
ORACLE®

ORACLE

ORACLE THIRTY YEARS 30

How green is red ?

Roland Rambau
Principal Sales Consultant

ORACLE

Where is red green

Roland Rambau
Principal Sales Consultant

Where does the Data Center Energy go ?

- VERY roughly (in a modern data center):

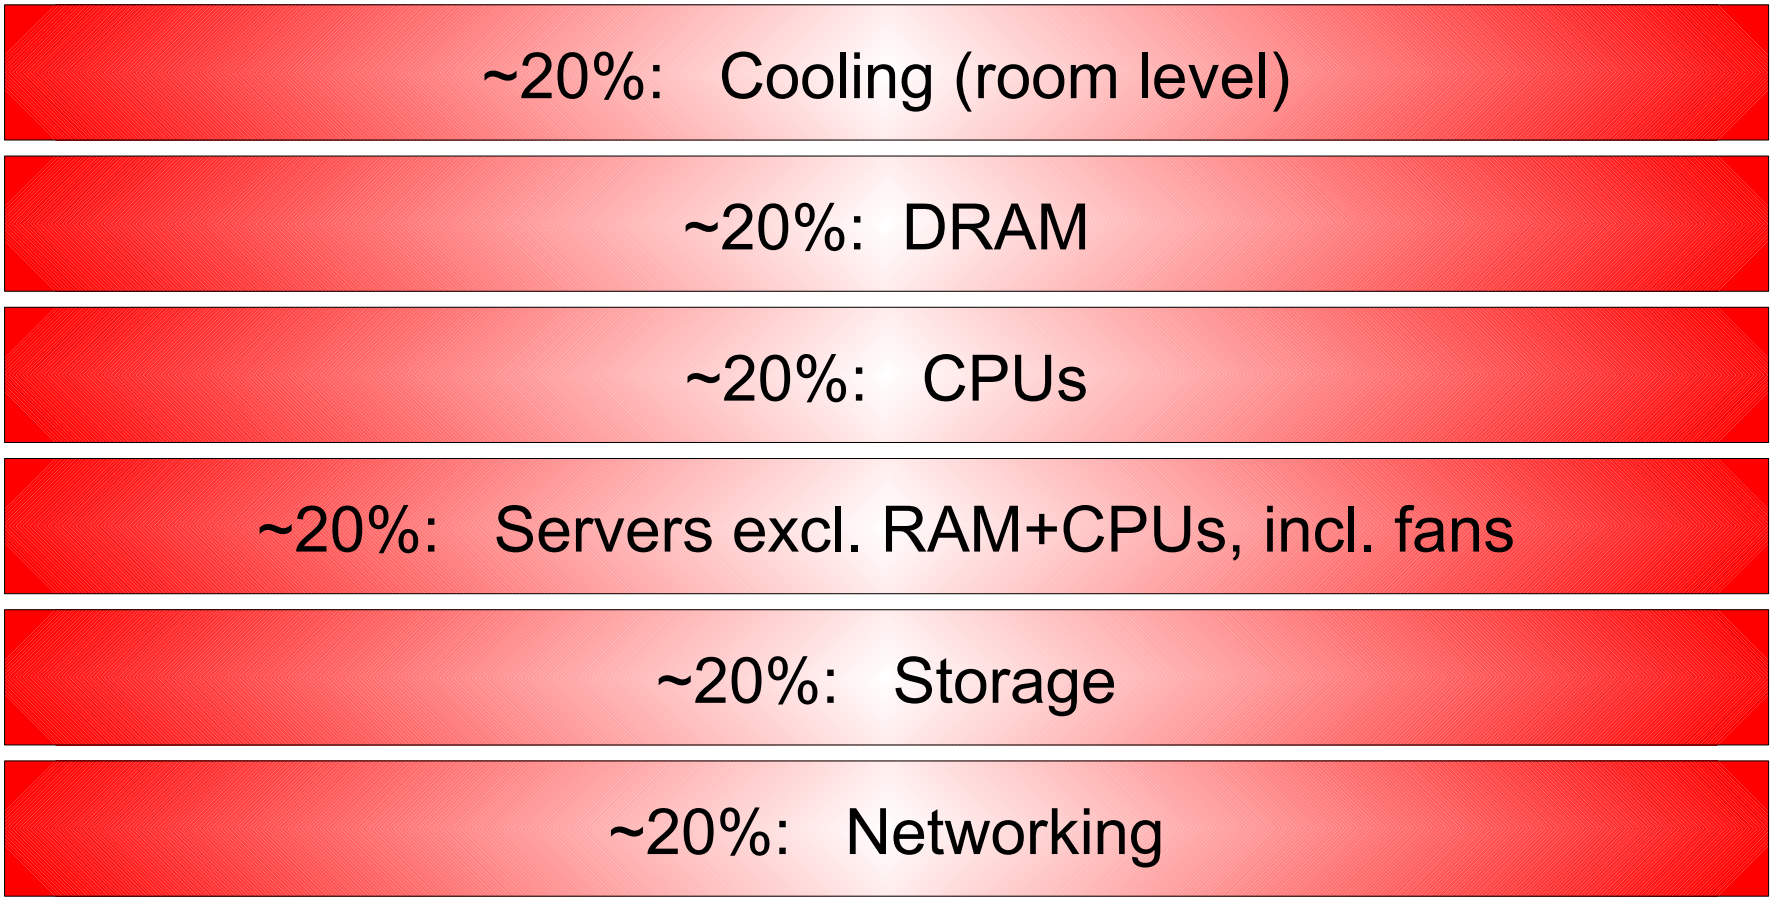

(~20% means 10% - 30% and YMMV quite a bit too)

Energy efficiency isn't just a chip or a hardware problem. It is a virtualization problem, an OS problem, a systems management problem, a networking problem, and a storage problem

Subodh Bapat, Sun 2009

Oracle x86 Cluster

Your Solution for Rapid Business Growth

- Engineered together
- Tested together
- Packaged together
- Certified together
- Deployed together
- Supported together
- Upgraded together

Oracle Virtualized eBusiness Suite Solution

Oracle Exadata Database Machine Version 2

Reduce management complexity

Performance and scale

Operational efficiencies

The Energy Efficient Stack

Energy efficiency through giving customers visibility into actual server power consumption, managing power draw across multiple servers, and providing optimization services

Energy efficiency through advising hypervisor of application workload requirements and supporting hypervisor directives to localize allocated resources

Energy efficiency through control of primary allocation of system resources (CPU, memory, I/O) to multiple guest Oss depending on need

Energy efficiency through control of memory DIMMs, I/O links, disks, and intelligence in firmware service processors

Energy efficiency through control of clock frequency, cores, threads, instruction pipeline and memory interfaces

Efficient Performance of partially utilized Resources

- Consequence: Oracle is really not that interested in optimizing the efficiency of partially loaded servers
  - Servers should always either be 100% utilized or powered down
  - Best energy efficiency is typically in highest performance system
  - The partially loaded server is not in our focus

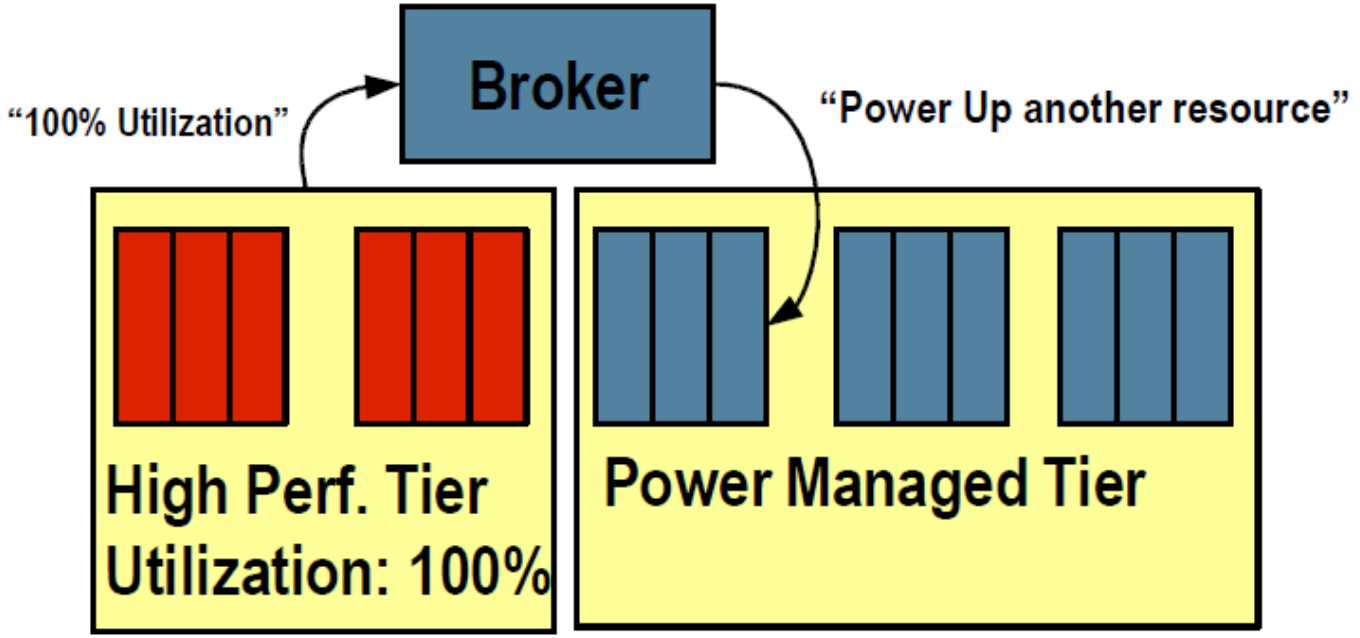

UltraSPARC Processor Roadmap

High Throughput, Energy Efficient, Fabless

2008 US-T2+	2010 US-T3	Future TBA	Future TBA	Future TBA
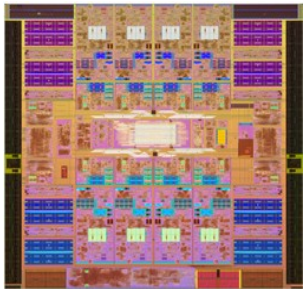	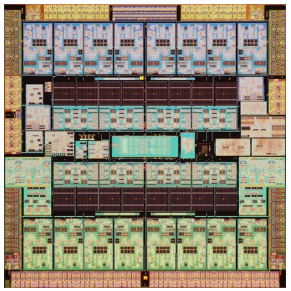			
• 3rd gen 8 Core CPU • 4 Socket Scalability • Faster Memory • High Speed Coherency Links • 65nm	• Doubled Cores • Improved FPU • Improved Security • Expanded Cache • Faster memory • Increased IO • 40nm	• New Advanced Processor Cores • Higher Frequency • System Compatible with US-T3 • 40nm	• More Cores • Higher Frequency • Larger Caches • Next Gen memory • Next Gen IO • Power Management • 28nm	• Higher Frequency • Much Larger Caches • Next Gen memory • Next Gen IO • Expanded Socket Scalability • 28nm
Expanded System Scalability	Increased System Throughput & Security	Improved Application Response Time	Increased System Throughput	Improved Everything

Binary Compatibility

The SPECpower standard benchmark

1 Load varied in 10% decrements from 100% to 0%

2 SSJ_ops calculated at each load point

3 Power measured at wall socket with approved external meter at each load point

4 SSJ_ops divided by Watts at each load point

Benchmark Results Summary

Performance			Power	Performance to Power Ratio
Target Load	Actual Load	ssj_ops	Average Power (W)	
100%	98.8%	84,913	288	295
90%	90.1%	77,489	280	277
80%	80.3%	69,012	268	257
70%	69.8%	59,971	258	232
60%	60.9%	52,386	251	209
50%	50.3%	43,226	243	178
40%	40.3%	34,638	236	146
30%	30.2%	25,952	230	113
20%	20.1%	17,275	224	77.1
10%	10.8%	9,326	220	42.4
	Active Idle	0	215	0
			Σssj_ops / Σpower =	175

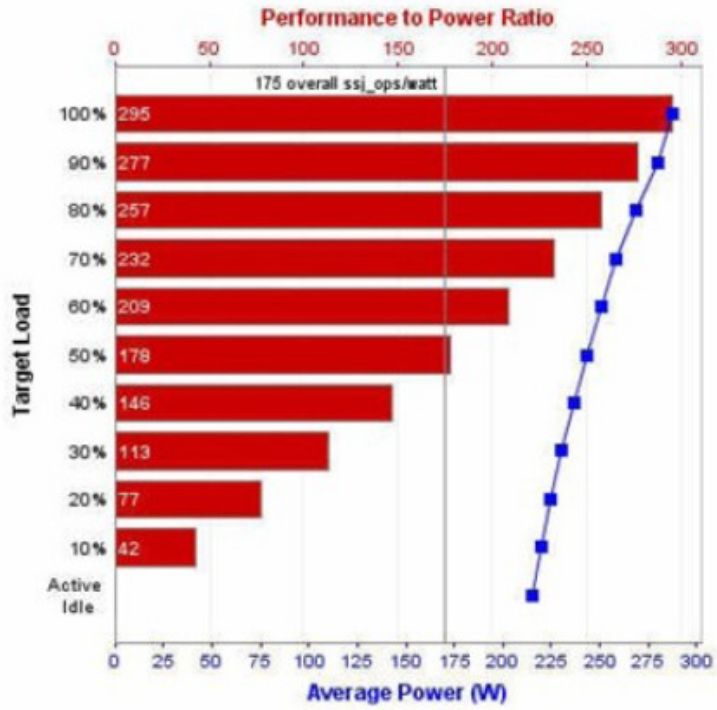

5 Total SSJ_ops at all load points (including 0%) divided by total Watts at all load points

6 Final SPECpower number for system (higher is better)

The SPECpower standard benchmark

1 Load varied in 10% decrements from 100% to 0%

2 SSJ_ops calculated at each load point

3 Power measured at wall socket with approved external meter at each load point

4 SSJ_ops divided by Watts at each load point

Benchmark Results Summary

Performance			Power	Performance to Power Ratio
Target Load	Actual Load	ssj_ops	Average Power (W)	
100%	98.8%	84,913	288	295
90%	90.1%	77,489	280	277
80%	80.3%	69,012	268	257
70%	69.8%	59,971	258	232
60%	60.9%	52,386	251	209
50%	50.3%	43,226	243	178
40%	40.3%	34,638	236	146
30%	30.2%	25,952	230	113
20%	20.1%	17,275	224	77.1
10%	10.8%	9,326	220	42.4
	Active Idle	0	215	0
			Σssj_ops / Σpower =	175

Total SSJ_ops at all load points (including 0%) divided by total Watts at all load points

5

6 Final SPECpower number for system (higher is better)

Sun Fire X4470

Best Mission-Critical Virtualization Platform

Prozessoren
- 4 oder 2 Intel Xeon 7500
  - X7560, X7550 oder E7530
    - 130 resp. 105 Watt

RAM
- maximal 512 GB mit 8GB DIMMS
  - max 64 DIMMs DDR3 registered mit ECC
  - 2GB, 4GB oder 8GB DDR3-1066

I/O
- 10 PCIe 2.0 Slots (x16, x8, x4)
- 4 GbE Ports
- 5 USB 2.0 (2 vorne, 2 hinten, 1 intern)

Interner Speicher
- 6 2.5” hotswap Festplattensteckplätze
  - 2.5” 146/300GB **SAS2** oder 32GB SSD
  - RAID HBA optional, erforderlich für SAS2

Hochverfügbarkeit
- 2 hotswap redundante PSUs (2060 W)
  - mit **Light Load Efficiency Mode (LLEM)**
- hotswap redundante Lüfter

Administration
- Oracle ILOM Service Prozessor

Advanced Thermal Engineering

- Pressurized chassis: Air pushed over components
- Cool components are more power efficient and last longer
- Two separate Air Flows to cool HDDs/PSUs and Motherboard

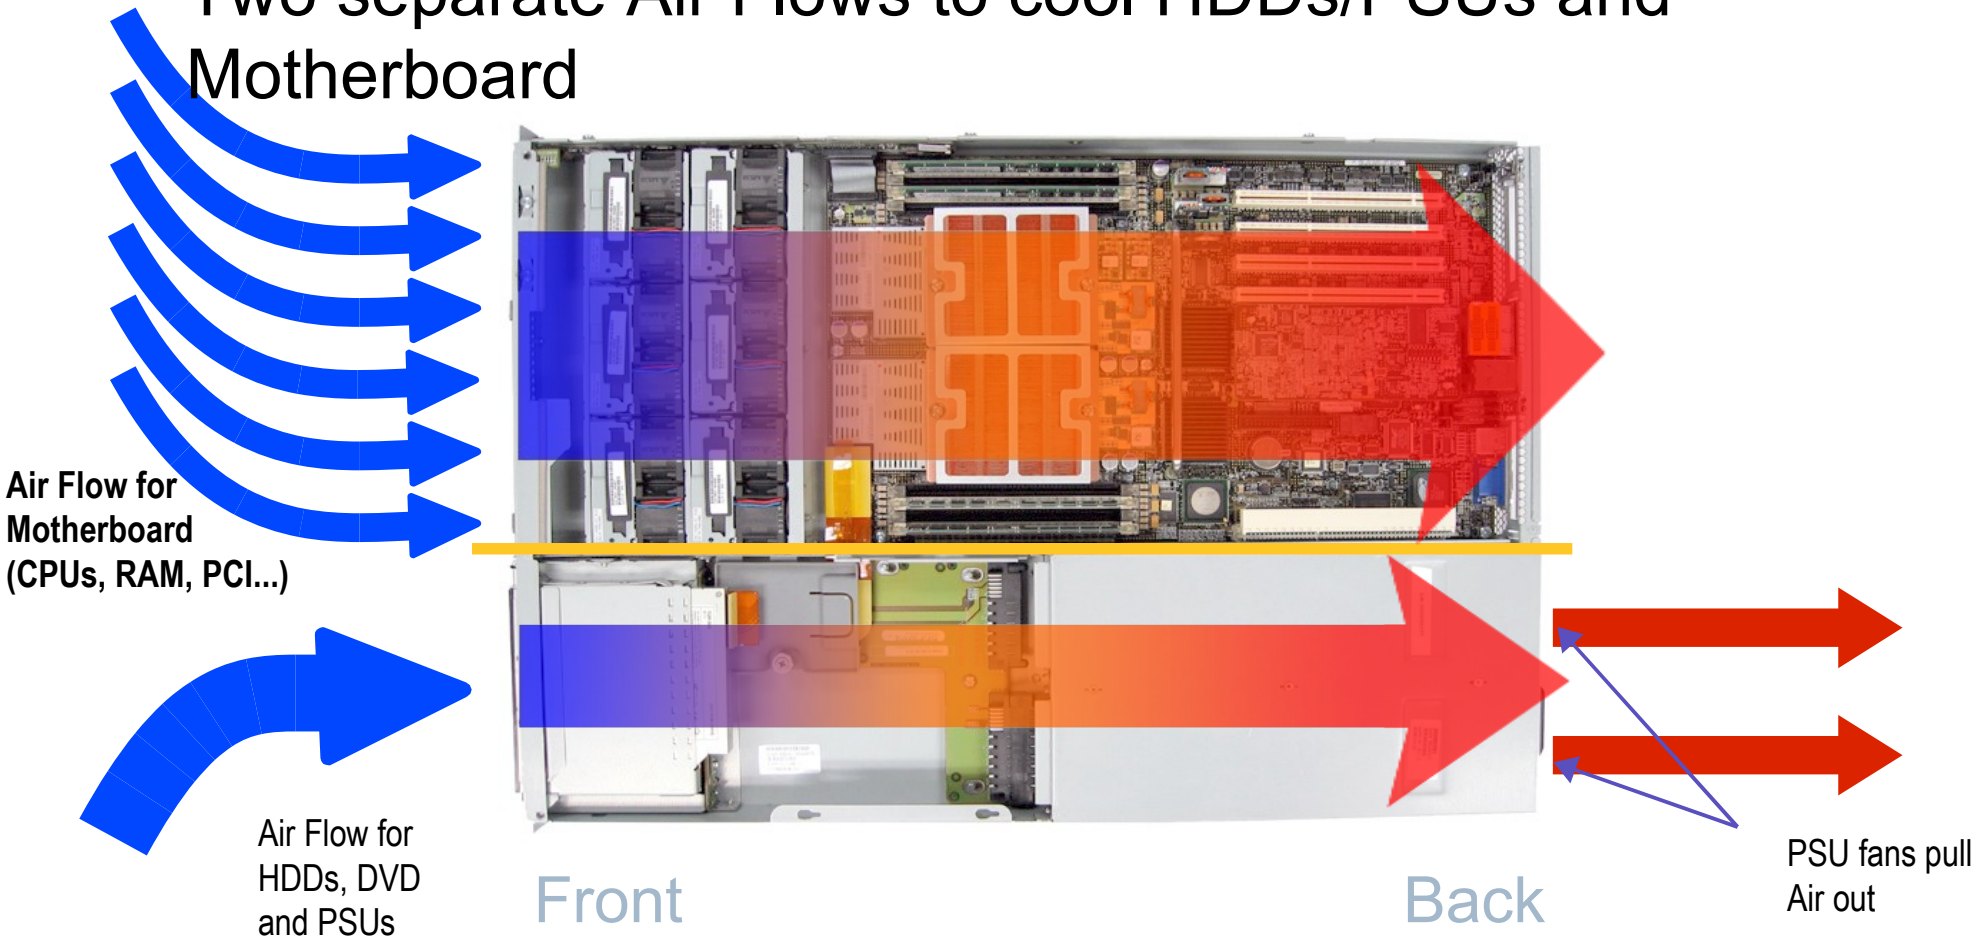

Sun Fire X4170 M2

Most Versatile
IT Infrastructure
Building Block

Prozessoren

- 1 oder 2 Intel Xeon 5600
  - X5670, E5620 oder L5640
    - 95, 80 resp. 60 Watt

RAM

- maximal 144 GB mit 8GB DIMMS
  - max 18 DIMMs DDR3 registered mit ECC
  - **4GB or 8GB DDR3-1333 low voltage (1.35V)**

I/O

- 3 PCIe 2.0 Slots (x16, x8, x8)
- 4 GbE Ports
- 5 USB 2.0 (2 vorne, 2 hinten, 1 intern)

Interner Speicher

- 8 2.5” hotswap Festplattensteckplätze
  - 146GB/300GB **SAS2**, 500GB SATA2, 32GB SSD
  - RAID HBA optional, erforderlich für SAS2
- DVD+/-RW

Hochverfügbarkeit

- 2 hotswap redundante PSUs (760 W)
- hotswap redundante Lüfter
- Oracle ILOM Service Prozessor

Low Voltage DIMMs

- New JEDEC standard summer 2010
- Can run at 1.35 Volt or 1.5 Volt
  - in principle also usable (useful) in older systems
- **significantly reduced power consumption**
  - Even when running 1.5V due to improved chip technology
  - at 1.35V operation about half the power of the older DIMMs
- Oracle exceeds the Intel PoR
  - Intel PoR for **1.35V** is:
    - only max 2 DIMMs per channel and only 1066 MHz max with 2
    - Other cases must switch to 1.5V
      - Can be a BIOS option “Performance” versus “Energy”
  - But Oracle has qualified with our LVDIMMs *more*:
    - No limitations compared to 1.5V, meaning 1.35V even with 2 DIMMs at 1333 MHz per channel, and also 3 DIMMs at 800 MHz per channel
      - **Never fall back to 1.5 V**

Sun Blade X6270 M2

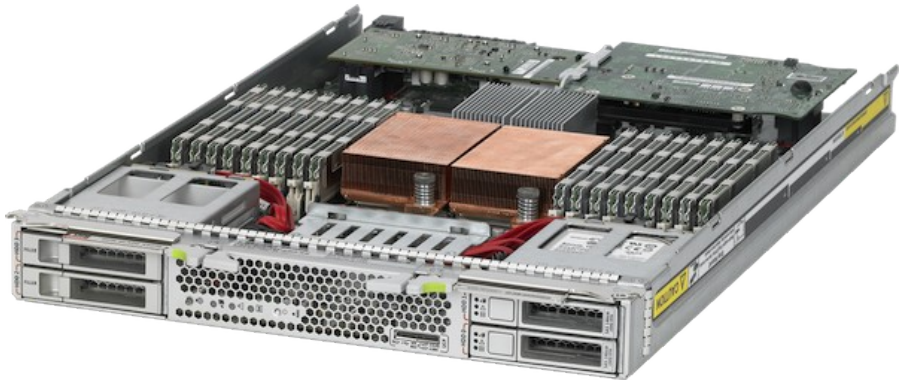

The Fastest Way
to Deploy
Business Applications

Prozessoren

- 1 oder 2 Intel Xeon 5600
  - X5680, X5670 oder E5620
    - 130, 95 resp. 80 Watt

RAM

- maximal 144 GB mit 8GB DIMMS
  - max 18 DIMMs DDR3 registered mit ECC
  - 4GB or 8 GB DDR3-1333 low voltage (1.35V)

I/O

- 2 PCIe 2.0 x8 ExpressModule Slots
- 2 PCIe 2.0 x8 to NEM slots
- 2 GbE Ports via NEM
- 3 USB 2.0 (2 via KVM Dongle, 1 intern)

Interner Speicher

- 4 2.5” hotswap Festplattensteckplätze
  - 146/300GB **SAS2**, 500GB SATA2, 32GB SSD
  - REM optional, erforderlich für SAS2

Hochverfügbarkeit

- **PSUs und Lüfter im B6000 Chassis**

Administration

- Oracle ILOM Service Prozessor

Sun Blade Power and Cooling Efficiency

- Low airflow resistance
  - All the components are aligned in the direction of the air flow
- Front to back cooling
  - No air intake from the hot isle or rack/cabinet
- Maximum cooling capacity
  - Full power processors (including 130W CPU's)
  - Maximum memory density per server
- Temperature based cooling
  - Increased fan speeds which consume more power are only used upon hardware monitor indicators
- Processor leakage power
  - Cooler CPU's draw less current. Increased airflow prevents CPUs from running too hot
- CMM manages the chassis power budget

Operating Expense Efficiency

	HP BL460c G6	IBM HS22	Sun Blade X6270 M2
Annual Power 	3,240 kWh	3,276 kWh	2,741 kWh
Annual Power Cost 	$334.40	$338.11	$282.87
Carbon Footprint 	1.95 tons CO_2	1.98 tons CO_2	1.65 tons CO_2

Up to 16% lower energy cost and environmental impact

Sun Cooling Door 5200/5600

Passive Rear Door Heat Exchanger Design

- No Additional Fans means greater efficiency
- 35KW Capacity – Room neutral cooling

Pumped Refrigerant Door (5600)

- Datacenter safe R134A refrigerant
- Compatible with Liebert XD systems
- Highest Energy Efficiency and smallest footprint

Chilled Water Door (5200)

- Low investment for those already with water in the datacenter
- Economical for smaller installations
- Connects to bottom (raised floor) or top (ceiling) water supply source

Oracle Solaris Studio

Build	Debug	Tune
Optimizing compilers and advanced math libraries	*Powerful debugging solutions*	*Advanced profiling and observability tools*
C, C++, Fortran Compilers Highest performance compilers for Solaris, with advanced code-generation technology resulting in record-setting performance across all supported architectures (SPARC, x86)	Debugger Scriptable, multithread aware, and provides memory error and leak information to track down difficult bugs	Performance Analyzer Identifies performance hotspots in serial and parallel applications, enabling you to quickly identify bottlenecks and tune for optimal performance
Performance Library Highly-tuned and parallelized scientific libraries for the newest Oracle Sun Systems (SPARC, x86)	Thread Analyzer Identifies hard-to-detect race and deadlock conditions in multithreaded code, allowing you to trace threading errors to source and fix issues	DLight System profiling tool that helps you utilize and visualize the power of DTrace on Oracle Solaris platforms

Next-generation IDE

Oracle Solaris Studio

20+ year history of delivering the best integrated development platform for building scalable, secure and reliable enterprise and ISV applications

Solaris Studio

Sun Studio

Sun ONE Studio

Forte Developer

Sun Workshop

SPARCWorks

Sun Compilers

1989 — New Name... Same Passion — 2010

ORACLE

Release History and OS Support

Product Name	Release Date	Solaris OS Support		Linux Support
		SPARC	x86 / x64	
Forte Developer 6	May '00	2.6, 7, 8	2.6, 7, 8	-
Forte Developer 6 update 1	Nov '00	2.6, 7, 8	2.6, 7, 8	-
Forte Developer 6 update 2	Jul '01	2.6, 7, 8, 9	2.6, 7, 8, 9	-
Sun ONE Studio 7	May '02	7, 8, 9	-	-
Sun Studio 8	Mar '04	7, 8, 9, 10	7, 8, 9, 10	-
Sun Studio 9	Jul '04	8, 9, 10	8, 9, 10	RHEL3; SLES8
Sun Studio 10	Jan '05	8, 9, 10	8, 9, 10	RHEL3; SLES8
Sun Studio 11	Nov '05	8, 9, 10	8, 9, 10	RHEL4; SLES9
Sun Studio 12	Jun '07	9, 10	9, 10	RHEL4; SLES9
Sun Studio 12 update 1	Jun '09	Solaris 10, OpenSolaris	Solaris 10, OpenSolaris	OEL5; RHEL5; SLES10
Oracle Solaris Studio 12.2	CY 2010	Solaris 10, OpenSolaris Solaris Express	Solaris 10, OpenSolaris Solaris Express	OEL5; RHEL5; SLES10

Bold denotes still orderable

The Sun Flash Storage Product Portfolio

Higher IO rates
At lower cost
And lower energy consumption

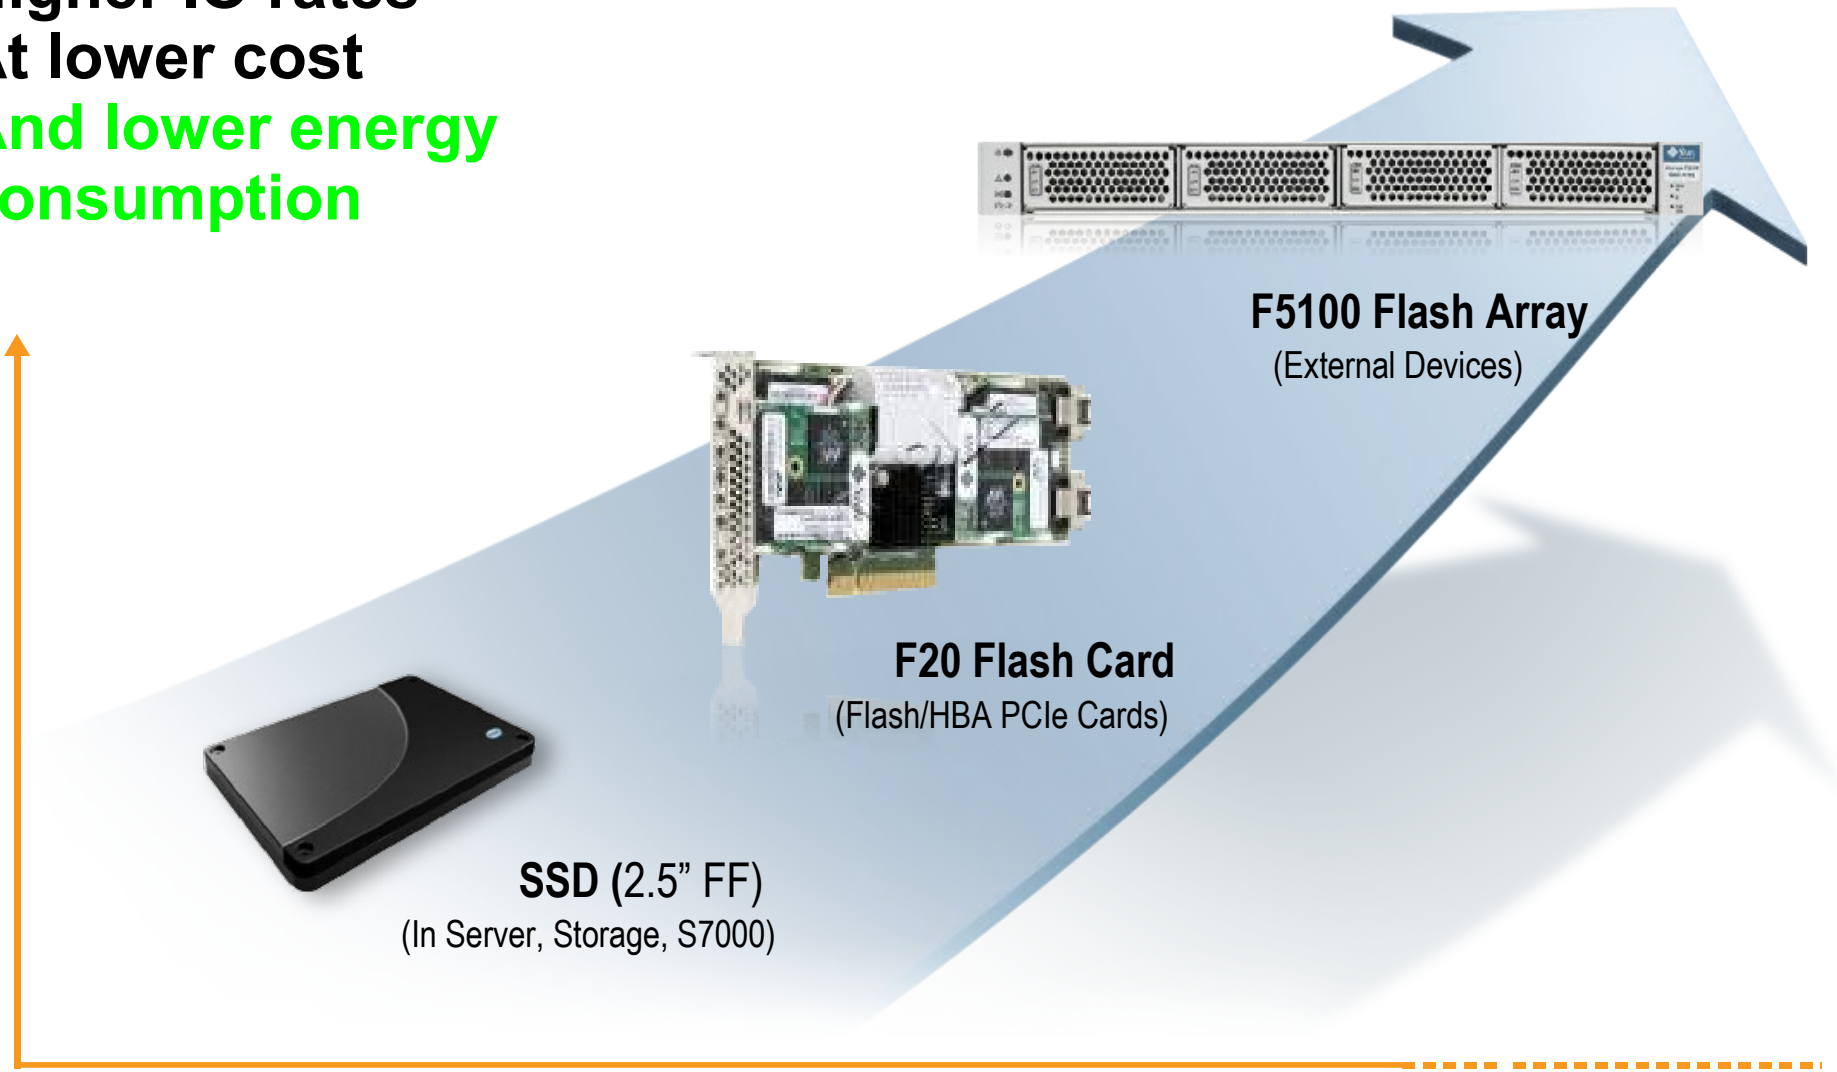

DB Flash Cache Scalability

- System Global Area (SGA) in Oracle DB is shared memory dedicated to an “instance”
  - SGA size is typically limited by available physical memory
- New Oracle Database 11g feature extends SGA by using the Sun Storage F5100 flash array
  - Benchmark shows up to ~5x Improvement by taking SGA into 100's of GB, at lower cost than system memory
  - High volume of SQL select transactions accessing a very large table in a typical business oriented OLTP database.
- Benchmark DB working set of about 3x the SGA size
  - Sun Storage F5100 Flash Array vs. without
  - Tests scale throughput with increasing Flash Cache size

Oracle Database & Flash

- Flash Enhances performance of Oracle Databases
  - New whitepaper details variety of possible improvements
- Read-only OLTP tests
  - 23% to 8x improvement - the lower the CPU utilization the higher the increase in performance
- Read-write OLTP tests
  - 13% to 4x improvement - the lower the CPU utilization the higher the increase in performance
- Decision Support query tests
  - 2% to 79 % improvement - improvements due to I/O channels bottlenecks
  - Index creation shows 20% improvement
  - Large Table Load shows 33% improvement
- Oracle cold backup shows 25% improvement

Solaris ZFS

Intelligent, Cost Effective Data Management

- End-to-end data integrity
  - 256-bit checksums
  - Copy-on-write transactions
- Easier administration
  - Pooled storage model—no volume manager
- Immense data capacity
  - World's first 128-bit file system
- Huge performance gains
  - Especially architected for speed
- **Compression**
- **Deduplication**
- **Snapshots + Clones**
- **Hybrid storage pools (flash + SATA discs)**

Best File System

Tiered Storage Architecture

Primary Tape Provides 40% Increase in Power Efficiency

Sun StorageTek SL8500 Tape Library

- 25-50% less space than comparable libraries
- 3x the performance of previous generation
- 2x drive capacity and performance with T10000 tape drive
- Zero cost static data storage
  - 0 kWh
  - 0 CO_2

The Most Complete Virtualization Portfolio

Desktop, Server, Storage Virtualization

DESKTOP VIRTUALIZATION

- Sun Ray Clients
- Oracle Virtual Desktop Infrastructure
- Oracle Secure Global Desktop
- Oracle VM VirtualBox

SERVER VIRTUALIZATION

- Oracle VM x86
- Oracle VM SPARC
- Solaris Containers
- Dynamic Domains

STORAGE VIRTUALIZATION

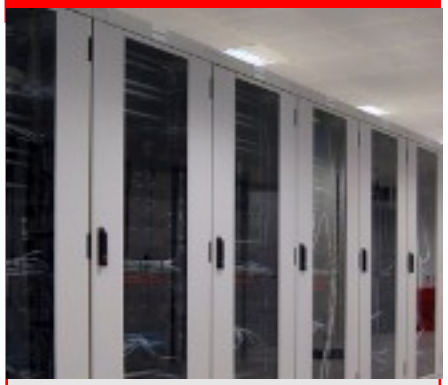

- Unified Storage
- Exadata
- ASM
- Storage Connect
- Storage Archive Manager
- Tape Virtualization

INTEGRATED SOLUTIONS

- Vertical Applications

Virtual Desktop Solution

Sun Ray and Secure Global Desktop

Sun Ray 2

- 4 watts
- Windows, Linux or Solaris desktop
- Standard monitor, keyboard, mouse

Uses Less Energy Than a Night Light

- Desktop runs on shared server
  - Accesses Solaris, Linux, Windows, mainframes and AS/400 apps
  - 10:1 – 50:1 Sun Ray to server ratio
  - No client upgrades required!

Customers save **$millions** annually in energy and systems costs

Servers	Qty	Racks	Power kWh/yr
V880	22	11	617,000
T2000	11	1	39,000
Savings	2x	10x	16x

Oracle's FCoE Benefits

Consolidation

- Reduced number of server ports
- Reduce number of switch ports
- Reduced cabling
- Reduced power consumption
- Increased speed of links
- Increased utilization of links

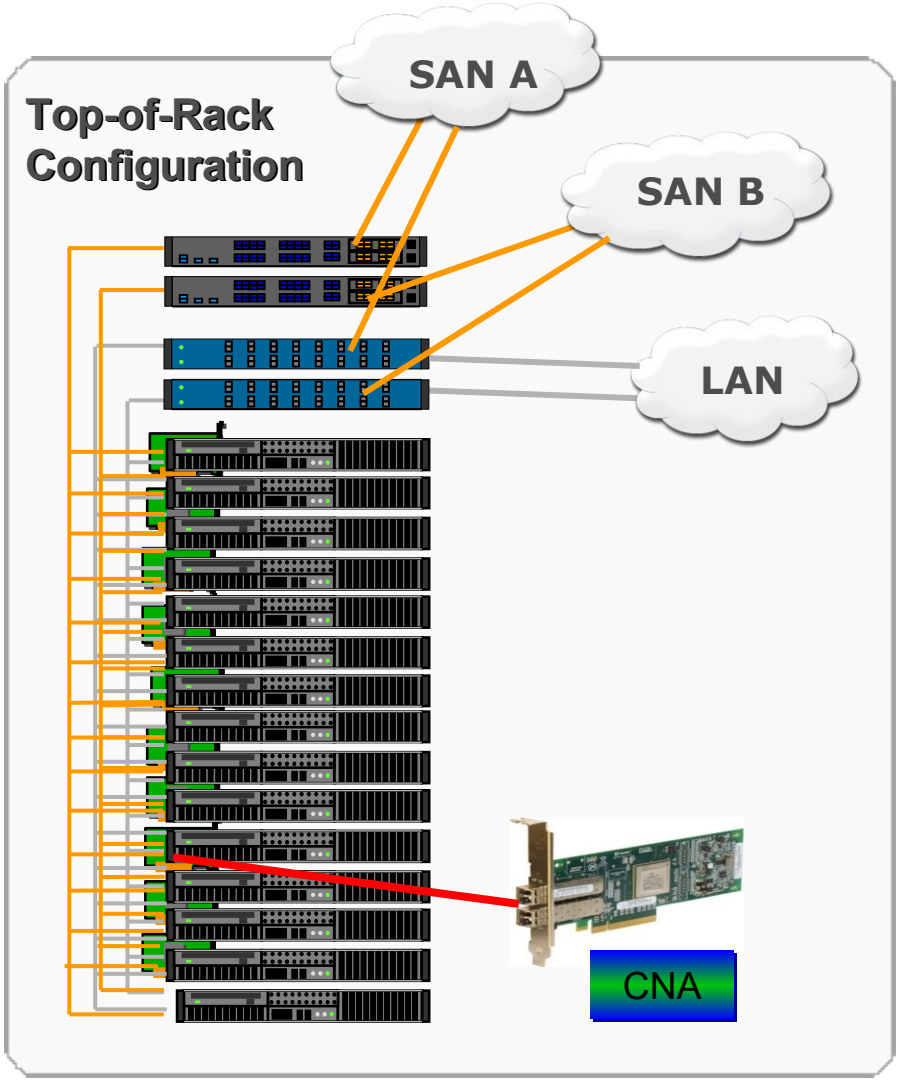

Consolidating I/O Interfaces Lowers CapEx and OpEx

Sun Blade 6000 Virtualized Multi-Fabric 10GbE M2 NEM

10:1 Cable Reduction - Zero Setup Management - Seamless Network Integration

Functionality	High bandwidth Network and Storage Virtualisation/Consolidation within the Sun Blade 6000 Chassis
Operating Systems	Oracle Solaris, Oracle Enterprise Linux, Oracle VM, Red Hat & SUSE Enterprise Linux, Windows, VMware
Blade Connectivity	Ten Virtualised 10GbE PCIe NICs, one per Blade Slot Two SAS2 6Gb/s Storage channels to each Blade Slot
External Connectivity	Two 10GbE Network Uplinks via SFP+ port Sixteen SAS2 Storage Uplinks (unsupported at launch) Ten 1GbE Ethernet Pass-Thru Ports
Management & Availability	Oracle ILOM Service Processor Up to two NEMs in each Sun Blade 6000 Chassis

Network Architecture Fabric Diagram

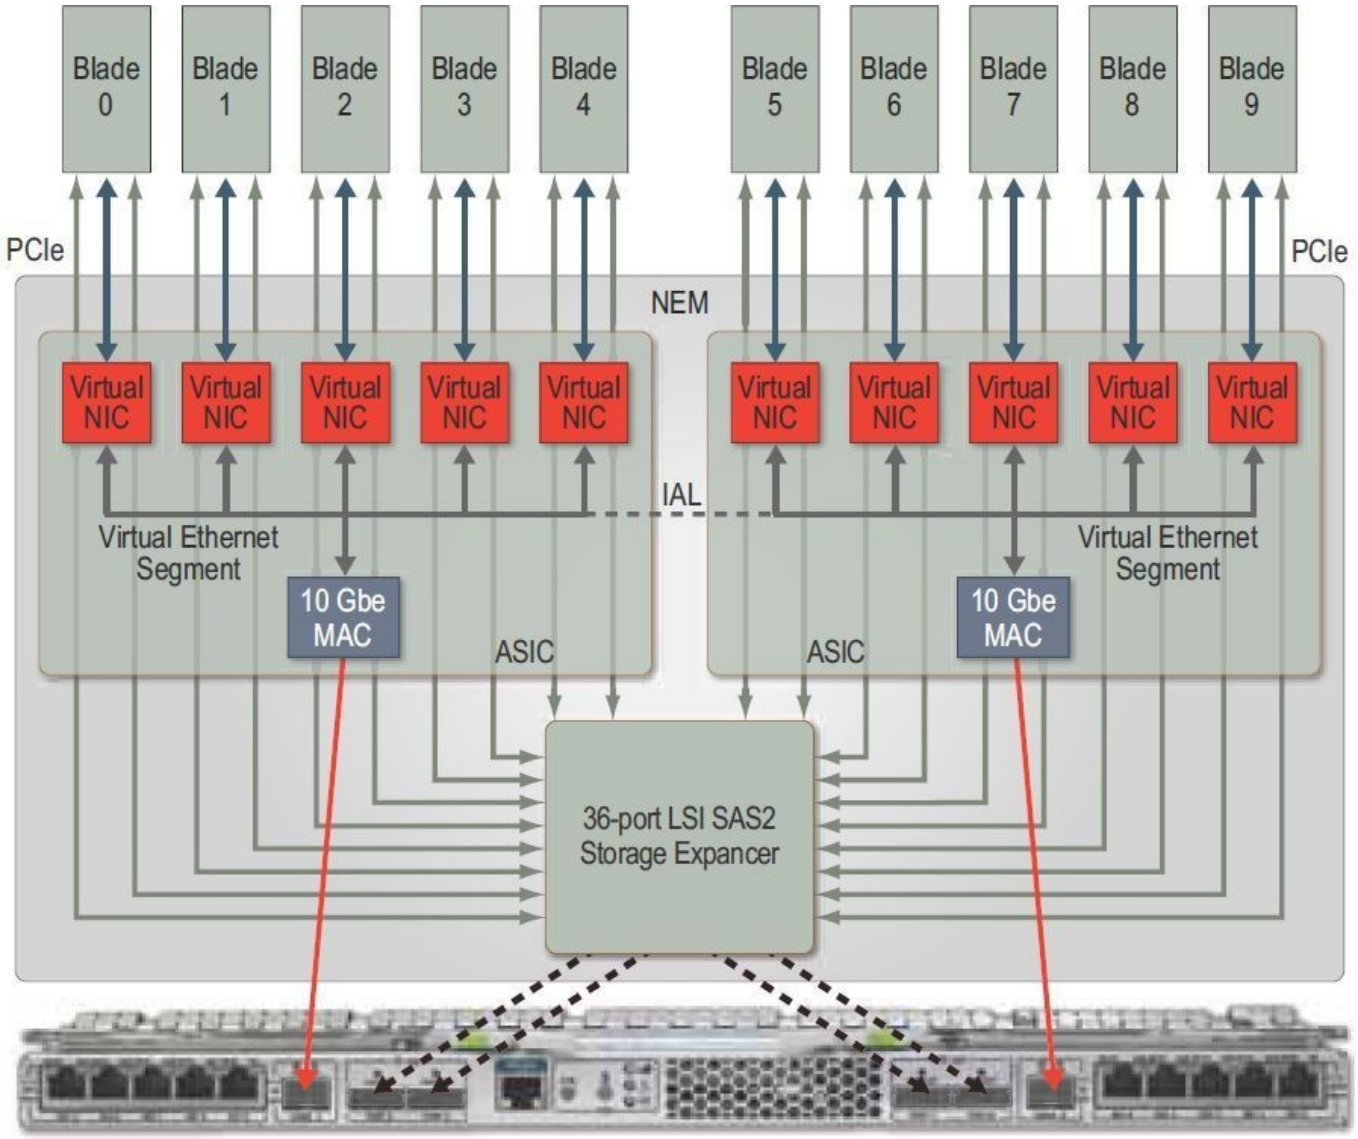

Hydra ASIC Overview

- 5 x8 PCIe endpoints, one to each blade (a 6th is unused)
- Five blades share a Hydra with one 10GbE uplink
- Leverages Sun Neptune 10GbE technology/drivers
- Each blade sees its own dedicated (virtual) NIC
  - Up to 16 MAC addresses per Blade slot
  - iSCSI, PXE protocols supported for boot
- Also a blade to blade loopback path, not a Switch/Hub
- 3 Modes of operation, configured through CMM GUI/CLI
  - Performance - Inter ASIC Link (IAL) disabled, 5-1 sharing of each 10GbE uplink (two separate broadcast domains)
  - Connectivity - IAL enabled, 10-1 sharing of one 10GbE uplink, second uplink is in “disabled” or “failover” mode (one broadcast domain)
  - Private - IAL enabled, uplinks disabled, traffic only visible to blades

Sun Colorado DC Consolidation

- Largest, most complex & costly consolidation in Sun's history
- 66% Datacenter compression
  - > 496k ft² to 126k ft²
- Scalable/Future Proof
  - > 7MW to 10MW
- First & Largest Liebert XD dynamic cooling install
- Water treatment saves 600k gallons a year, eliminates chemicals
- Waterside economizer, free cooling > 1/3 of year.
- Compressed 165k ft² raised floor to <700 ft² ($4M Cost Avoidance)
- Flywheel UPS, eliminates batteries.
- Chillers 32% more efficient at avg load than ASHRAE std
- 2 ACE Awards
- Removed 1M kWh per month

Oracle Utah Compute Facility (UCF)

Oracle's Newest Green Data Center

another center in California currently testing 380V DC as std power

Innovation Is Everything


Sun
ORACLE

SOFTWARE. HARDWARE. COMPLETE.


ORACLE®