



**ORACLE®**

# From SuperScale to WallStreet

Nicolas Dube, Ph.D., Technical Director, HPC, XSO

# Oracle + Sun?

Leading Software Company

+ Leading Hardware Manufacturer

---

= #1 Integrated IT Solutions Vendor

# What is SuperScale?

- 2010: 100 TeraFlops +
  - Thousands of Sockets
  - High-Throughput Interconnect (IB...)
  - GBytes/s Filesystems
  - Massive Scaling, Jitter Minimization

# Sun's Top 500


- TACC: #9, 433 TFlops, 7872 nodes
- Sandia: #10, 424 TFlops, 5202 nodes
- Juelich: #13, 275 TFlops, 3288 nodes
- Kisti: #14, 275 TFlops, 3288 nodes
- ANU, BoM, CLUMEQ, Galileu, CHPC, TiTech, Leeds...

# Juelich's Juropa


- #1 3, 308 TFlops (peak)
- 27x SunBlade 6048 racks
- 6x M9 Infiniband switches
- 0.8 PB Lustre FS (18 GB/s)



3x






=



# Sandia's RedSky

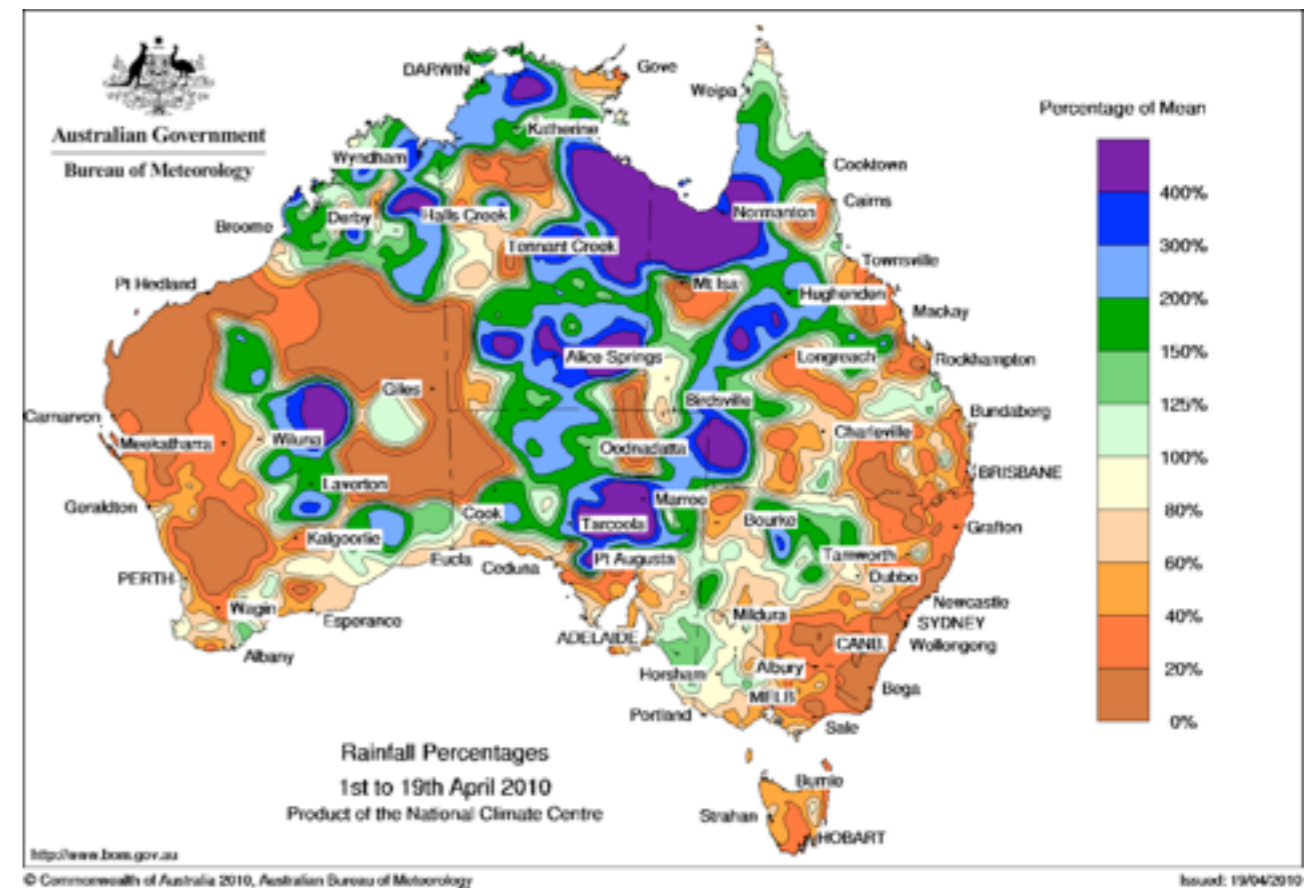
- #10, 487 Tflops (peak)
- 5202 nodes, 41 616 sockets
- Torus Network Topology
- Red-Black isolation



<http://www.youtube.com/watch?v=mNW9cYY4tqc>

# Bureau of Meteorology

- 54 TFlops, 576 nodes
- Full HA system
- Operational Weather
- 4 Hours Forecasts



# Project Galileu (Petrobras)

- 3 Sites: Sao Paulo, Rio de Janeiro, Maceio
- 126 Tflops total
- Torus Network Topology
- Ever built a DataCenter in the Rainforest?



# CLUMEQ's Colossus

- 86 TFlops
- One of a kind DC
- Boot over IB
- Dynamic Provisioning



<http://www.youtube.com/watch?v=2rwOyUP0sOM>

# Energy Efficiency

- Driving PUE to 1
- rPUE/ ERF < 1



CLUMEQ Silo PUE Data						
		summer	free-air	free-water	hours	energy (kWh)
Operating hours					8760	8760
Chilled water hours					720	720
Free air cooling hours					2160	2160
Free water cooling / re-use hours					5880	5880
	kW					
IT load	330.00	X	X	X	8760	2890800
Fans load	3.60	X	X	X	8760	31536
Chiller load	56.30	X			720	40536
Water tower	9.38	X			720	6756
Water pumps	3.75	X		X	6600	24772
Ultrasonic humidifying	1.75	X	X	X	240	420
Transformer loss	6.60	X	X	X	8760	57816
UPS loss (60 kVA)	7.20	X	X	X	8760	63072
Lighting (daily average)	1.50	X	X	X	8760	13140
<b>Total</b>	<b>420.09</b>					<b>3128848</b>
<b>Energy re-use</b>	<b>330.00</b>				<b>5880</b>	<b>1940400</b>
<b>Net Energy consumption</b>						<b>1188448</b>
Summer PUE						1.27
Free air-side PUE						1.06
Free water-side PUE						1.07
<b>PUE (year average)</b>						<b>1.08</b>
<b>rPUE / ERF (year average)</b>						<b>0.41</b>

# Why Infiniband?

- 40 Gbps bi-directional
- 1.6 $\mu$ s base latency + 0.2  $\mu$ s / switch hop
- Fully redundant (if required)
- Sun-Oracle solution:
  - 3x cable reduction
  - Leaders in IB integration + deployment
  - Torus has fewer switch hops

# Boot over IB (BoIB)

- Single fabric to the node
- Hierarchical boot servers
- How about cold-booting 960 nodes in less than 5 minutes? :-)
- First node shows up around 3:30
- Last nodes came in at 4:56

Initial Power on:	0:00
Bios bring up:	1:30-1:45
IB light-up from HCA:	1:45-2:15
Fetch Kernel:	2:00-2:30
Kernel Boot:	2:30-3:00
IB HCA re-init:	2:45-3:30
System image Fetch:	3:00-4:00
System Boot:	3:15-5:00

# Dynamic Provisioning

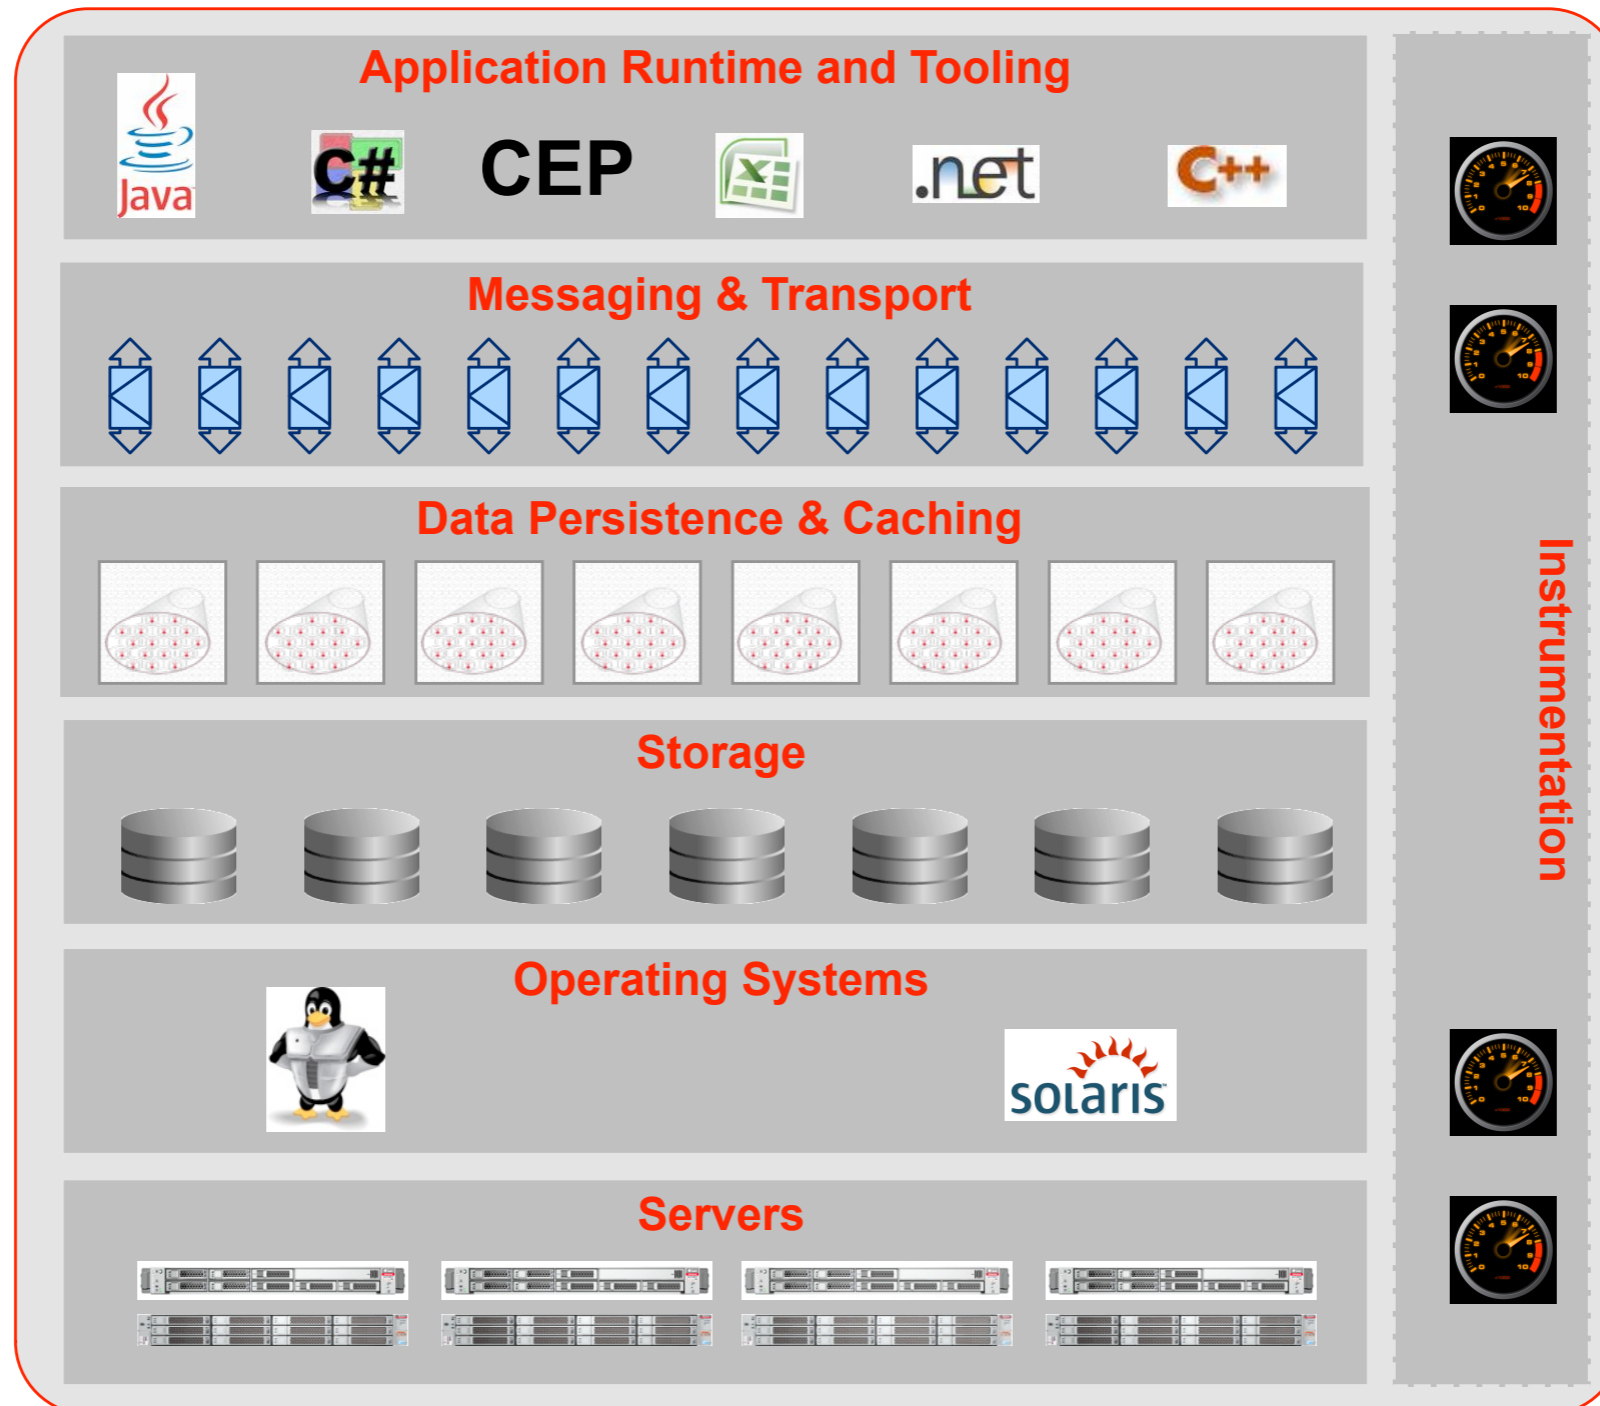
- Hardware “virtualization”
- Switching OS images “on-demand”
- SGE’s (scheduler) built-in feature
- This is the “cloud”.
- Real energy savings

# Lustre

## The World's Fastest and Most Scalable File System

- Parallel, scalable shared POSIX filesystem
  - Global name space up to 32 Petabytes of storage
  - Aggregate 240 GB/second IO for filesystem
  - 15,000 metadata operations/second
- High-performance heterogeneous networking
  - 50,000+ clients
  - 2GB/s+ IO for individual clients & servers
- High Availability
- Object-based architecture


# Low Latency Trading



# Software stack

- Sun HPC strategy Achilles' heel: CentOS
  - Support Plan ?
- Solution: Oracle Enterprise Linux
  - Oracle Unbreakable Linux






**ORACLE®**

**Nick Kamparosyan, Linux & Virtualization, Oracle**

***Realize True Value and Cost Savings with Oracle Unbreakable Linux***



The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.




# Oracle's Linux Leadership

- #1 Linux Database with 76% market share (Gartner, 2009)
  - More than 40% of Oracle Real Application Clusters deployments are on Linux
- More than 30% of Oracle E-Business Suite applications ship on Linux platform (Oracle internal)
- More than 50% of recent middleware shipments on Linux (Oracle internal)

# Oracle Unbreakable Linux

Enterprise-class support for the Linux operating system with premier backports, comprehensive management, indemnification, testing and more – all at significantly lower cost.



- 24x7 global coverage, 145 countries, 27 local languages
- Dedicated engineering and delivery team
- Backed by world's largest support team
- Enterprise-quality, Lower cost
- Service Excellence in Integrated Services (TSIA)
- 2008 Service Excellence in Mission Critical Support Award (SSPA)
- 2006 J.D. Powers and Associates Global Certification for Outstanding Customer Support

ORACLE



# Why Does Oracle Provide Linux Operating System Support?

- Customers demand enterprise-quality support
  - No pressure to update to the latest Update release
  - Timely bug fixes
  - Testing like traditional Unix vendor
- Linux is strategic to Oracle and its customers
  - We sell a lot of our software running on Linux
  - Linux provides an open standards alternative
- Oracle understands enterprise support
  - More than 30 years supporting enterprise customers
  - Efficient, low cost provider



# Oracle Unbreakable Linux

- Linux = open source software
- **Oracle Enterprise Linux = Software**
  - Tracks Red Hat Enterprise Linux product releases
  - Freely available source and binaries
  - Freely distributable binaries
- **Oracle Unbreakable Linux = Support Program**
  - For both Red Hat Enterprise Linux and Oracle Enterprise Linux
- x86 and x86-64 (AMD and Intel), Intel IA-64 architectures

**Making Linux Better**

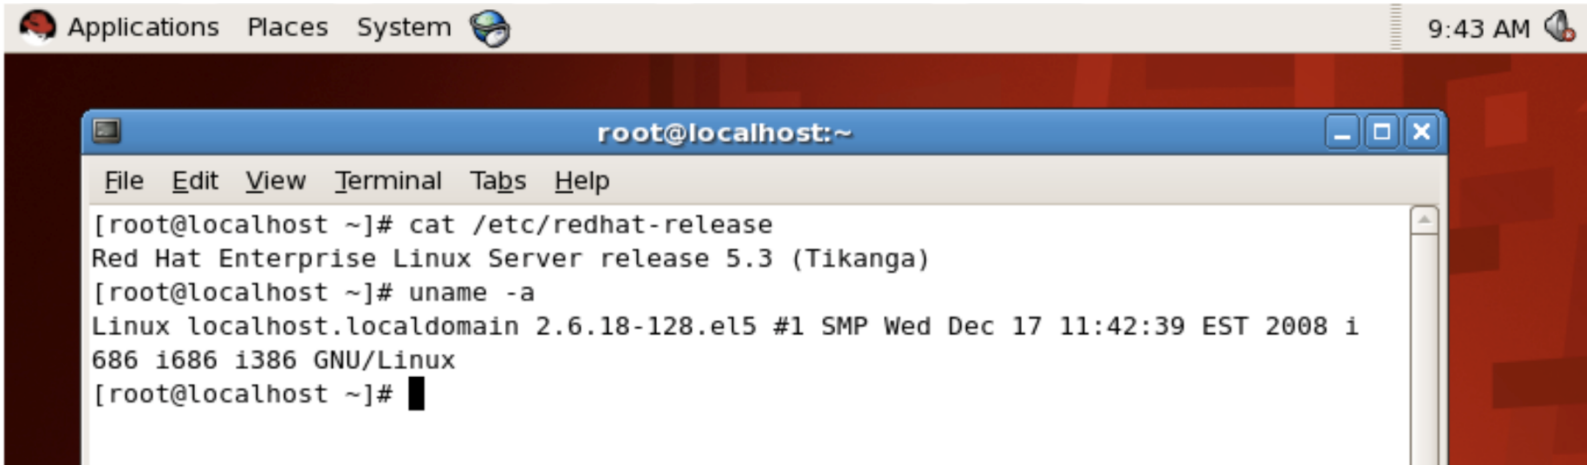


# How Does Oracle Maintain Compatibility

- Oracle maintains a repository containing the RHEL RPMs exactly as issued by Red Hat
- If any RHEL RPM is released by Red Hat, the corresponding OEL RPM is created from it **anew**
- Oracle re-releases every RPM released by Red Hat
- If OEL modifications exist for such an updated RPM, they are **reapplied** to the newly imported RPM
- **No branches of Red Hat code exist for OEL**: Oracle **does not** branch off from major RHEL releases
- Oracle **freezes** released OEL RPMs unless a custom backport is requested by a Premier level support customer

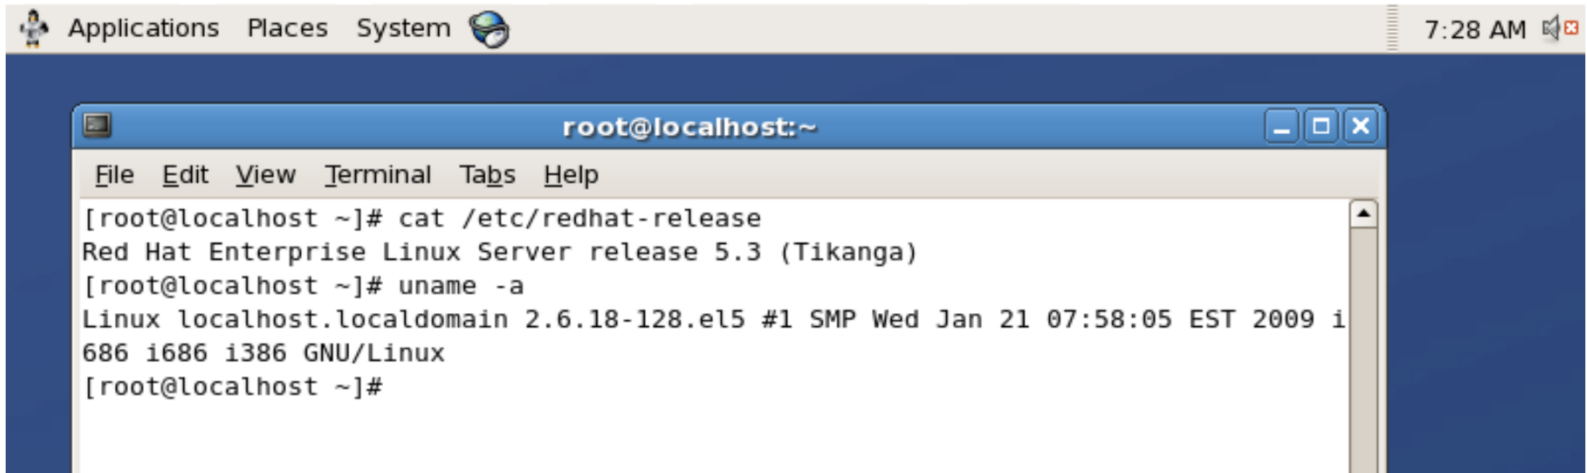
# OEL and RHEL Compared

## RHEL5 Update 3

A terminal window titled 'root@localhost:~' with a menu bar (File, Edit, View, Terminal, Tabs, Help) and window controls. The terminal shows the output of 'cat /etc/redhat-release' and 'uname -a'. The system is identified as Red Hat Enterprise Linux Server release 5.3 (Tikanga) on localhost.localdomain, dated Dec 17 11:42:39 EST 2008.

```
root@localhost:~  
File Edit View Terminal Tabs Help  
[root@localhost ~]# cat /etc/redhat-release  
Red Hat Enterprise Linux Server release 5.3 (Tikanga)  
[root@localhost ~]# uname -a  
Linux localhost.localdomain 2.6.18-128.el5 #1 SMP Wed Dec 17 11:42:39 EST 2008 i  
686 i686 i386 GNU/Linux  
[root@localhost ~]#
```

## OEL5 Update 3

A terminal window titled 'root@localhost:~' with a menu bar (File, Edit, View, Terminal, Tabs, Help) and window controls. The terminal shows the output of 'cat /etc/redhat-release' and 'uname -a'. The system is identified as Red Hat Enterprise Linux Server release 5.3 (Tikanga) on localhost.localdomain, dated Jan 21 07:58:05 EST 2009.

```
root@localhost:~  
File Edit View Terminal Tabs Help  
[root@localhost ~]# cat /etc/redhat-release  
Red Hat Enterprise Linux Server release 5.3 (Tikanga)  
[root@localhost ~]# uname -a  
Linux localhost.localdomain 2.6.18-128.el5 #1 SMP Wed Jan 21 07:58:05 EST 2009 i  
686 i686 i386 GNU/Linux  
[root@localhost ~]#
```



# OEL Compatibility

- Since October 2006:
  - No incompatibilities found by Oracle
  - No incompatibilities reported by Red Hat
  - No incompatibilities reported by ISVs
- No Red Hat Enterprise Linux installed at Oracle
- All development, testing and certification occurs on Oracle Enterprise Linux
- RHEL is certified for Oracle products on paper only
- Independent, third-party study from the Edison Group proves binary/source compatibility for OEL and RHEL  
<http://www.oracle.com/us/technologies/linux/ubl-edison-066204.pdf>



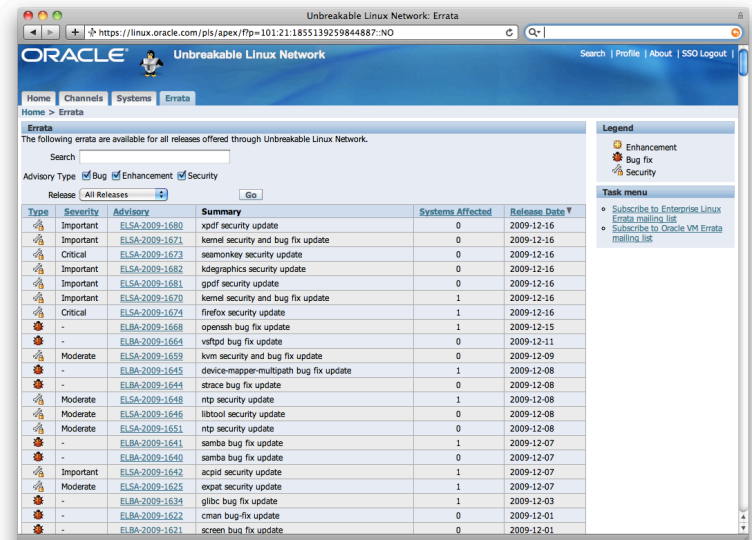


# Switching to Unbreakable Linux Support

- Existing Red Hat customers do not need to re-install Linux
- Switch to Oracle Linux support in just minutes
  - Purchase support
  - Download Oracle's up2date – [linux.oracle.com](http://linux.oracle.com)
  - Run `up2date --register` to create an account using your CSI for Linux support
  - System will be registered with the Unbreakable Linux Network. up2date will then download updates built by Oracle

# Software Channel: Unbreakable Linux Network (ULN)

- linux.oracle.com
- Updates, bug fixes, security fixes
- Access via up2date and Web browser
- Errata mailing list
- Search packages
- Download packages
  - Binary and source
- Manage channel subscriptions
- See how current your registered servers are





# Three Support Levels

- **Network Support includes:**
  - Access to updates and patches
- **Basic Support includes:**
  - Access to updates and patches
  - Log service requests via Web or phone, 24x7
  - Free Software Included:
    - Oracle Management Pack for Linux
    - Oracle Clusterware for Linux
    - Oracle Cluster File System (ocfs2)
- **Premier Support includes:**
  - Access to updates and patches
  - Log service requests via Web or phone, 24x7
  - Premier Backports of fixes to currently supported software versions
  - Lifetime support
  - Free Software Included:
    - Oracle Management Pack for Linux
    - Oracle Clusterware for Linux
    - Oracle Cluster File System (ocfs2)



**Equal to Red Hat's  
Best Support**



**Equal to Oracle's  
Database Support**

**Note: Only count physical sockets; no limit on cores or number of virtual guests**

ORACLE



# What is Premier Backporting?

- **Traditional Backporting** = A specific bug fix produced for the latest version of a package may be retroactively created and introduced as part of an earlier release or update level (e.g. a bug fix released in RHEL5 or OEL5 is also released as part of RHEL4 or OEL4)
- Only Oracle offers **Premier Backporting**, which goes far beyond traditional backporting. For example:
  - Customer runs RHEL5 or OEL5 Update 3 release and encounters a bug;
  - Although EL5 Update 4 is already released, the customer prefers not to upgrade;
  - With **Premier Backporting**, Oracle will provide a specific bug fix for the version the customer is running without forcing an upgrade;
  - In sharp contrast, a Red Hat support customer must wait to upgrade to the entire Update release to get just the one bug fix they need.
- Modeled after the way we support traditional Oracle products

**No pressure to upgrade to the latest Update release**

# Complete Management Included

- Oracle Management Pack for Linux  
Complete Linux server lifecycle management
  - Provisioning
  - Patching
  - Monitoring
  - Administration
  - Configuration Management
- **Free** for Basic and Premier support customers





# Oracle Clusterware for Unbreakable Linux

- Cluster software groups individual servers to cooperate as a single system:
  - Enables high availability and business continuity;
  - Helps ensure the protection of an application, Oracle or third-party;
- Included with Oracle Unbreakable Linux Basic and Premier support levels at no additional charge
- Clustering technology used in every RAC installation
- No need to pay for third party clustering software

[oracle.com/technology/software/products/database/clusterware](http://oracle.com/technology/software/products/database/clusterware)



# OCFS2

- General purpose, extent-based clustered file system
- POSIX compliant
- In Linux mainline since 2.6.16
- Built in cluster stack with distributed lock manager (DLM)
- Included in Oracle Enterprise Linux 4 and 5
- Also included in SLES
- Widely used

# Price Comparison

	Oracle	Red Hat
Installable Binaries	FREE	Fee-Based
Network Access (Patches and Updates)	\$119 (unlimited virtual guests)	\$349
Basic Limited (2 or less CPUs, sockets)	\$499 (unlimited virtual guests)	\$799 (4 virtual guests; limited hrs.)
Basic (More than 2 CPUs, sockets)	\$1,199 (unlimited virtual guests)	\$2,499
Premier Limited (2 or less CPUs, sockets)	\$1,399	No equivalent support
Premier (More than 2 CPUs, sockets)	\$2,299	No equivalent support
Management and Provisioning	Included (ULN, Oracle Management Pack for Linux)	\$10,143 ((\$9,951 Satellite Server + \$192 for 2 modules)
Cluster File System, Cluster Software	Included (OCFS2, Oracle Clusterware for Unbreakable Linux)	\$4,400 (GFS \$2,200 per node, min. 2 nodes for new customers)

ORACLE

# Linux Pricing Example 100 servers (2CPU)

	Oracle	Red Hat
Management	ULN + Included Oracle Management Pack for Linux License - \$0 Support - \$0	RHN License+Support - \$9401
Servers	Unbreakable Linux (Limited) Support - \$499	Standard (per 2 CPU) Support - \$799 RHN add-on \$192 Total : \$991
Clustering	ocfs2 and Oracle Clusterware license and support - \$0	gfs - \$2,200 per node
Total Support Cost (Annual)	\$49,900	\$319,100
Total Support Cost (3 Years)	<b>\$149,700</b>	<b>\$957,300</b>
<b>Save \$807,600 with Oracle Unbreakable Linux</b>		

ORACLE



# Linux Deployment at Oracle

- Oracle Enterprise Linux is the base development platform
  - More than **20,000** developers at Oracle use Oracle Enterprise Linux to develop key Oracle products – Database, Middleware, E-Business Suite
- Key internal systems run on Linux
- Oracle Global IT runs Oracle Enterprise Linux with more than **42,000** servers and supporting **4+ million** external users and **84,000** internal users
  - Oracle.com
  - Oracle Files Online
  - Oracle demo systems
  - Oracle financial systems
  - Oracle development organization

# Oracle On Demand Data Center Runs Oracle Enterprise Linux



- 3.6 Million users
- 10,000 servers
- 3.5 Petabytes storage
- 64-bit Linux



# Significant Kernel Contributions

- **Oracle Cluster File System 2** — The first default cluster file system in Linux (kernel 2.6.16 onwards)
- **Btrfs** — New file system featuring error detection, correction, and tolerance which will provide enhanced scalability and simplified management for large storage configurations
- **T10 DIF** — A first-of-its kind initiative for Linux that will bring enterprise-class data integrity (T10 Data Integrity Field standard) to the Linux platform.
- **Asynchronous I/O** — New interface reduces complexity
- **NFS over IPv6** — IPv6 networking support in the Linux NFS client and server
- More information at [oracle.com/technologies/linux/linux-tech-leadership-contributions.html](http://oracle.com/technologies/linux/linux-tech-leadership-contributions.html)



# Focus on Linux Testing

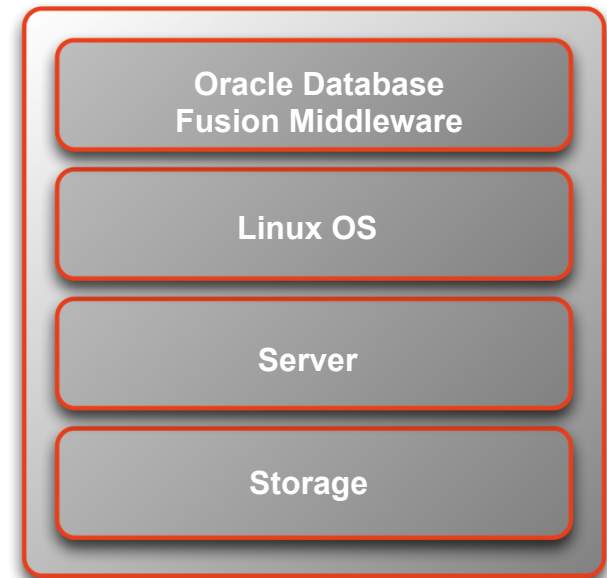
- Real-world regression and stress testing
- Customer-centric testing:
  - Test Linux features that matter to Oracle customers
  - Oracle and non-Oracle workloads (e.g. backup) running concurrently
  - Adverse conditions (low memory, low disk space, etc.)
  - Long, continuously running stress tests (detect memory leaks)
  - Check for performance regression and degradation

# Testing: Oracle Validated Configurations

Pre-tested, validated, and supported Linux architectures, including

- Software, hardware, storage, drivers, networking components
- Best practices for Linux deployment
- Real-world testing of complete stack
- More than 100 configurations published, freely available for download

[oracle.com/linux](https://oracle.com/linux)



**Oracle Validated Configurations offer faster Linux deployments while lowering infrastructure costs**

# More Than 4,500 Customers...and growing

Rely on Oracle Unbreakable Linux and Oracle VM



ORACLE



# Partner Endorsements

700+ Partners Endorse Oracle Unbreakable Linux

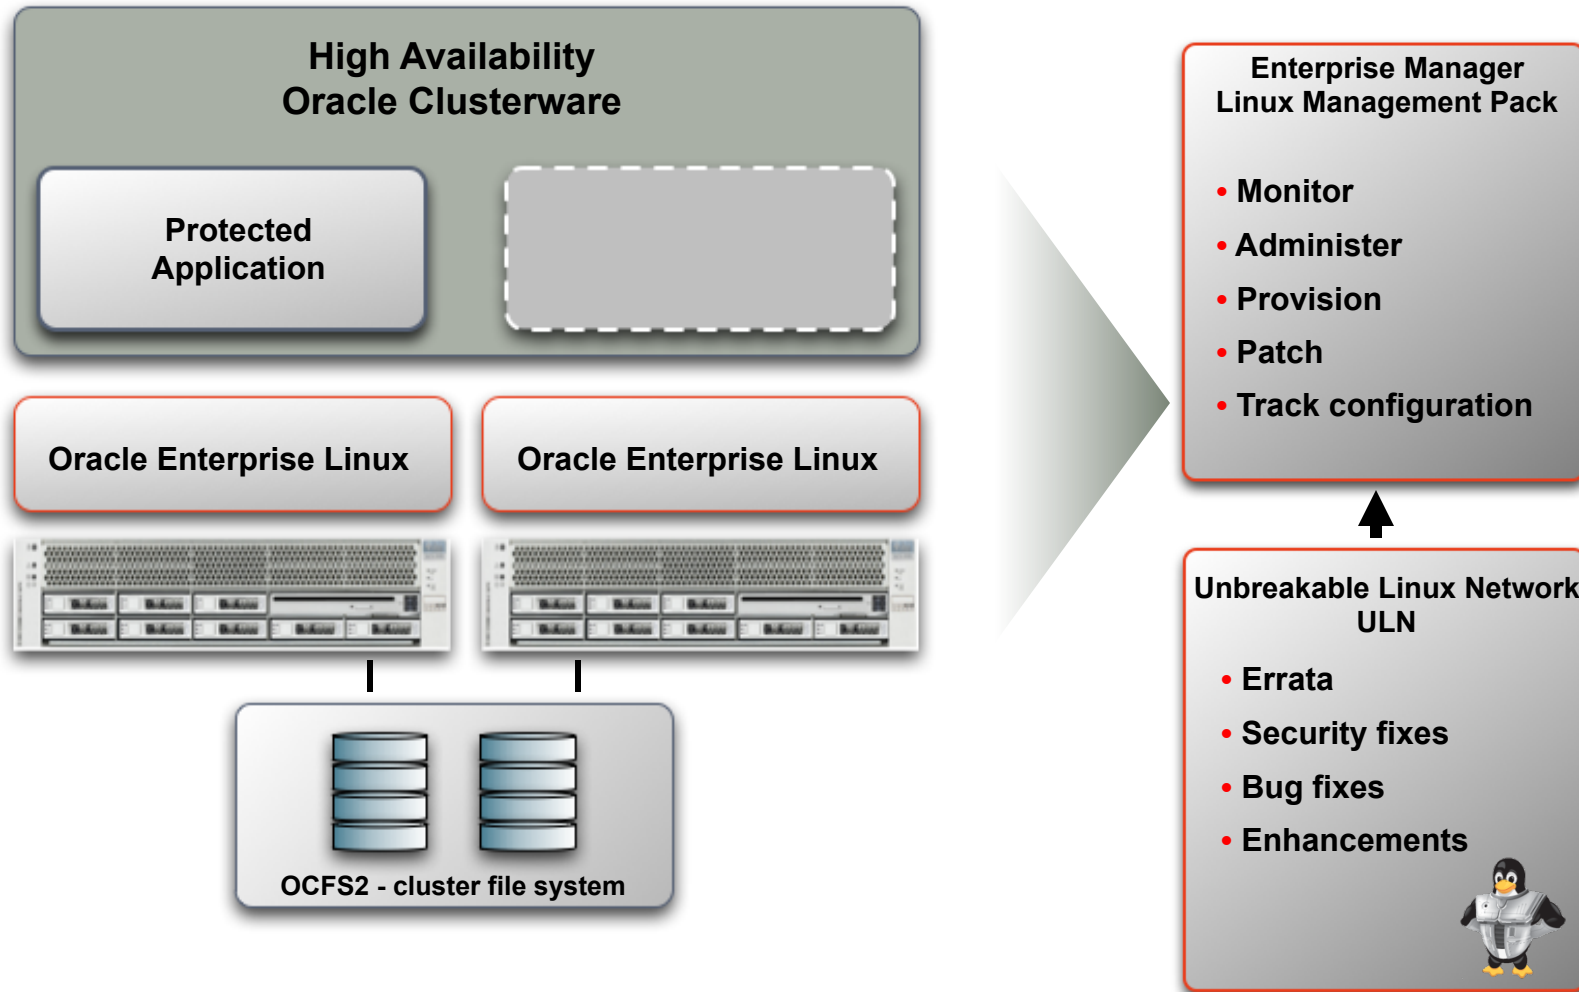




# Why Oracle for Linux Support?

- Industry leading enterprise software company with 30+ years of data center experience
- Offers a complete software stack that runs on Linux
- Linux leader; More than a decade of Linux experience; industry's first commercial DB on Linux in 1998
- More than 7,000 professionals offer 24x7, large scale, global support in 145 countries, in 27 languages
- Large-scale Linux testing with the largest and most complex data center workloads
- Significant and ongoing contributions to Linux community
- Free Enterprise Linux binaries, free download
- Free value added software for clustering and server management
- Focused on lowering Linux support cost for customers

# Complete Solution




ORACLE



# Cut Costs with Oracle Unbreakable Linux

- Free Enterprise Linux software; **no license cost**
- **Free** Oracle Management Pack for Linux software
- **Free** Oracle Clusterware for Oracle Unbreakable Linux Customers
- **Affordable** Support subscription; starts at just \$119/yr. per server; at least 50% off Red Hat list
- Large scale 24x7, global support; improved response times; **faster time to market**
- Consolidated support contact **eliminate finger pointing**
- Validated Configurations **speed Linux deployments**
- Consolidate support contracts; **lower TCO**

# Complete, Open, Integrated




## Customer Benefits:

- Open and standards-based
- Innovation delivered faster
- Better performance, reliability, security
- Shorter deployment times
- Easier to manage and upgrade
- Lower cost of ownership
- Reduced change management risk
- One-stop support
- Reduced down time

## Resources

- **Linux Home Page**  
[oracle.com/linux](http://oracle.com/linux)
- **Free Download:**  
**Oracle Enterprise Linux**  
[edelivery.oracle.com/linux](http://edelivery.oracle.com/linux)



**Unbreakable**  
**Linux**

ORACLE®

ORACLE®



# Thank you

- [nicolas.dube@oracle.com](mailto:nicolas.dube@oracle.com)
- [nick.kamparosyan@oracle.com](mailto:nick.kamparosyan@oracle.com)

**SOFTWARE.**  
**HARDWARE.**  
**COMPLETE.**

**ORACLE®**