
ORACLE®

ORACLE®

Overview: Oracle Linux

Lenz Grimmer
Senior Product Manager

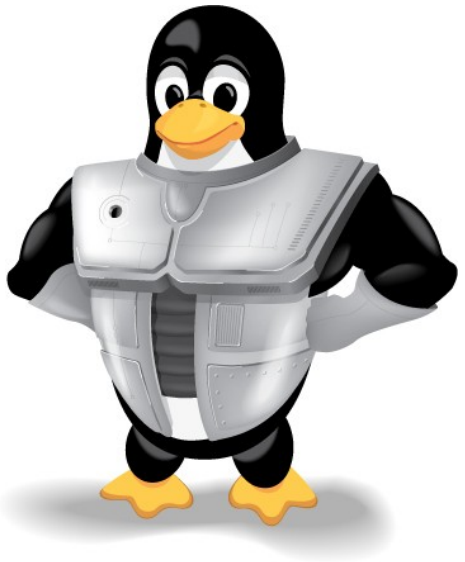

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.

The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

$ whoami

1998

2002

2008

2010

Program Agenda

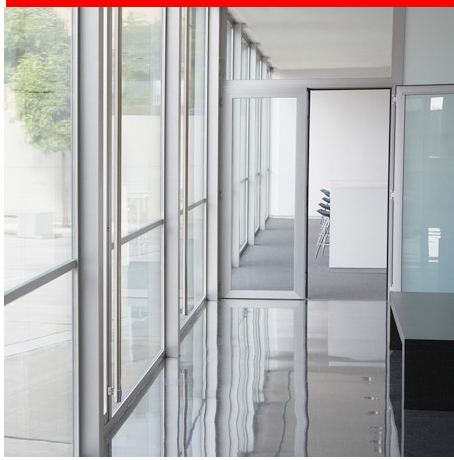

- Oracle's commitment to Linux
- Overview: Oracle Linux and the Unbreakable Enterprise Kernel
- Oracle Linux Support Program

Larry Ellison
CEO, Oracle
December 2010

"Solaris is clearly the number one Unix, and we're working very hard at making Oracle [Enterprise] Linux the number one Linux"

ORACLE®
LINUX

ORACLE®
SOLARIS

Oracle Linux – Overview

- Launched at Oracle Open World in 2006
- Compatible with Red Hat Enterprise Linux
- Freely available and distributable (source and binaries)
- Oracle offers Linux support for Oracle Linux and RHEL
- Oracle's base Linux development platform
- Customers can switch in minutes – no reinstall needed
- Applications run unchanged

Certification: Unbreakable Enterprise Kernel

- No need to (re-)certify unless your application has any kernel module dependencies
- Using the Unbreakable Enterprise Kernel changes only the box labeled “Kernel.”
- Installing a kernel does not change system libraries such as glibc
- The glibc version is 2.5 whether you choose to run the Unbreakable Enterprise Kernel or Red Hat compatible kernel in Oracle Linux

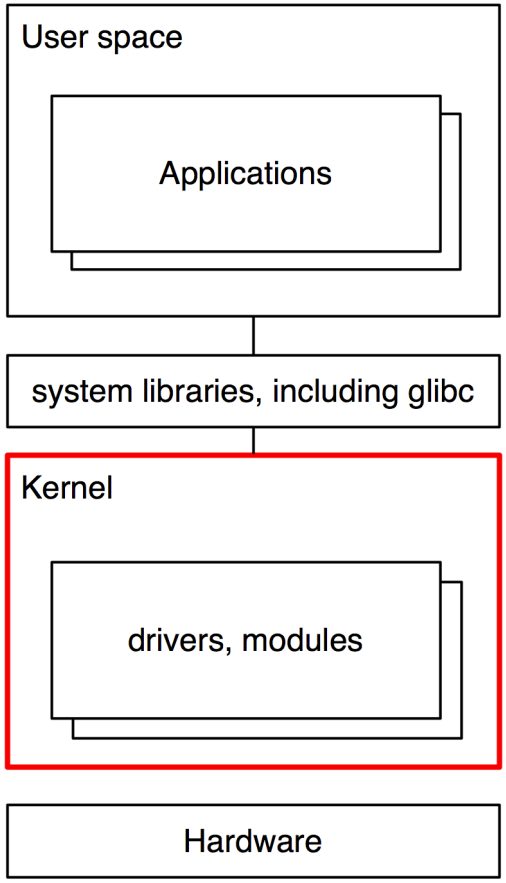

Oracle Linux Compatibility

- Since October 2006:
  - No incompatibilities reported by a customer
  - No incompatibilities reported by Red Hat
  - No incompatibilities reported by ISVs
- All development, testing and certification occurs on Oracle Linux
  - No Red Hat Enterprise Linux installed at Oracle

More Than 6,000 Customers Use Oracle Linux

The Unbreakable Enterprise Kernel

- Fast, modern, reliable
- Used by Exadata and Exalogic for extreme performance
- Allows Oracle to innovate without sacrificing compatibility
  - Oracle Linux now includes **both** the Unbreakable Enterprise Kernel **and** our existing Red Hat Compatible Kernel
  - You choose at *boot time*: a system optimized for running Oracle software **or** strict Red Hat compatibility.

Oracle now recommends only the Unbreakable Enterprise Kernel for all Oracle software on Linux

The Unbreakable Enterprise Kernel

- Large x86 systems (64bit)
- Improved power management
- Fine grained CPU and memory resource control
- Hardware Fault Management
- Built in data integrity
- Diagnostics Tools
- Upstream innovation tested and supported by Oracle
  - More than 80,000 test-hours every day

Unbreakable Enterprise Kernel Facts

- Based on the stable 2.6.32 kernel
  - Includes a number of enhancements already in 2.6.32 contributed by Oracle
  - Plus new optimizations from Oracle that are all open source
- Existing applications run unchanged
- Default on Oracle Linux 6 and 5.6
- Easy installation on top of Oracle Linux 5 or RHEL
- Free download via public yum server
- Covered as part of Oracle's Unbreakable Linux support program – No change in pricing

The Unbreakable Enterprise Kernel: Fast

Benchmark	RHEL 5 compatible kernel	Unbreakable Enterprise Kernel	Gain
8kb flash cache reads (IOPS)	197 thousand	1 million	**400%**
Solid State Disk access	4GB/second	9.5GB/second	**137%**
Infiniband RDS messages, single card (IOPS)	89 thousand	273 thousand	**200%**
8 socket database OLTP (transactions per minute)	**1.8 million**	**3.2 million**	**75%**

The Unbreakable Enterprise Kernel: Modern

- Bigger servers
  - Up to 4096 CPUs and 2 TB of memory
  - Up to 4 PB (petabyte) clustered volumes with OCFS2
  - Advanced NUMA support
- Power management
  - CPUs to stay in low power state when the system is idle
  - ACPI 4.0
- Fine grained CPU and memory resource control

The Unbreakable Enterprise Kernel tracks mainline Linux – users get community and Oracle enhancements faster

The Unbreakable Enterprise Kernel: Reliable

- Eliminates silent data corruption using Data Integrity – stops corrupt data from being written
- Reduces system crashes and improves application uptime via Hardware Fault Management
- Improved Diagnostics Tools

Improved InfiniBand and RDS Performance

- OFED stack updated to 1.5.1
- Reduced lock contentions
- Spread interrupts over CPUs
- RDS

Task Control Groups

- Fine grained control over CPU, memory
- Subset the resources of a larger system
- Limit CPU and memory available an application or group of applications
- Control access to devices
- Works inside virtual guests

Improved Power Management

- Tickless kernel
  - Timer interrupts are performed on demand rather than at a predetermined frequency
  - Enables CPUs to stay in low power state when the system is idle
  - Reduced overall power consumption
- ACPI 4.0
- Powertop
  - Reduce server power usage by identifying power hungy processes

OCFS2 1.6

- Reflink
  - Writeable snapshots
  - Unlimited snapshots of snapshots
- User space cluster stack support
- JBD2 support
- POSIX ACL support
- Quota support
- Extended/Security attributes
- Metadata checksums
- Indexed directories

Data Integrity

ORACLE LINUX

EMULEX

- Data Integrity Field
  - Protects path between HBA and storage device
- Data Integrity eXtensions
  - Protects path between application and HBA
- **Detect** in-flight data corruption
- **Prevent** corrupt data from being written
- Works with DIF/DIX aware Host Bus Adapter
- Data integrity-enabled ASM kernel driver will protect against data corruption from application to disk platter

Data Integrity

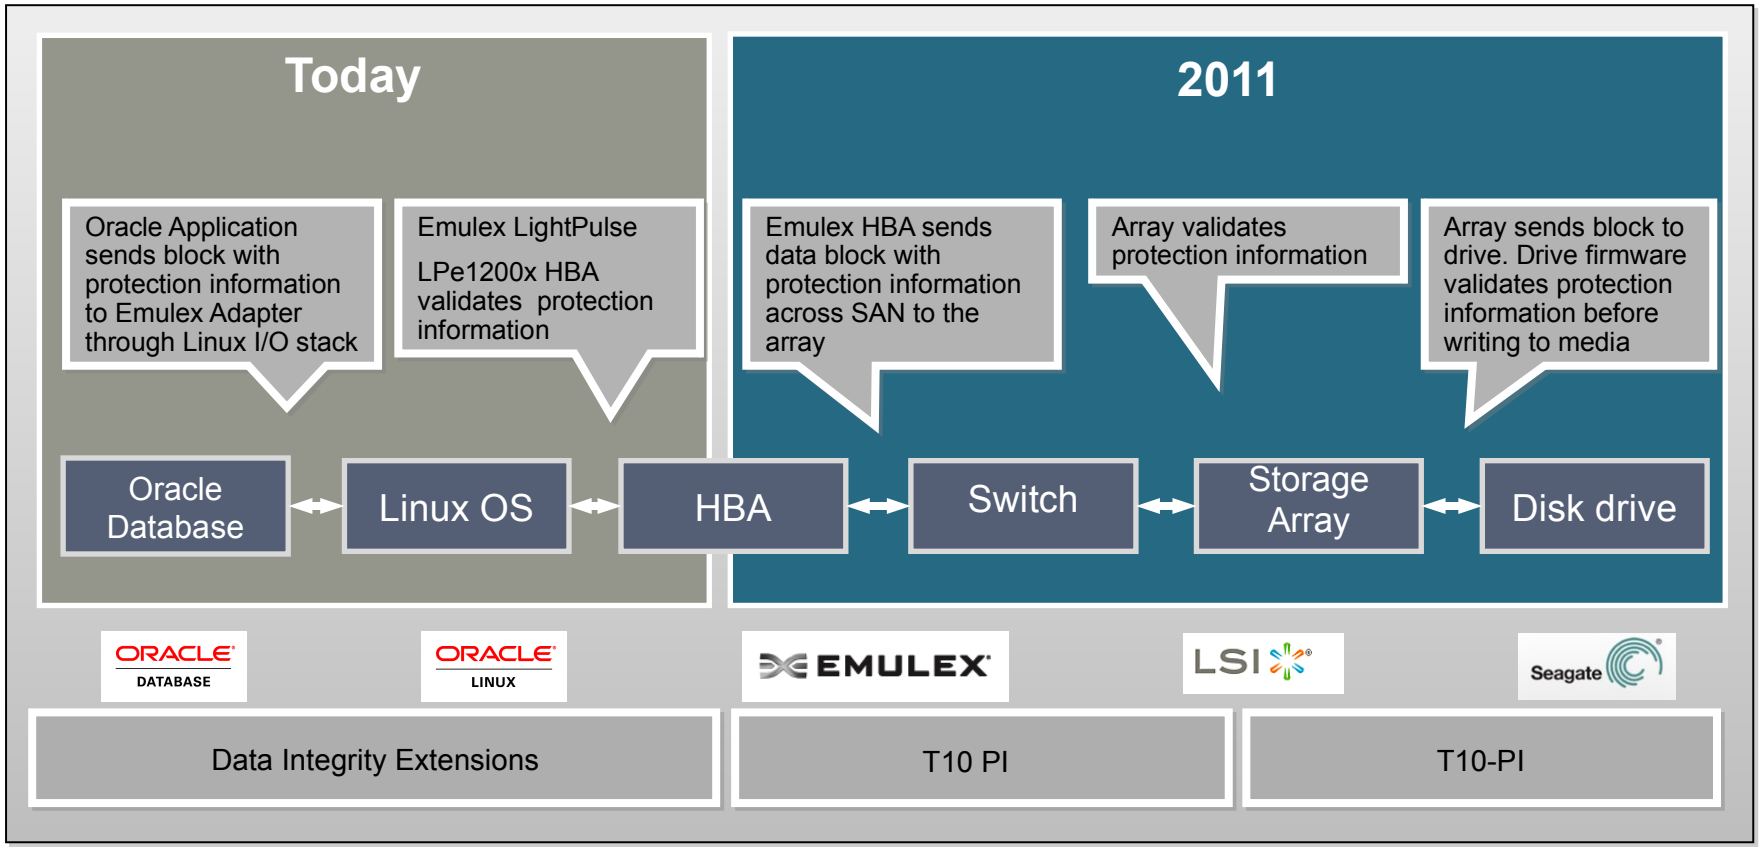

Application-to-Disk Data Integrity: current and future support

Improved Hardware Fault Management

- Hardware errors detected and logged before they affect OS or application
- Automatic isolation of defective CPUs and memory
- Avoids system crashes
- Improves application uptime

New Diagnostic Tools

- Performance Counters for Linux (PCL)
  - Kernel subsystem keeps track of hardware and software events
  - Tracing and analysis without affecting system performance
  - Find application and kernel CPU bottleneck
- Latencytop
  - Find what actions or operations that are causing latency in applications or in the kernel
- ftrace
  - Tracing framework for analyzing performance and latency

Performance Improvements

- Improved asynchronous write-back performance
  - Keeping up with fast storage
  - Improved buffered write accounting
  - Reduces stalls and inefficient writeback when mixing devices of different speeds
- Improved scalability on fast storage such as solid state drives
- NUMA improvements
  - Reduced page cache contention
  - Improves performance for large systems under load

Performance Improvements (2)

- I/O Affinity
  - Ensures processing of a completed IO is handled by the same CPU that initiated the I/O
- Receive Packet Steering (RPS)
  - Distributes the load of received packet processing across multiple CPUs

Other improvements

- Initial NFS IPv6 support
- RAID5 to RAID6 restripe support
- I/O topology support
  - Kernel tells application what drive requirements are
  - Improves write performance
- SSD detection
  - Block layer will try harder to dispatch I/O when it knows storage device is fast
- fallocate()
  - Speed up reserving space for large files

New Contributions

All Linux kernel enhancements described earlier for The Unbreakable Enterprise Kernel are open source and have been made available to the Linux community.

Other Oracle Open Source Contributions

- gcc: UTF16 and UTF32 support
- libstdc++: bugfixes, C++0x implementation
- Binutils: GOLD bugfixes and features
- Languages: PHP extensions for Oracle
- QA: testing and stabilization
- Tracing: systemtap
- Yast ports to RHEL/OEL
- Bugzilla plugin for OracleDB
- Some projects hosted at: http://oss.oracle.com

Oracle Linux: Summary

- Choice of two kernels
  - Oracle's Unbreakable Enterprise Kernel (default in Oracle Linux 6)
  - Red Hat Compatible Kernel
- Both kernels are Open Source Software and are included – Difference is in functionality
- Oracle Linux Support Program
  - Support pricing is same for both
- **Oracle recommends using the Unbreakable Enterprise Kernel**

Installing Unbreakable Enterprise Kernel

- Only needed for Oracle Linux 5.5 or RHEL 5.5 x86-64 (and above)
- Register system with ULN and subscribe to Oracle Linux 5 Latest channel
- Get Unbreakable Enterprise Kernel and recommended packages: **`up2date oracle-linux`**
- Unbreakable Enterprise Kernel and recommended packages for Oracle Database installation **`up2date oracle-validated`**
- Alternative download method (free for anyone)
  - Source and binary both available for download
  - Go to public-yum.oracle.com

Unbreakable Enterprise Kernel Roadmap

- 2.6.32 UEK
- NUMA improvements
- Infiniband & RDS improvements
- App-to-SAN DI

- New mainline kernel
- btrfs
- Vswitch
- Zcache
- Storage Connect
- Linux containers

- Continue to stay close to mainline kernel
- Test real world workloads

Roadmap / direction

- Btrfs production
  - Checksum on data and metadata (use hardware accel)
  - Data and metadata COW
  - Snapshots and subvolumes (versioned with rollback)
  - SSD optimizations
  - Small files are packed in as metadata (very efficient)

- Compressed page cache (zcache)
  - Cache that helps to keep more pages of the (filesystem) page cache longer in memory
  - Provides additional cache using LZO compression and thus results in fewer disk I/O operations

Zcache Iozone performance (1GB machine)

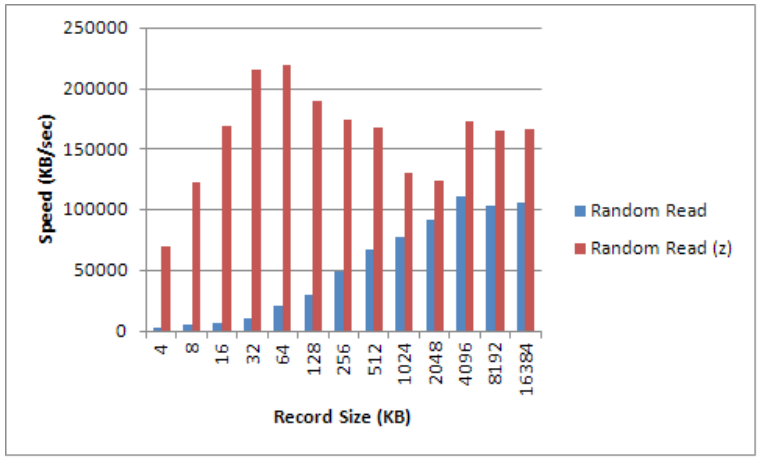

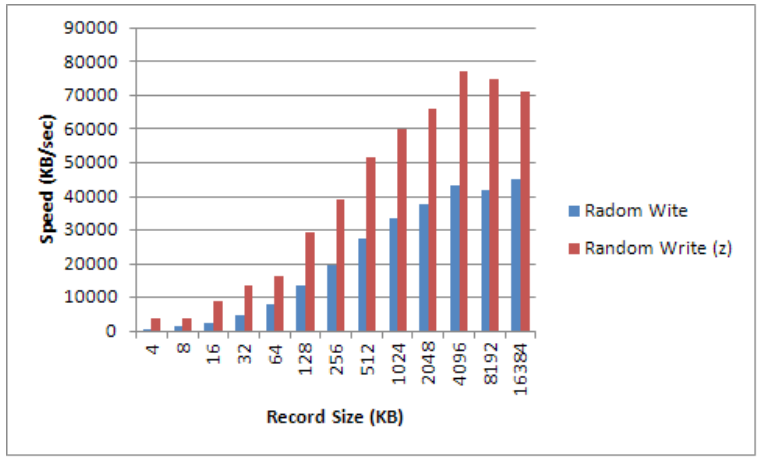

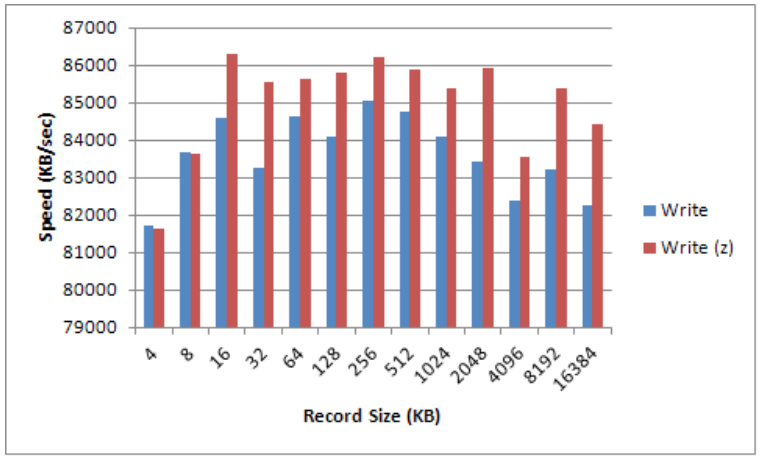

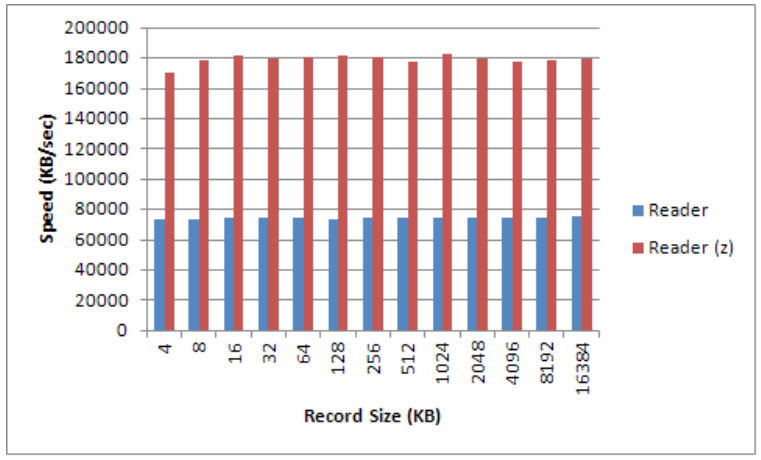

Other planned features

- Linux Containers
  - Resource management
  - Resource isolation (solaris zones/containers)
  - No overhead
- OpenVswitch
  - Built in virtual switch (similar to crossbow)
  - Flow monitoring (openFlow)
  - Per flow ACL, VLAN support, QoS

Linux Containers

- Isolated virtual program execution environments (VEs)
- Resource management
- Resource isolation
- Low overhead

Create
Reboot
Freeze
Destroy

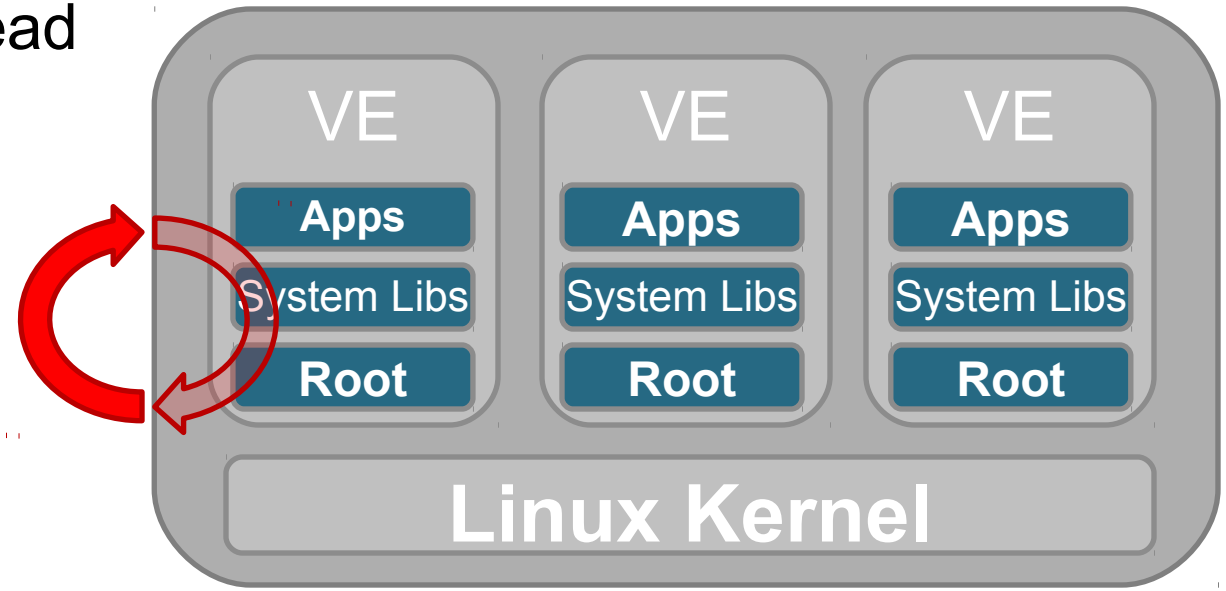

Oracle Linux Support Program

Enterprise-Class, Global, 24x7 Linux Support

Oracle Linux Support

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

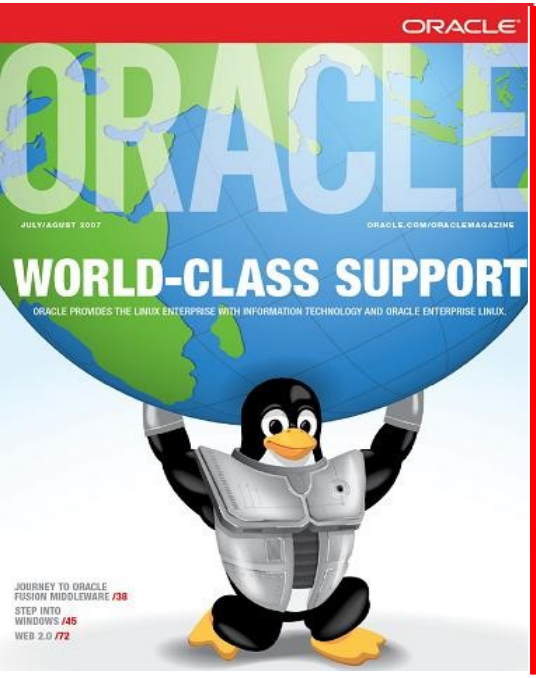

- 24x7 global coverage, 145 countries, 27 local languages
- Dedicated engineering and delivery team
- Backed by world's largest support team
- Enterprise-quality, Lower cost
- Service Excellence in Integrated Services (TSIA)
- 2008 Service Excellence in Mission Critical Support Award (SSPA)
- 2006 J.D.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

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 120 configurations published, freely available for download

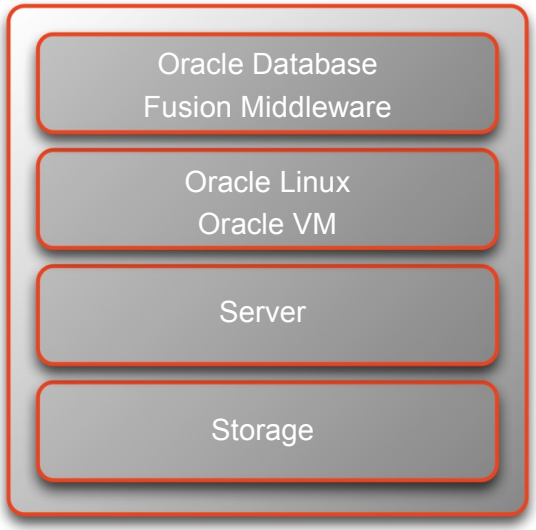

Oracle Validated Configurations offer faster Linux deployments while lowering infrastructure costs

Key Takeaways

- Oracle recommends Oracle Linux + Unbreakable Enterprise Kernel for all enterprise servers
  - Fast, Modern, Reliable; latest innovation in Linux
  - Supported as of Oracle Linux 5.5
  - Included in Oracle Linux 5.6 and Oracle Linux 6
- Oracle will continue to release and support the Red Hat compatible kernel
- No OS reinstall required to benefit from the new kernel and improvements

External Resources

- **Oracle Linux Home Page**
  oracle.com/linux
- **Follow us on Twitter**
  @ORCL_Linux
- **Free Download: Oracle Linux**
  edelivery.oracle.com/linux
- **Oracle Technology Network**
  oracle.com/technetwork/

Hardware and Software

ORACLE®

Engineered to Work Together


ORACLE®