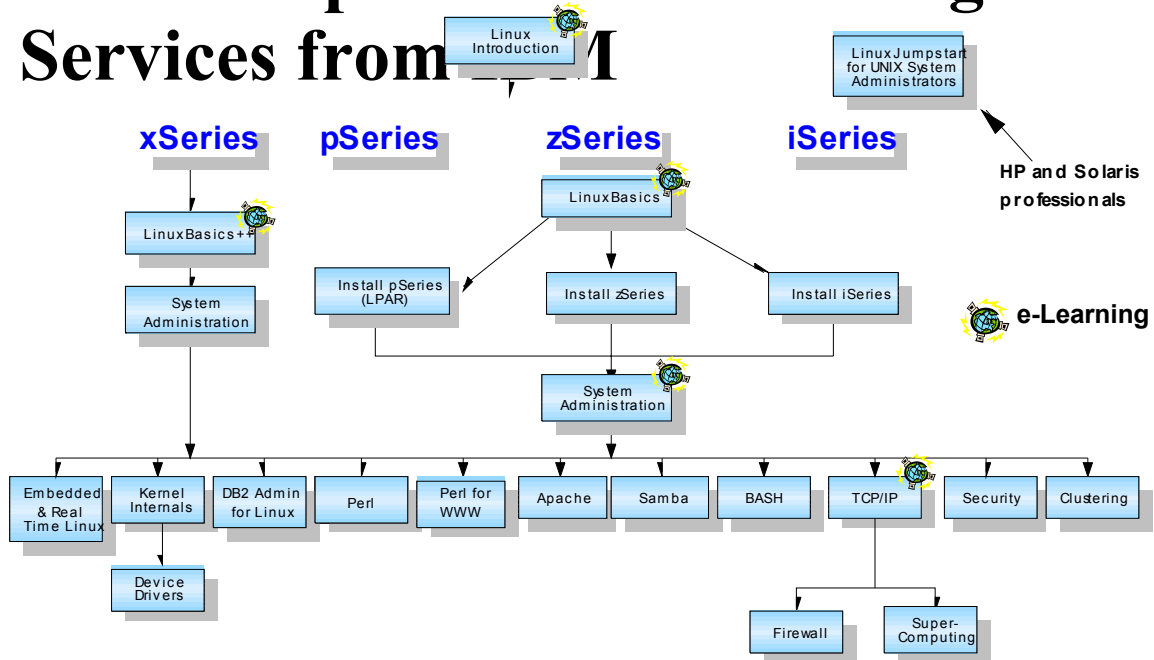
The New York Times

"... now companies are looking for commercial applications for the technology, such as creating a network for electronic games. In this case, IBM's database software will be part of the analysis process."

REUTERS  NEWS AND FINANCIAL INTELLIGENCE FROM THE WORLD LEADER

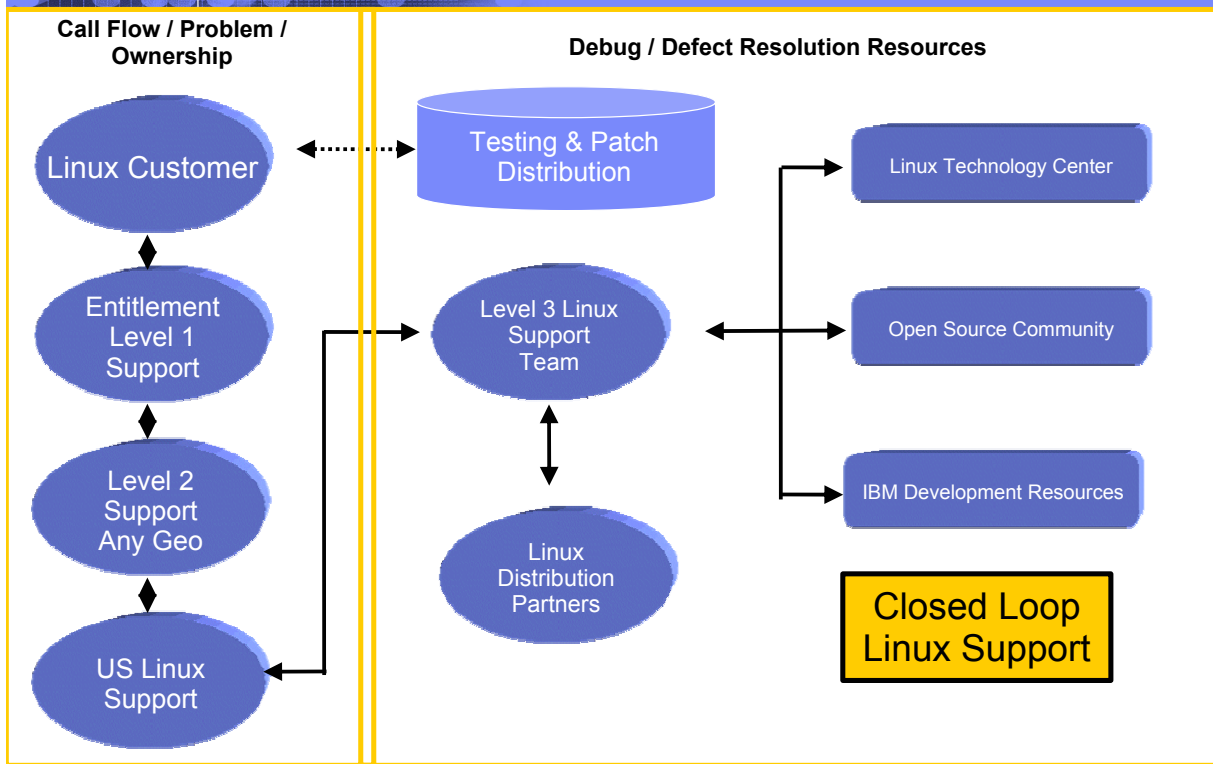
IBM CORNER

Linux / Open Source Learning Services from IBM



IBM Linux Operational Support Line Services

- How to/usage questions/ defects/ fixes/ patch
 - 800 number or electronic access
 - Defect support by IBM team, "down to the source code"
 - Response time: 2 hours Primeshift, 2-4 hours Offshift
 - Available on all IBM eServers, zSeries, xSeries, iSeries, pSeries
 - 9x5, 7x24
- Advanced Support
 - Customized Service Solution
 - Focused Support Team
 - Accelerated Response
 - Customized Reporting
- Proactive Recommendations/support
- Account Advocate (**Assigned Technical Specialist/Team**)
 - Track, review, and escalate problems as required




IBM Linux Technology Center


- **Mission: Accelerate maturation of standard, architecture-independent Linux into Enterprise**
 - Or "Make Linux Better!"
- **Extend value-net of System vendors, Linux companies, and existing Linux open source community to drive improvements back into Linux**
 - Drive Standards in the Linux space
 - Team of 250+ - small part of overall community
 - Open source - working within the shared vision of the community
 - World-wide virtual team - Adelaide, Austin, Bangalore, Beaverton, Boeblingen, Boston, Boulder, Canberra, Chicago, Denver, Haifa, Hawthorne, Hursley, Mount Laurel, Poughkeepsie, Raleigh, Rochester, San Francisco, San Jose, Seattle, Somers, Yamato, & Yorktown.
- <http://www.ibm.com/linux/ltc>


How the LTC Operates


- LTC is organized by teams focusing on various areas of Linux, much like any operating system development team.
- All development is done in open source and using open community methods (e.g. projects hosted in open forums, all contributors are welcomed)
- We participate in projects originated by IBM and others.
- Contributions are submitted to maintainers and Linux distributors.
- IBMers become maintainers and leaders in the community as they gain credibility as individuals.

LTC Accomplishments and Direction













 **Contributions Delivered**







 **Contributions Already Accepted**

 **More Work Underway**













 **Working with the community on the ecosystem**

OS Services

- | | |
|--|---|
|  Print |  System Management |
|  Internationalization |  Embedded |
|  User Libs |  Availability/Failover |
|  POSIX Threading |  Accessibility |
|  Interoperability |  Usability |
|  Cluster Management |  Network Security |

-  Test Verification
-  Defect Mgmt
-  Change Mgmt
-  Perf Eval
-  Standards
-  Documentation

Kernel

- | | |
|---|---|
|  Scalability |  Volume Management |
|  NUMA |  Availability |
|  Journalled Filesystem |  I/O & Storage |
|  Networking |  Scheduler |
|  Serviceability |  Security |
|  Problem Determination |  Reliability |