



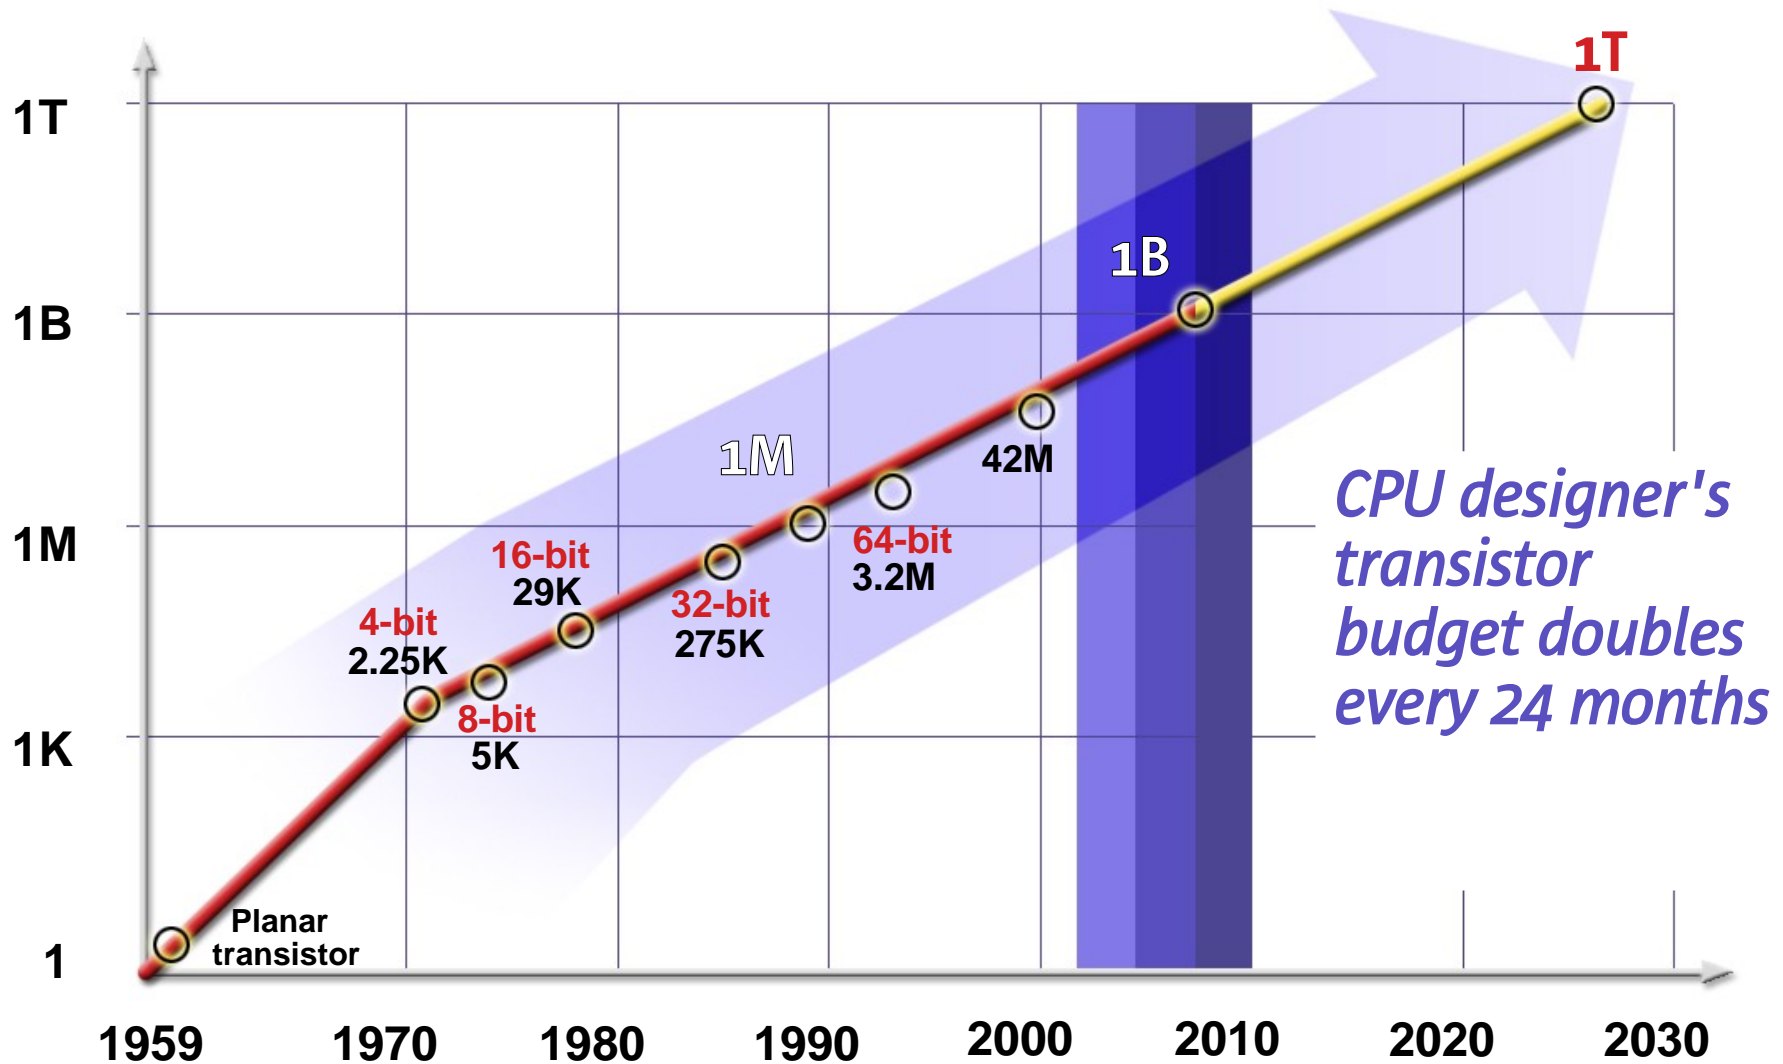
Oracle 10g CoolThreads and Containers

Michael O'Connor
Engagement Architect
Sun Microsystems, Inc.

Agenda

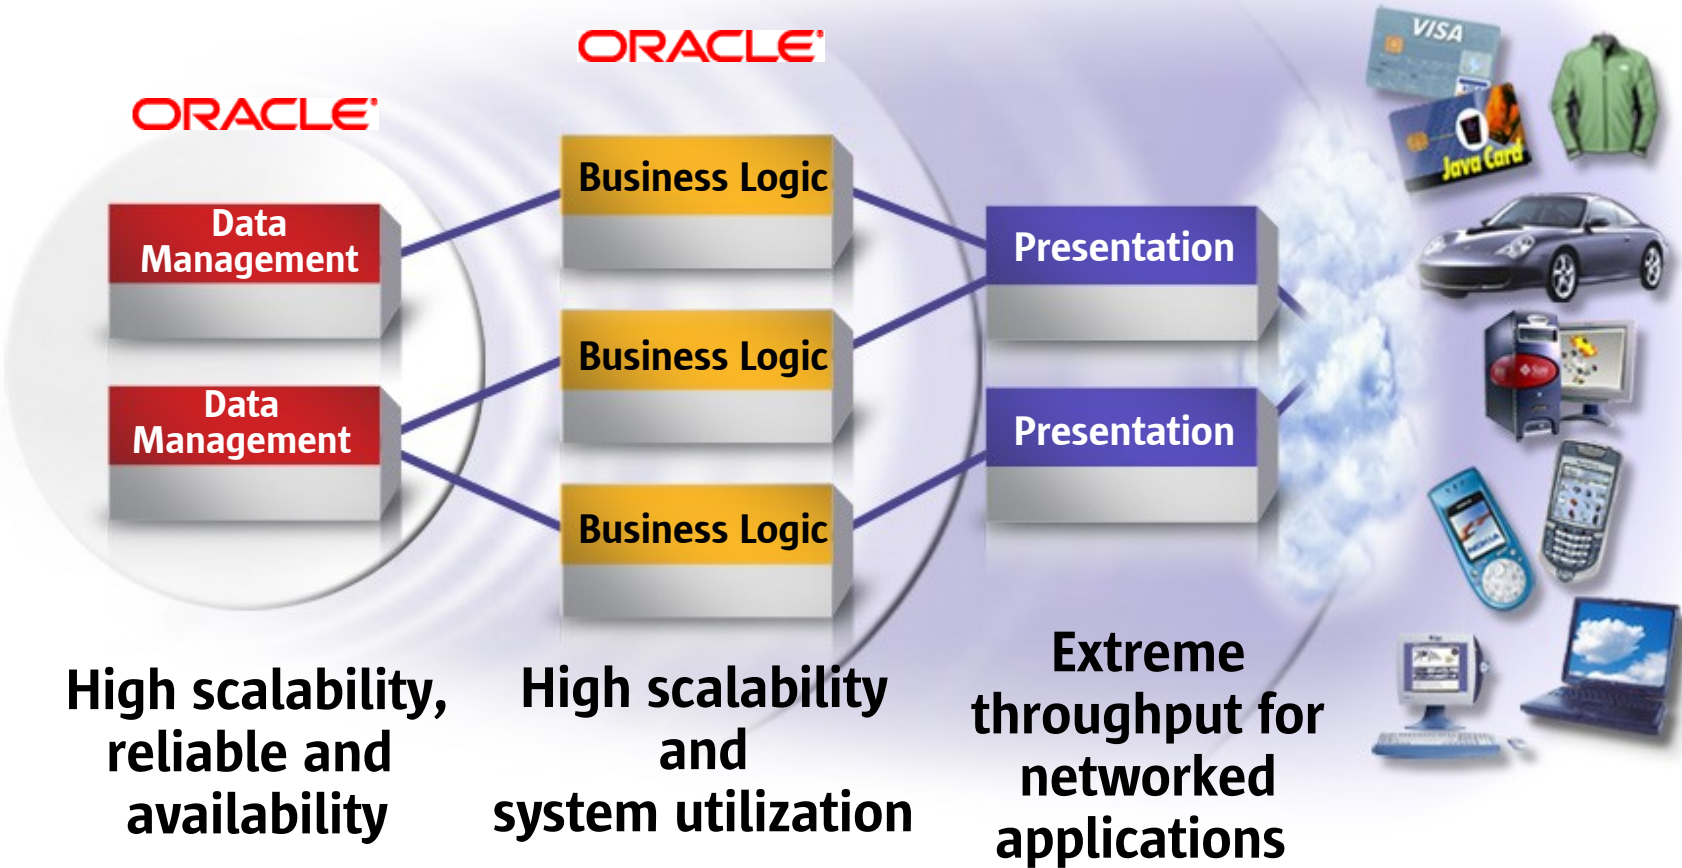
- Motivation for Niagara
- Solaris virtualization with Containers
- 10g on Niagara Performance Results

Moore's Law

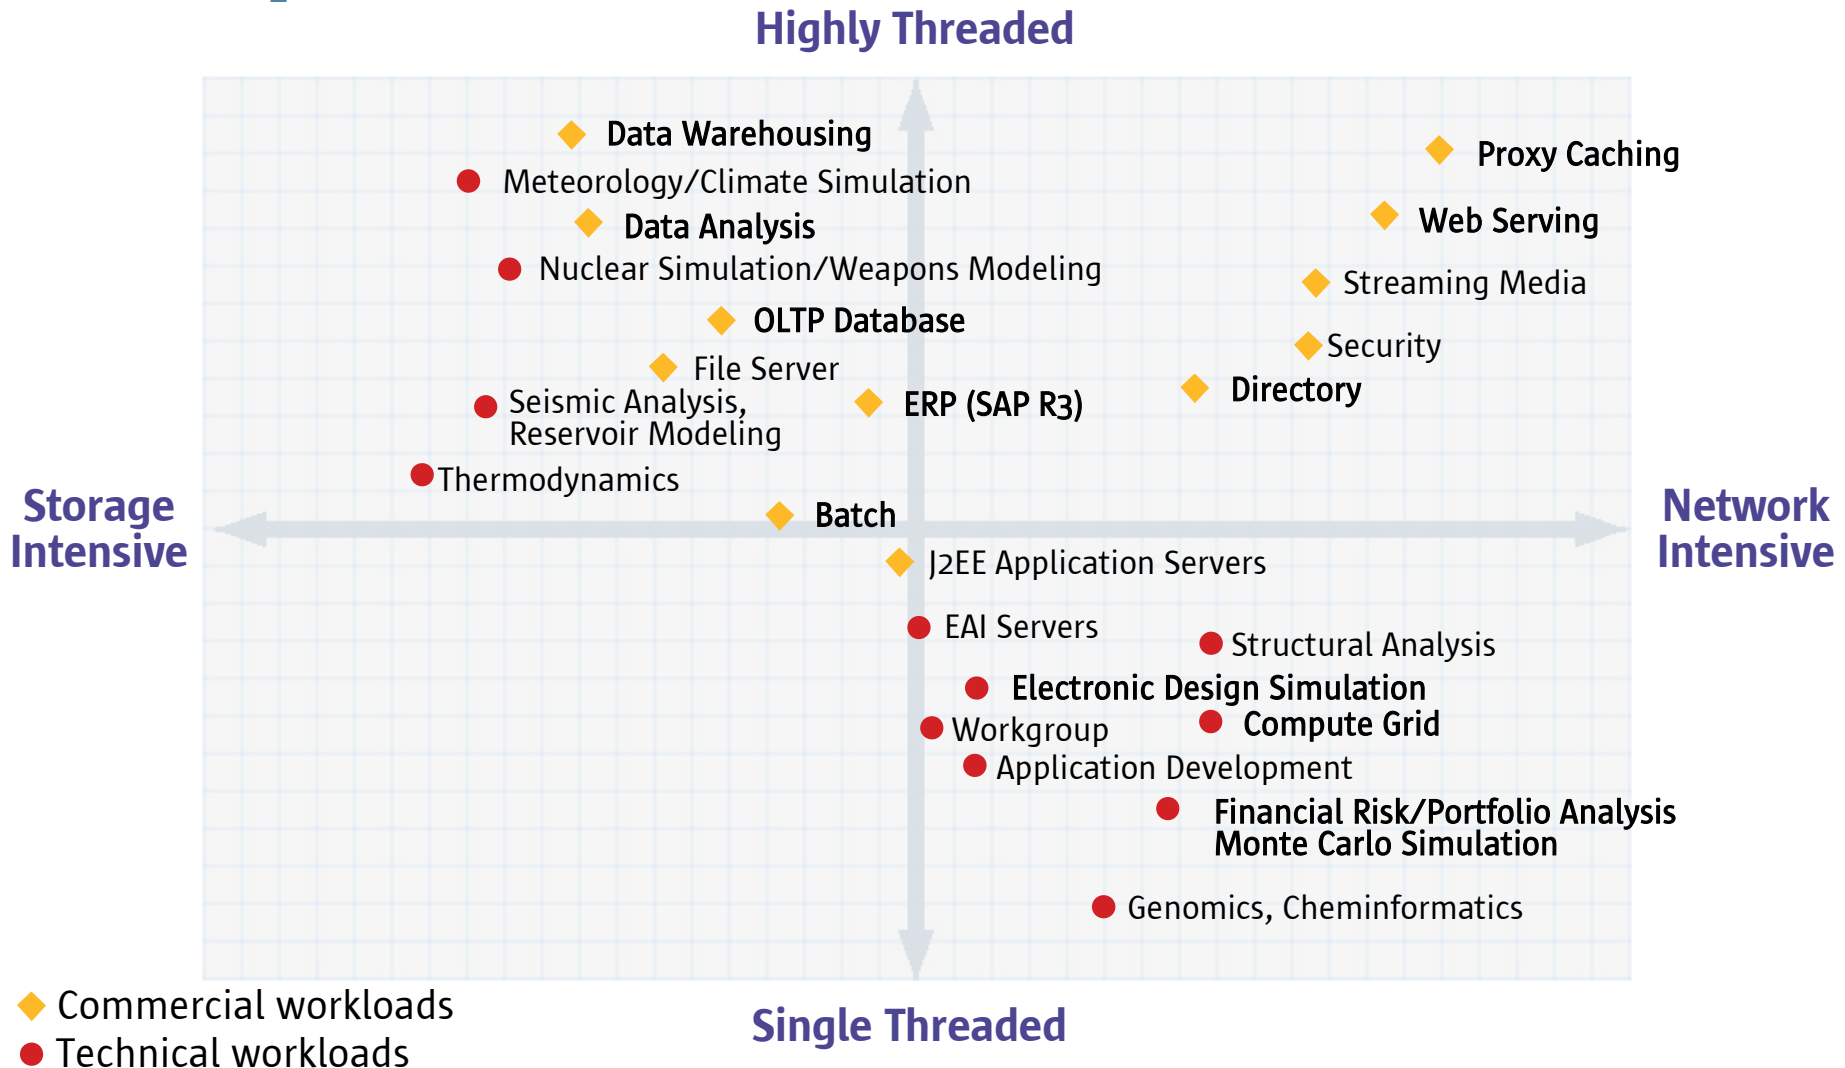


Sun Fire™ Servers

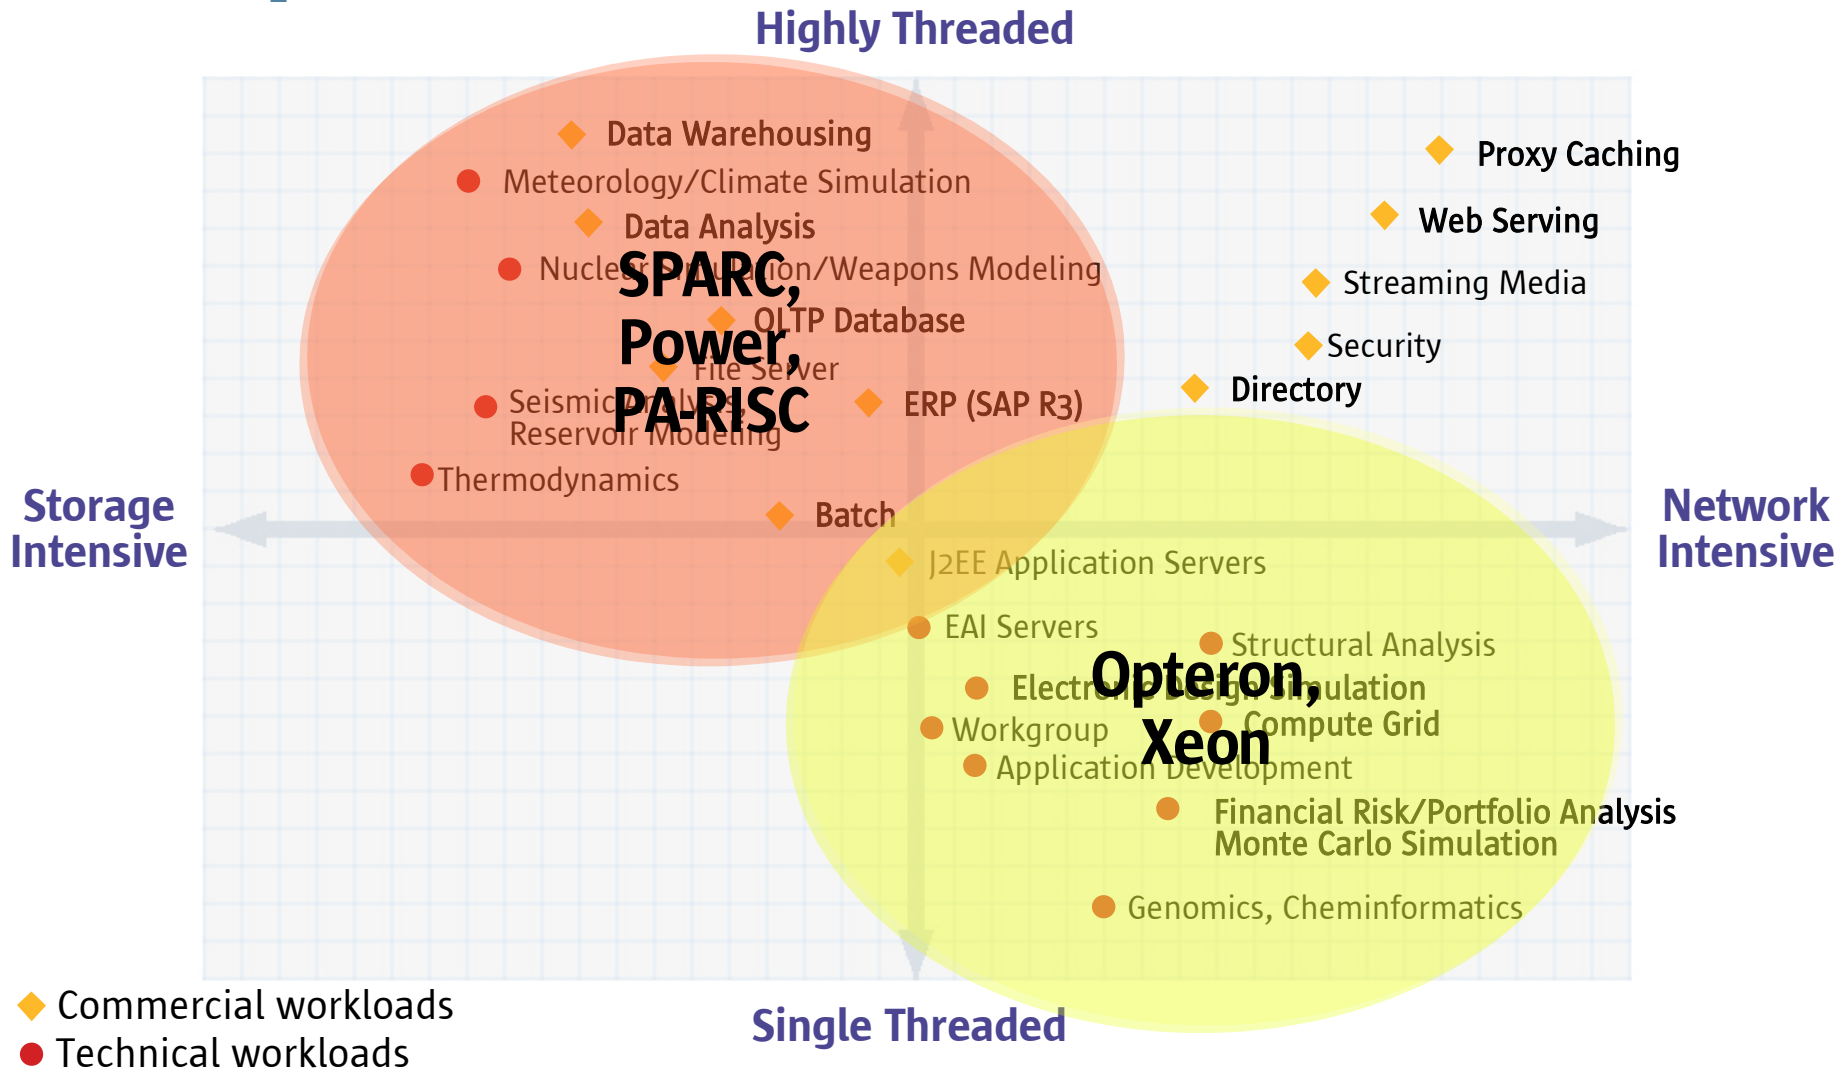
Optimized for Applications Throughout the Data Center



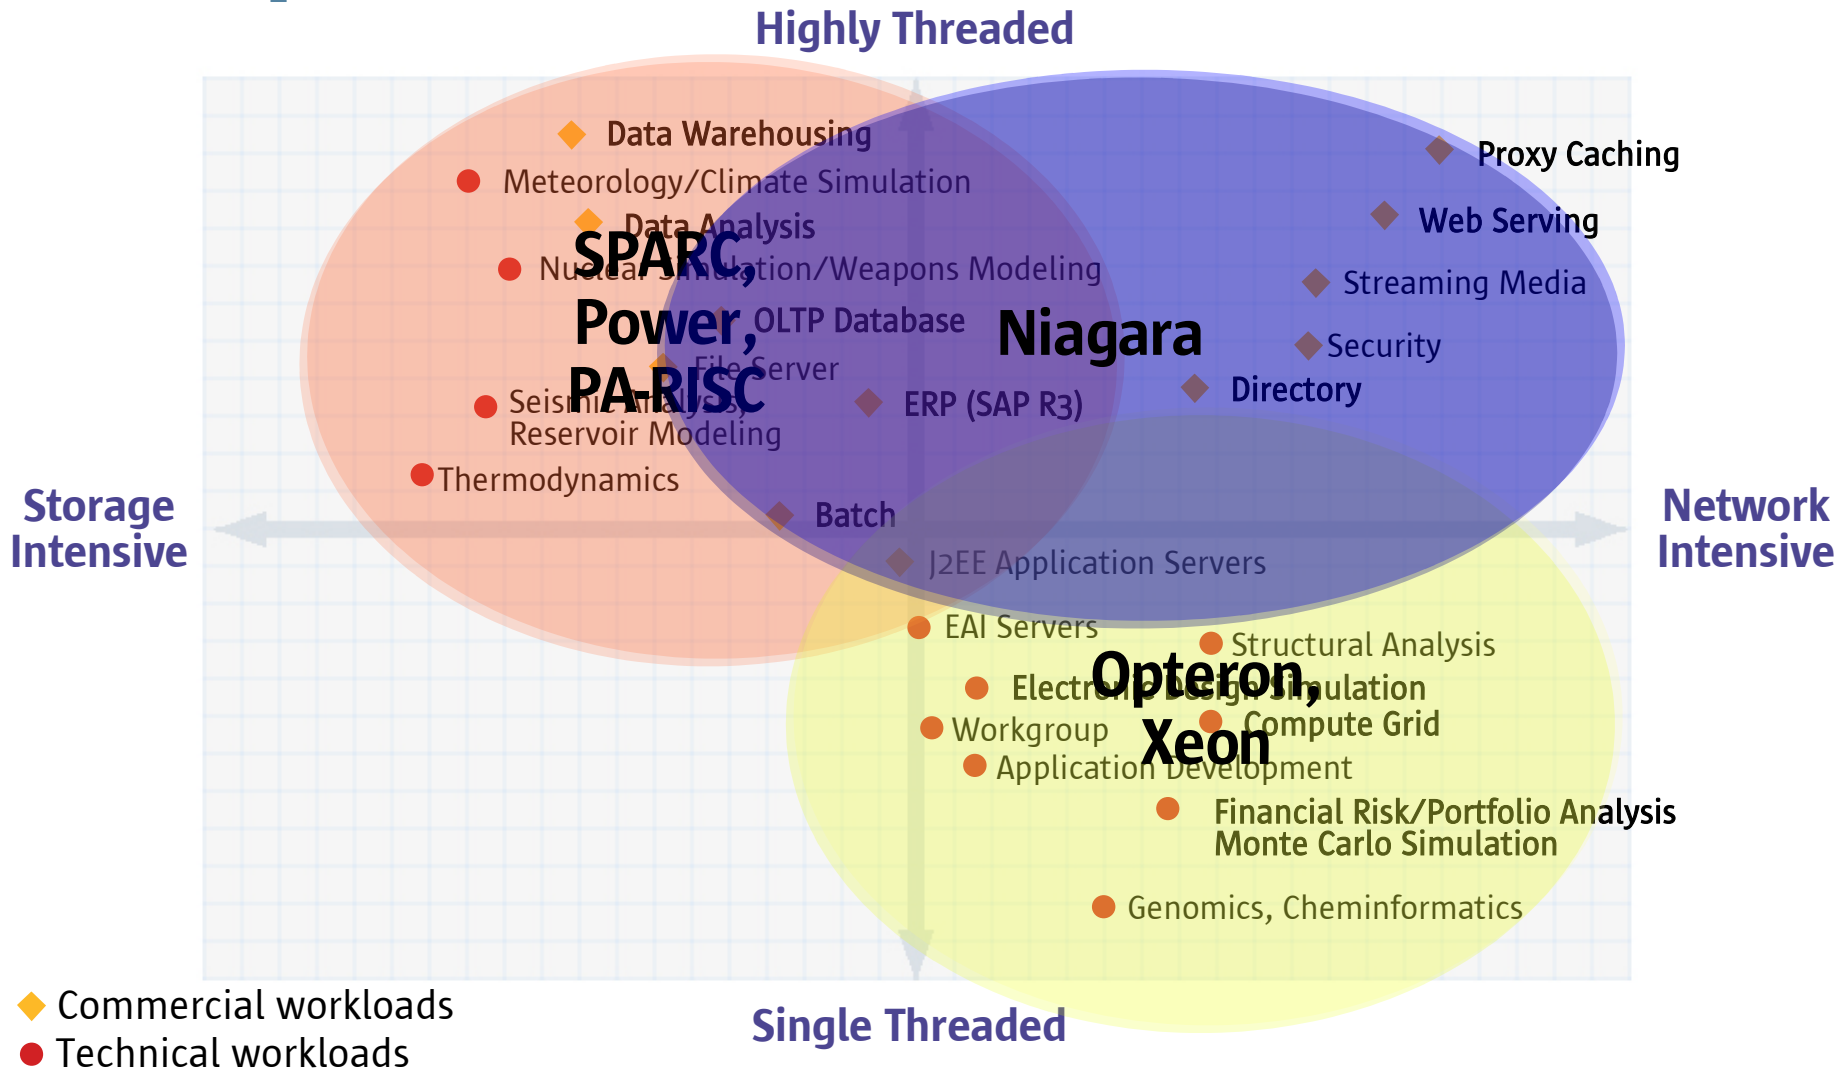
Mapping Workloads to System Requirements



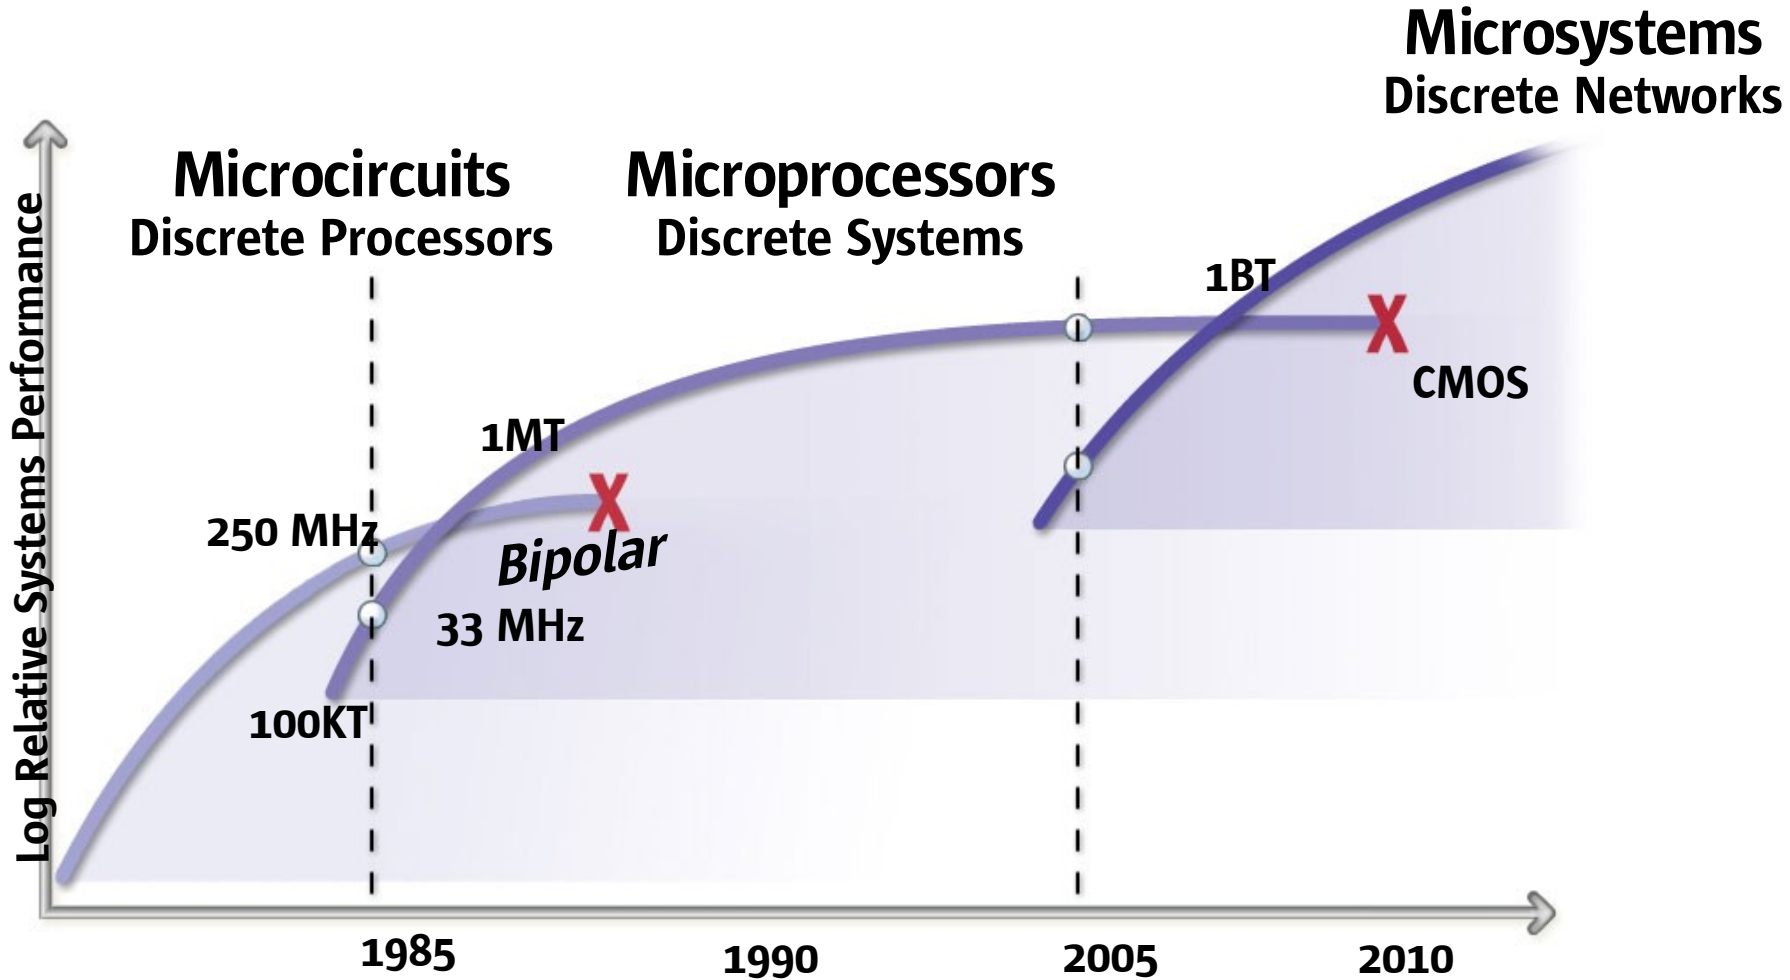
Mapping Workloads to System Requirements




Mapping Workloads to System Requirements



Semiconductor Technology Advances Enable Disruptive Change




Typical Complex High Frequency Processor using Instruction Level Parallelism (ILP)




Note: Up to 75% Cycles Waiting for Memory

Why faster freq. isn't producing corresponding performance increases

Relative Performance

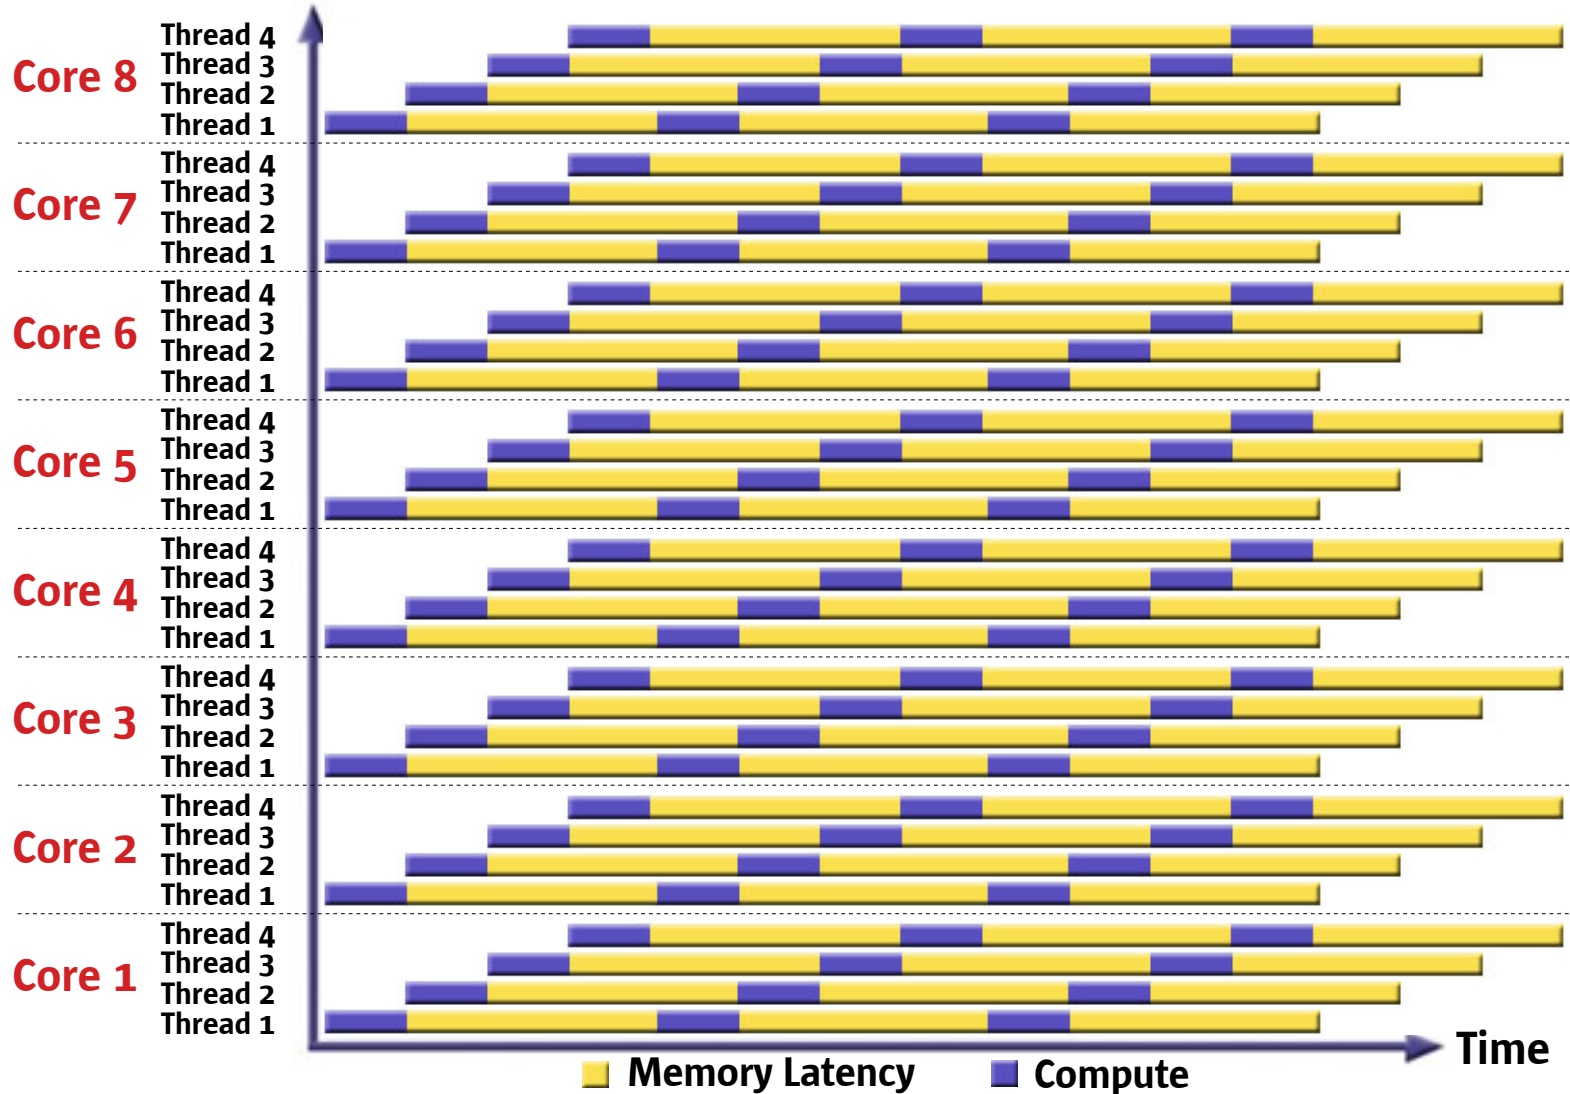


A new Approach: Chip Multi-Threading (CMT)

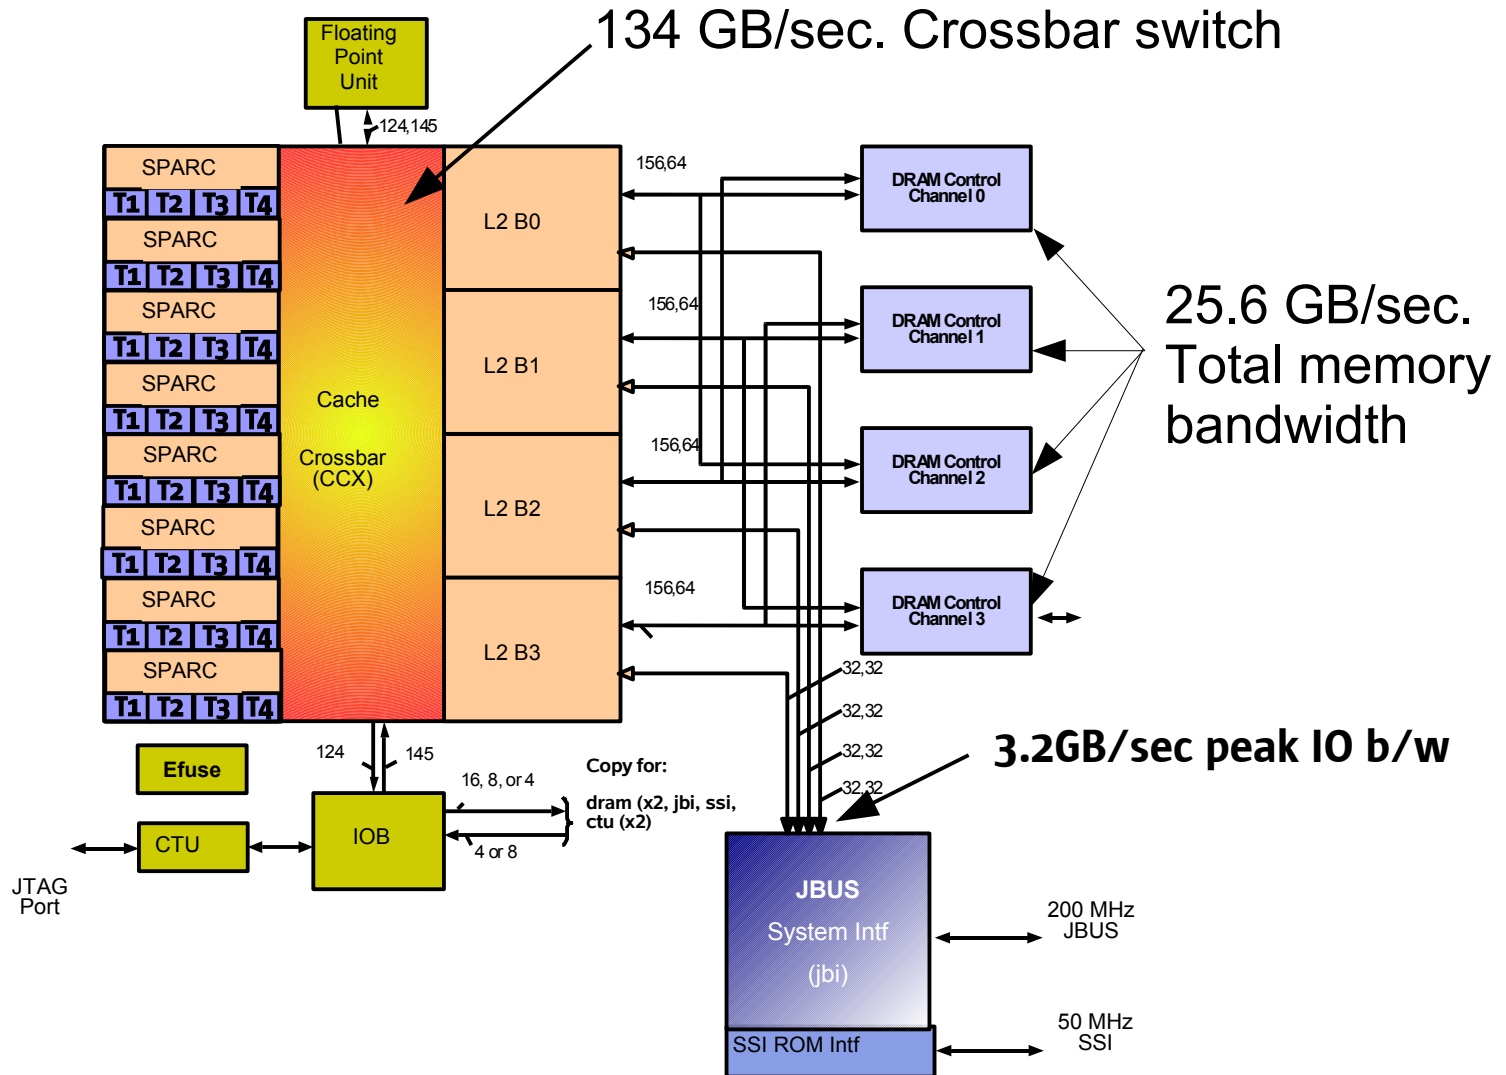


Processor Core busy 100%!

Sun's CMT – Massive Parallelism




The UltraSPARC T1 Multicore Processor




Network Facing CMT Processor Thermal Advantage

“Cool Threads”

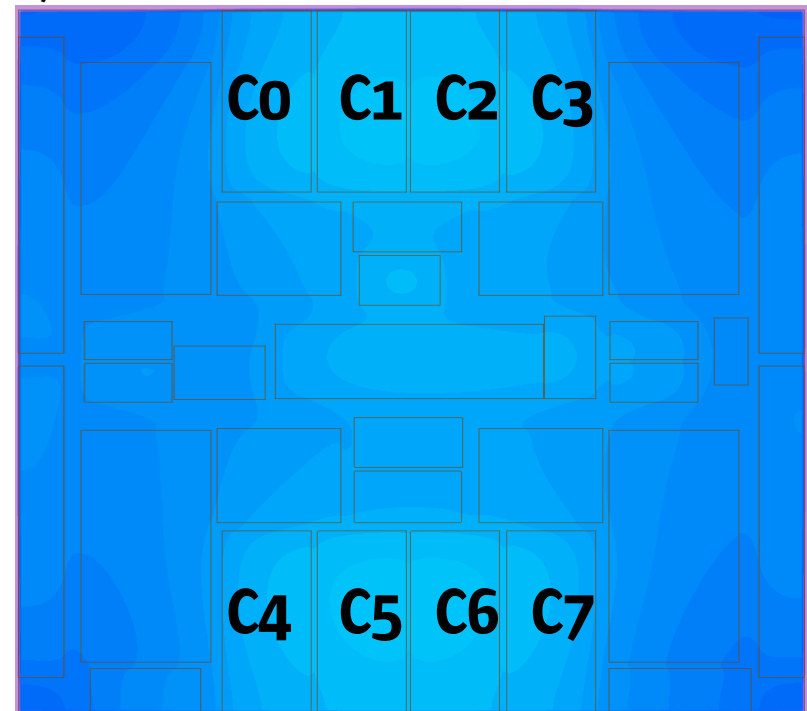


- Cool – improved performance, power & reliability
- Uniform – improved thermal management, clock distribution & reliability

(size not to scale)



Single-core Processor



CMT Processor

What Industry Analysts Say

“...the performance per watt is dramatically better than the industry standard... if Niagara delivers what Sun has said it will, I think it could be a real game changer.”

*Nathan
Brookwood
Analyst, Insight64*


“This announcement of yet another no-hassle SPARC system upgrade provides great asset protection to Sun’s huge base of applications and customers.”

*Gordon Haff
Senior Analyst, Illuminata*


Magnitude of Throughput Gains

Significantly Improves Real-World Application Performance

Web/App Throughput



Data Throughput



Our Customers Say it Best


Fiducia IT AG

On Energy Efficiency

“...we couldn't believe a 32 thread server could be so small...With the Sun Fire™ T2000 servers we'll be able to replace our current servers at a rate of **4:1** while doubling our performance. And...it looks like we'll be able to cut our energy use from **5.2 to .35 kw/hour**. That's a factor of 14. Amazing!”

*Matthias Schorer, Chief Architect
Fiducia IT AG*

The most Advanced OS on the Planet



Over \$500M, 3000 engineering years
600 new features, Guaranteed Binary
Compatibility

OpenSource, \$0 RTU license

“free as in beer”

DTrace – Diagnose production systems

Predictive Self Healing

Support for over 700 Platforms -
x86/x64/SPARC

5M+ Registered Users

14,000 OpenSolaris Community Members

Solaris Containers

One license - unlimited partitioning

Hundreds of applications on
one system

Ultimate consolidation tool

Oracle recognized licensing boundary


Restart in seconds

Improves Security

SPARC / Intel / AMD



Oracle 10g in Containers




- Complete light-weight virtual Solaris instance
 - > Requires hostname, IP, rootpath
 - > Independent user, system and process address spaces
 - > Created in minutes, rebooted in seconds
- No need to set system V semaphores in /etc/system
 - > Structures & arrays dynamically allocated
 - > Set Sys V IPC params on per-process or per-project basis w/o rebooting system
- Higher default values mean less tuning
 - > process.max-sem-ids (10 -> 128)
 - > process.max-sem-nsems (25 -> 512)
 - > project.max-shm-memory (0x800000 -> ¼ physical memory)

Who is Using Solaris 10 Today?

- 10 of the top 11 Retailers
- 5 of the top 5 Petro companies
- 5 of the top 5 Manufacturing companies
- 7 of the top 8 Financial Services companies
- 3 of the top 3 Shipping companies
- 5 of the top 5 IT companies
- 5 of the top 5 Insurance companies
- 5 of the top 5 Food/Beverage companies
- 6 of the top 6 Healthcare/Pharmaceutical companies

Solaris 10 OS License Growth



What's New in Solaris 10 6/06?

- Solaris ZFS
- Postgres for Solaris
- Networking Enhancements
- Fault Management support for AMD64
- Additional Enhancements
 - > Multimedia & Desktop Enhancements
 - > Additional x64/x86 Support
 - > Remote Technology Services

Oracle 10g on Niagara Performance

<some contents omitted>

iGenOLTP

- The iGenOLTP benchmark v2.1 is an OLTP benchmark based on a real customer workload. This workload simulates a **Global Order System** that includes four continent-based schemas all accessed via central views. This benchmark is easily portable on any database that supports the ANSI v2 standard and SQL views. For these tests we are using Oracle 10.2.0.1
- For this benchmark, we are using the schema for a single continent representing a database size of about 16 Gigabytes with 7.7M customers, 90,000 locations, 10,000 products & 77M orders
- Each JDBC transaction is executing the following :
 - (a) Obtain a random customer id number
 - (b) Perform query 1 - *Customer cost summary by product*
 - (c) Perform query 2 - *Customer zipcode statistics*
 - (d) Perform query 3 - *Delete activity*
 - (e) Perform query 4 - *Insert activity*
 - (f) Perform query 5 - *Complete customer summary*

Oracle Licensing

- An 8-core T2000 requires only a 2 CPU license
 - > Oracle uses a special .25 x multiplier for systems with T1 CPU
 - > http://www.oracle.com/corporate/pricing/multicore_faq.pdf
 - > Qualifies for SE license (when not using RAC)
- Oracle recognizes Solaris Resource Pools
 - > Effectively limits number of processors or cores available for use by Oracle
 - > Use to combine non-oracle workloads on same system while controlling license costs
 - > Example: On an 8-core T2000, create an Oracle resource pool with only 4 cores and bind multiple Oracle zones to that resource pool – result is only one SE/EE license is required

Customer Feedback

<http://sun.com/servers/coolthreads/t2000/ratings.jsp>

- The little box that could
- By John Verified Purchaser from Colorado on 6/7/2006
- Best Uses: Mass consolidation Server
- *We recently picked up two T2000's (8 core @ 1.2GHz, 32G of ram) for use as cluster nodes for our non-production Oracle databases.*
- *With the help of Solaris 10's new zones feature, we were able to consolidate 9 database environments onto this single system (for a total of 45 oracle instances). We then migrated the netbackup master server, a few nfs shares and just about anything else that we could find.*
- *All of this easily fit onto a single node. (During our testing, we wanted to make sure that in the event of a node failure we did not overload the remaining node)*
- *The throughput on this thing is absolutely amazing. We can easily saturate our FibreChannel links during the netbackup DSU destage operation to the tape drives, and the server is still extremely responsive and able to handle requests for the other applications.*
- *The only thing limiting further consolidation is the amount of RAM on the system. I would love to see it hold 64G of RAM.*

Want to learn more?

- **Get Solaris.**
 - > sun.com/solaris/get
- **Get Trained.**
 - > Free Training: sun.com/solaris/freetraining
 - > Solaris Learning Paths: sun.com/training/solaris
- **Get Datasheets & Whitepapers.**
 - > sun.com/solaris/reference_materials
 - > sun.com/servers/coolthreads/
- **Get started with Solaris How-To Guides.**
 - > sun.com/solaris/teachme