Corporate Update

Peter Ungaro
President & CEO

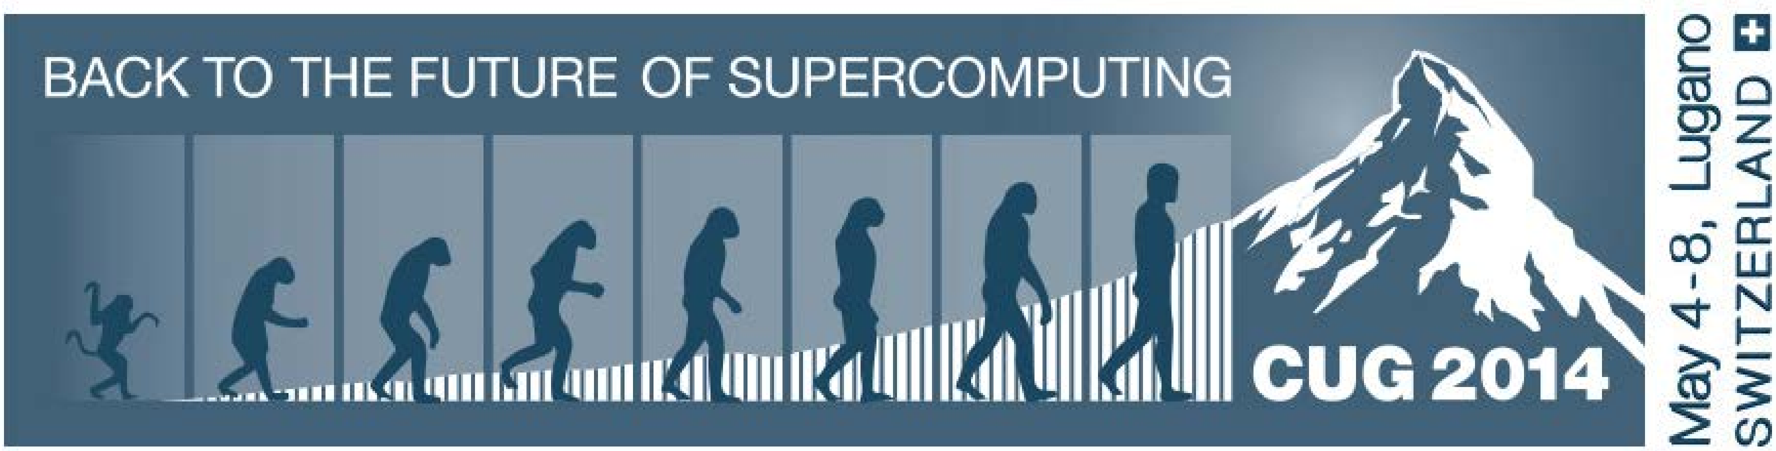

Thanks!

A Special Thanks to our Host!

CSCS

Swiss National Supercomputing Centre

… And CUG Sponsors

New CUG Members

PGS

ECMWF

DWD

STATE

HLRN

養 和
HKSH

CU NY
The City University of New York

DoD SUPERCOMPUTING RESOURCE CENTER
NAVY DSRC

My First 10 Years at Cray…

One thing that hasn't changed…

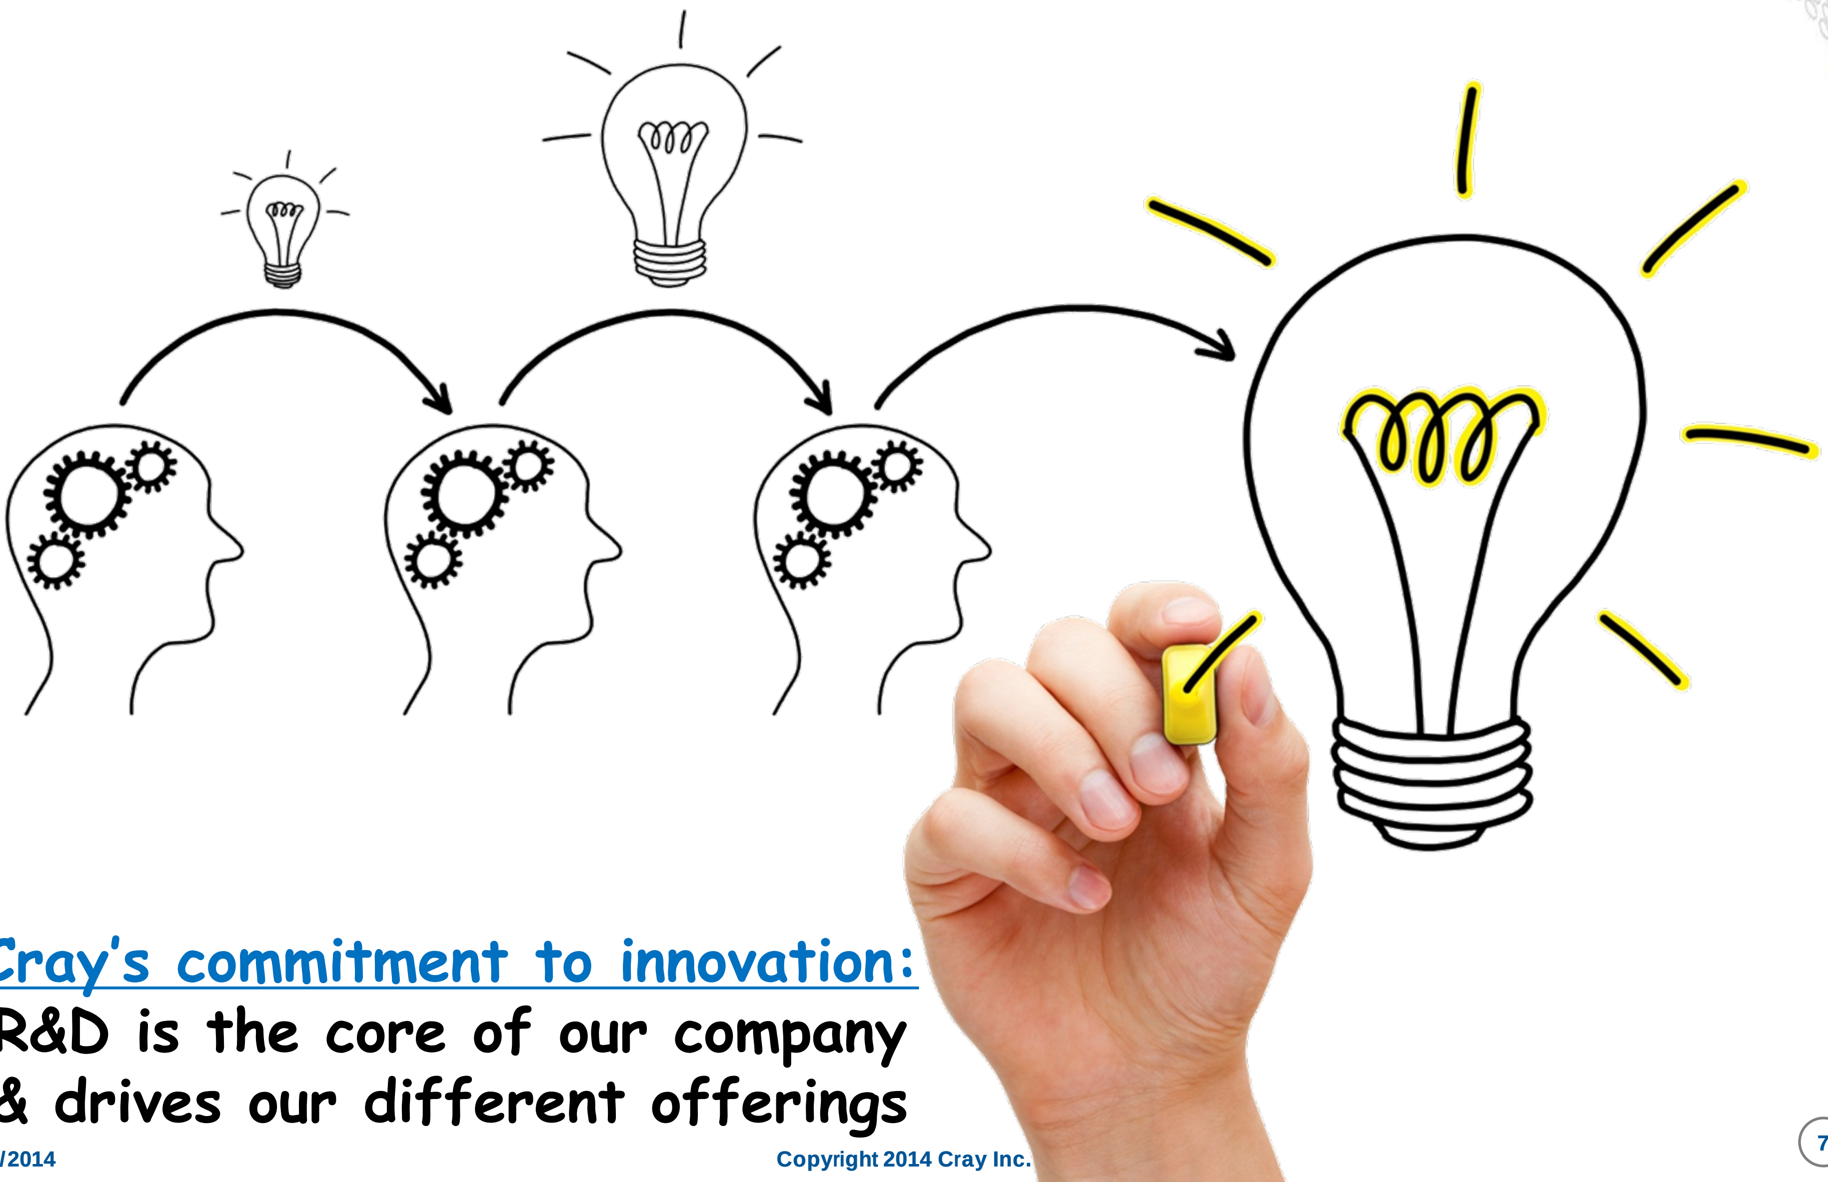

Cray's commitment to innovation:

R&D is the core of our company & drives our different offerings

Adaptive Supercomputing Innovations

System Interconnect

Systems Management &
Performance Software

Packaging

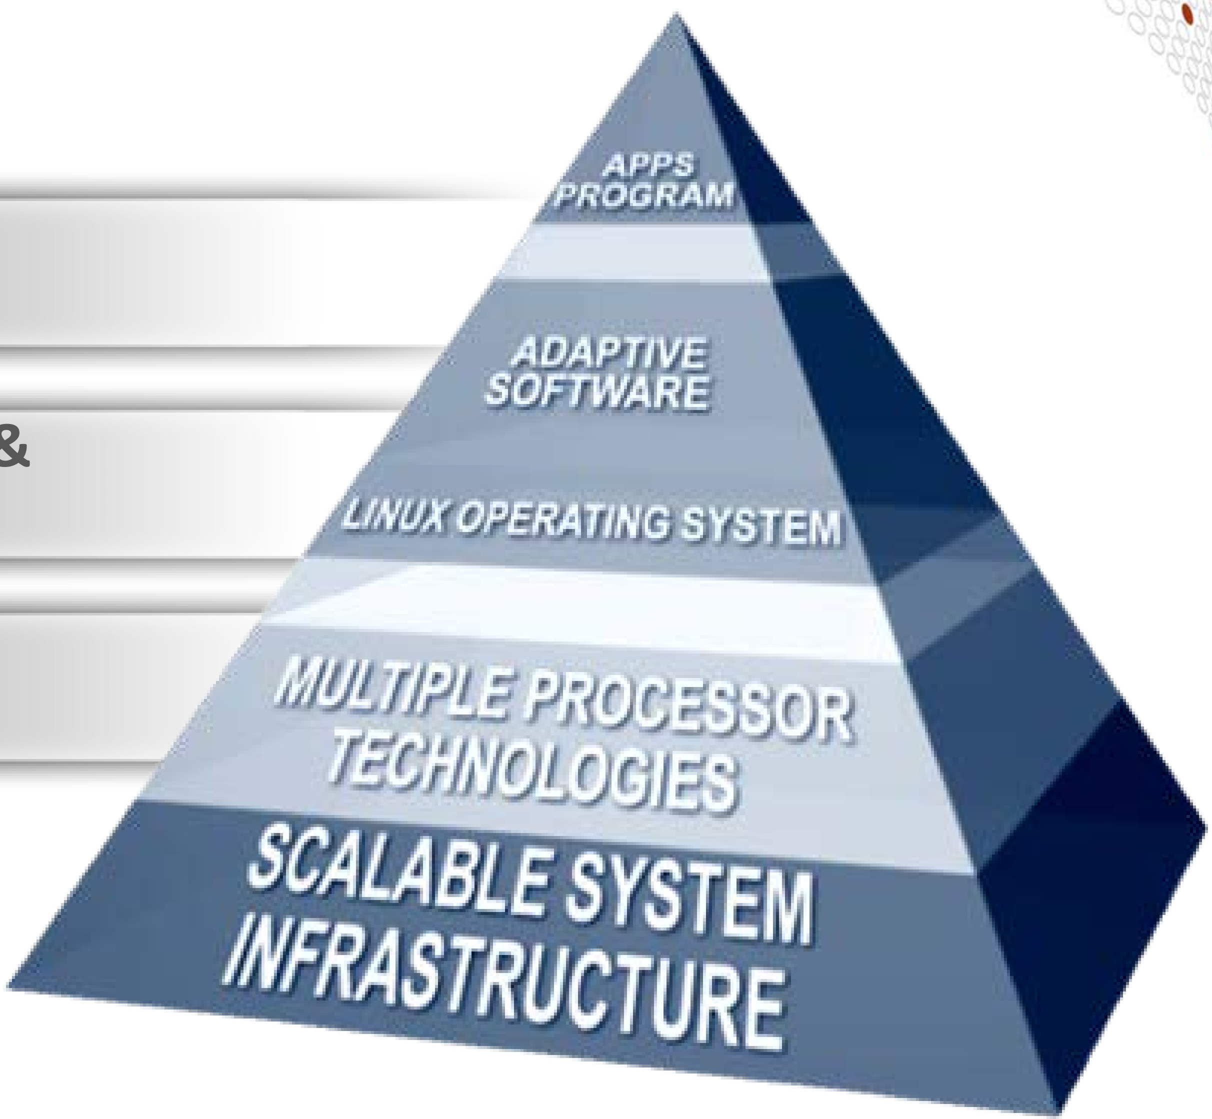

Cray's Vision:
The Fusion of Supercomputing and Big & Fast Data

Modeling The World
Cray Supercomputers solving "grand challenges" in science, engineering and analytics

Data Models
Integration of datasets for search, analysis, predictive modeling and knowledge discovery

Data-Intensive Processing
High throughput event processing & data capture from sensors, data feeds and instruments

Math Models
Modeling and simulation to provide the highest fidelity virtual reality results

Compute | **Store** | **Analyze**

Supercomputing Big Data

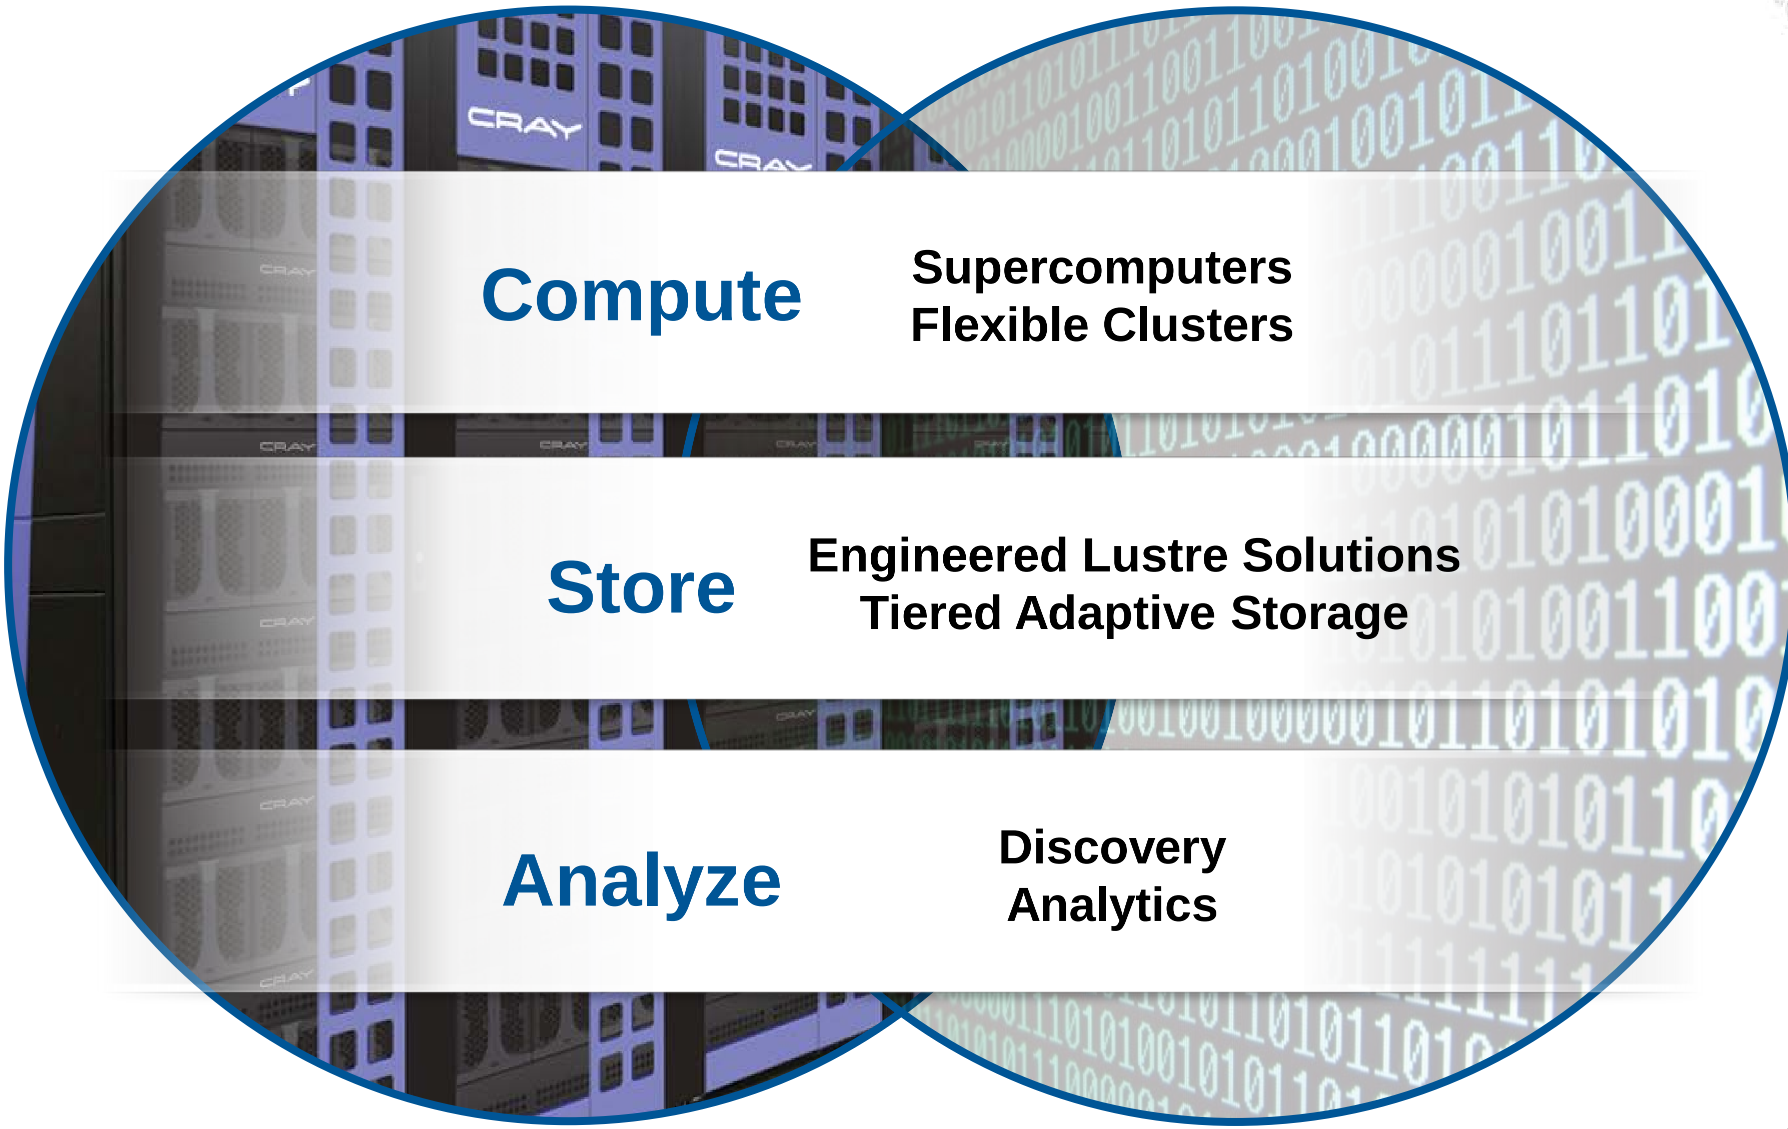

2X in 2


CRAY
1955 Olson Drive

15%

One Last Baby Picture…

Welcome Cori!!!

Integrated HPC Environments are the capability that will turn data into insight and discovery

Our Vision…

Build a world-class integrated supercomputing environment that enables transformational computing across a broad set of science, engineering and advanced analytics (big data) applications

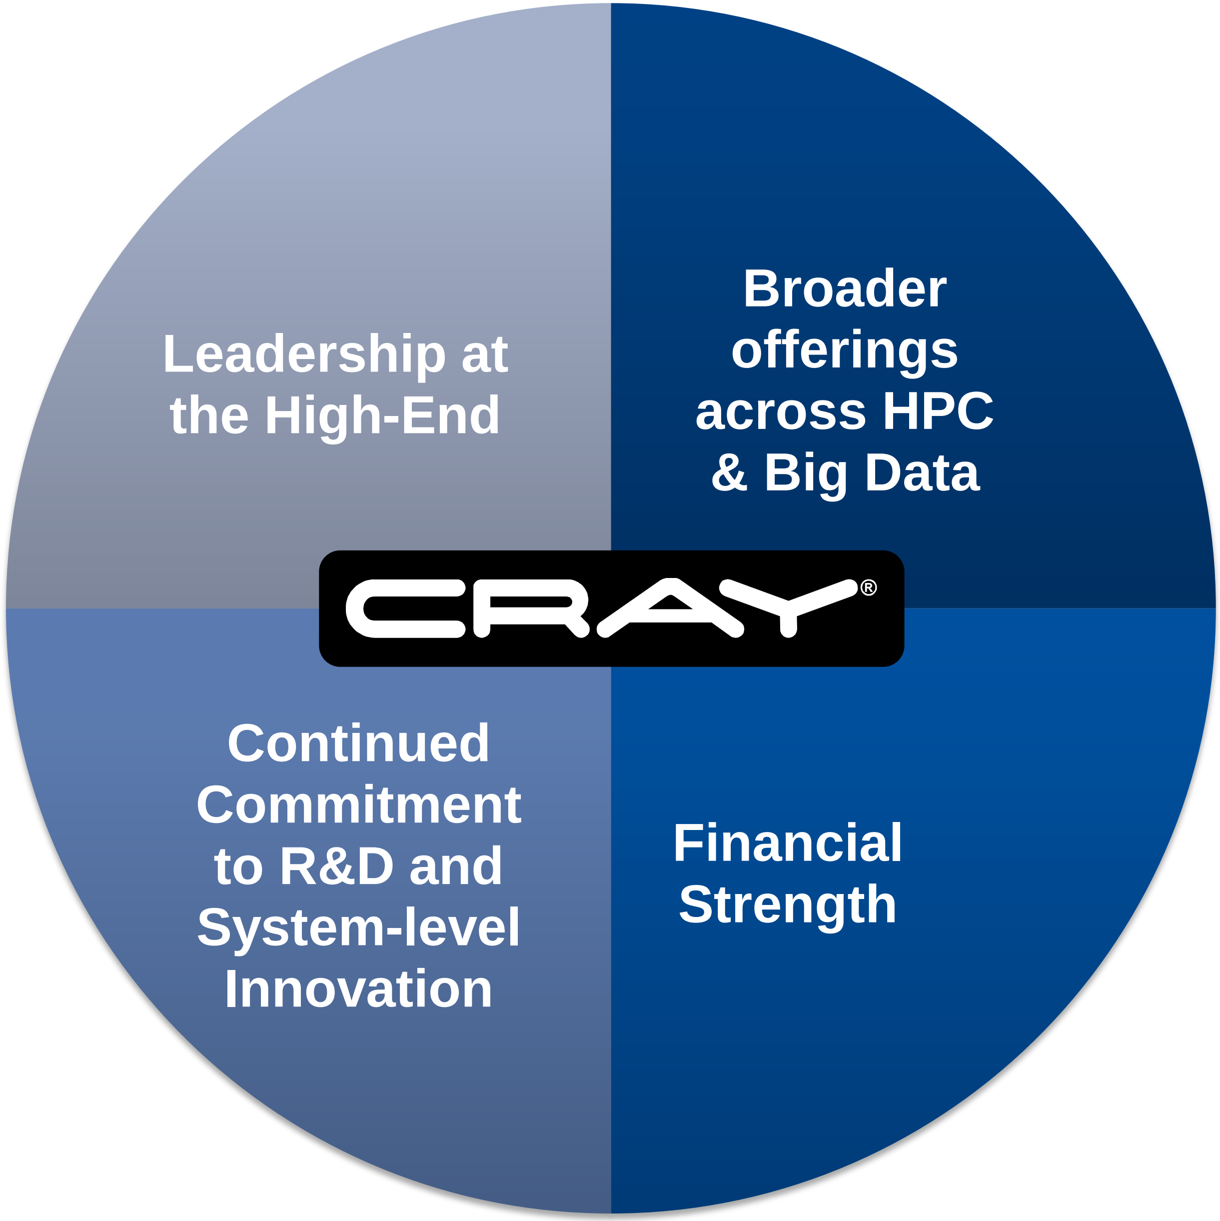
Leadership at the High-End
Broader offerings across HPC & Big Data
CRAY®
Continued Commitment to R&D and System-level Innovation
Financial Strength

Key Technology Trends

Bill Blake
Chief Technology Officer

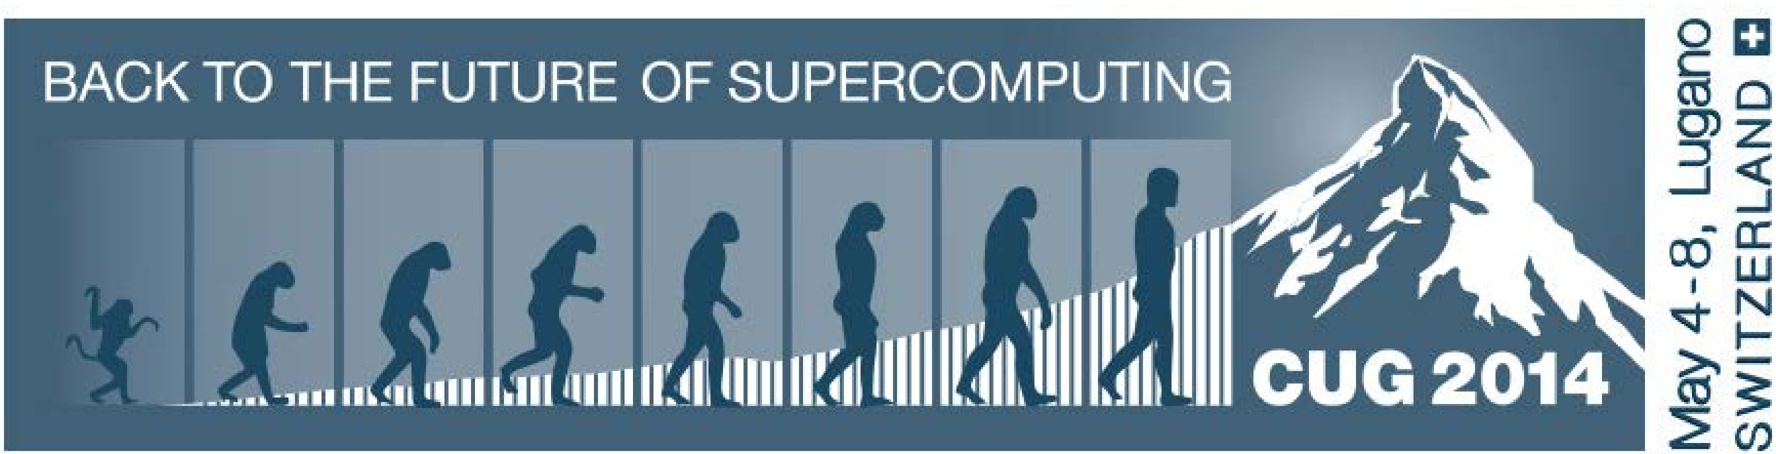

The Fusion of Supercomputing with Large Scale Data Analytics

Key Trends

- The cost per FLOP is dropping to near zero!
- The application needs for memory size and IO performance are insatiable!
- New levels of memory and storage hierarchy are emerging
- The MPP interconnect must support atomic Global Memory Operations
- Future systems are being influenced by the converged infrastructure of compute & storage
- Converged infrastructure demands management of O(10000) nodes for Exascale or Hyperscale

The Challenge of Exascale

- **The U.S. Government House Science Committee, May 2013:** ***America's Next Generation Supercomputer: The Exascale Challenge***
  - We must reduce power consumption by at least a factor of 50.
  - We must increase the parallelism of our applications software and operating systems by at least a factor of 1,000.
  - We must develop new programming methods to increase dramatically the number of programmers that can develop parallel programs.
  - We must improve memory performance and cost by a factor of 100.
  - We must improve systems reliability by at least a factor of 10.

The Cost per FLOP is Dropping Like a Spent Rocket!

0.000000001 cents per FLOP

TALE OF THE TAPE:
SUPERCOMPUTER
VS. GAME CONSOLE

	SANDIA LAB'S ASCI RED	SONY PLAYSTATION 3
DATE OF ORIGIN	1997	2006
PEAK PERFORMANCE	1.8 teraflops	1.8 teraflops*
PHYSICAL SIZE	150 square meters	0.08 square meter
POWER CONSUMPTION	800 000 watts	<200 watts

* For GPU; CPU adds another 0.2 teraflops

SANDIA LAB'S RED STORM XT	NVIDIA-powered Game Console
2005	2019
40 teraflops	40 teraflops[2]
280 square meters	0.08 square meter
1 000 000 watts	<200 watts

Maximum Scalability: System Node Specialization of Red Storm

Key to Cray's MPP scalability is system node o/s specialization combined with the Cray's high bandwidth, low latency interconnect

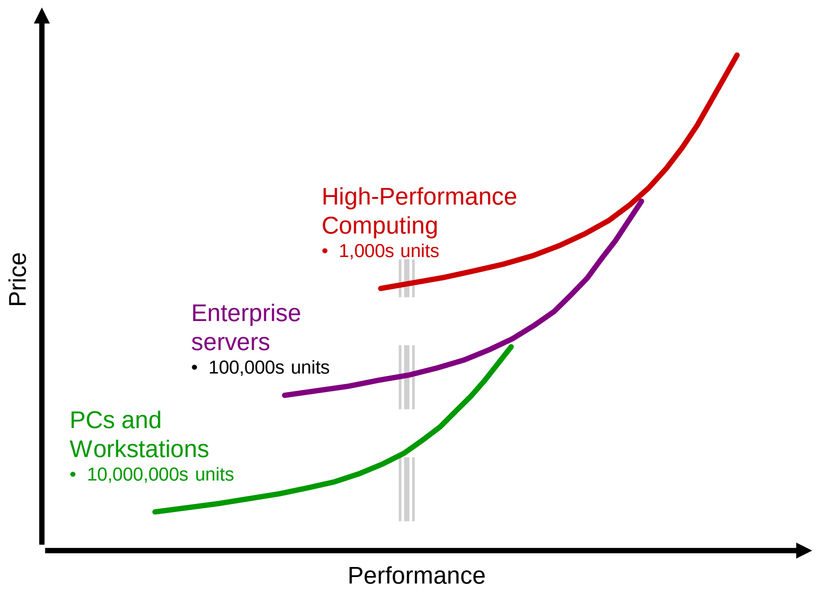

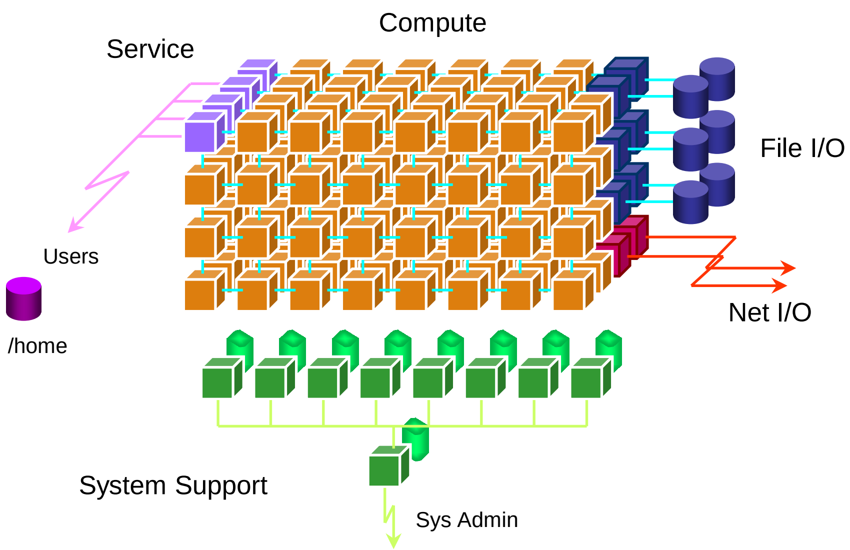

Courtesy of Dr. Bill Camp, Sandia National Laboratories circa 2005

Adapt the System to the Application not the Application to the System

Cray's Adaptive Supercomputing combines multiple processing architectures into a single scalable system—CPU, GPU, or Multi-threaded

Our focus in on the user's application where the adaptive software, the compiler or query processor, knows what types of processors are available on the heterogeneous system and targets code to the most appropriate processor

The next step is to evolve Adaptive Supercomputing to Big Data workloads

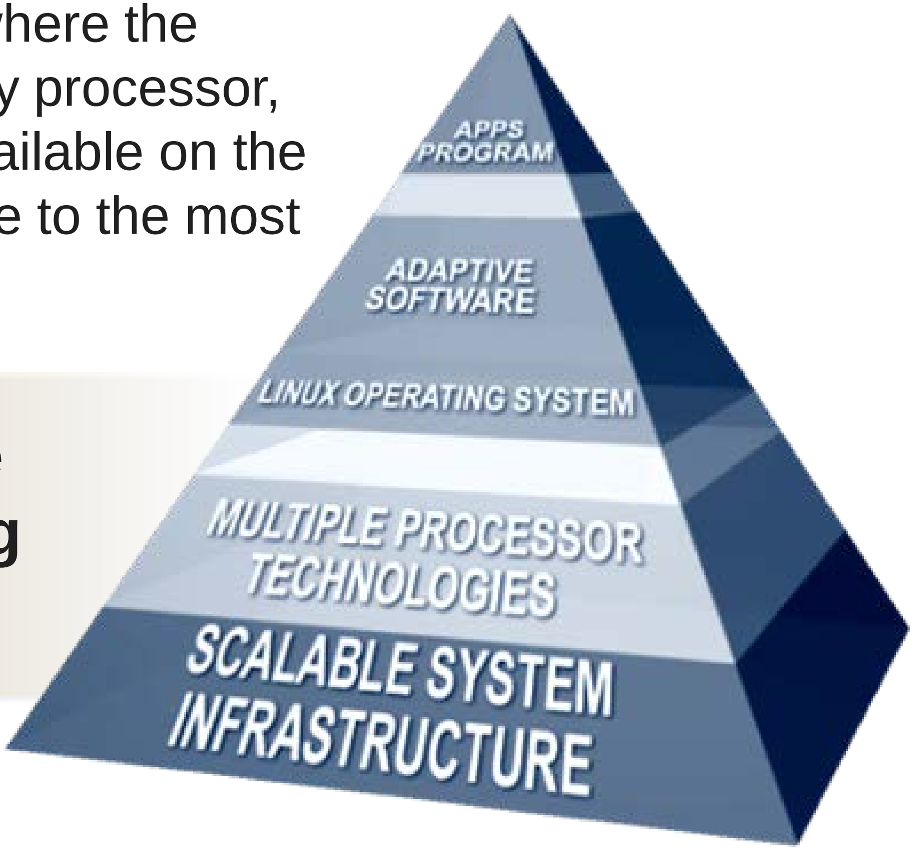

Motivation For "Cascade" XC30

Why are HPC machines unproductive?

- **Difficult to *write* parallel code (for example, MPI)**
  - Major burden for computational languages
- **Lack of programming tools to *understand* program behavior**
  - Conventional models break with scale and complexity
- **Time spent trying to modify code to fit *machine* characteristics**
  - For example, clustered machines have relatively low bandwidth between processors, and cannot directly access global memory
  - Programmers then try hard to reduce communication, resorting to bundling communication up in messages instead of just accessing shared memory

Cray's XC30 system and tools provide the needed help!

- The Aries network provides hardware assist for MPI operations and atomic Global Memory Operations
- The entire programming tool kit is optimized for parallel programming with runtime analysis allowing best library/kernel to be used dynamically
- Continuing R&D to establish new auto-tuning/optimization approaches

System Architecture Differences…

Exascale Supercomputing

- Scalable computing w/high BW, low-latency, Global Mem Architectures
- Highly integrated processor-memory-interconnect & network storage
- Ability to apply all compute power to one highly parallel application
- Low data movement – load the “mesh” into memory and compute
- Move data for loading, defensive check-pointing or archiving
- “Basketball court sized” system that consume <20 MWatt

Hyperscale Computing – aka the Cloud

- Distributed computing at largest scale
- Divide-and-conquer approaches on Service Oriented Architectures
- Ability to apply compute power to many apps with multi-tenancy
- High data movement-- Scan/Sort/Stream all the data all the time
- Lowest cost processor-memory-interconnect & local storage
- “Warehouse sized” systems that collectively consume >260 MWatt

Need to Create a “Virtuous Cycle”

Cloud provides new distributed programming models that utilize “divide and conquer” approaches with massive scale-out Service Oriented Architectures using local storage and low cost hardware, and **new data analytics algorithms where data scientists claim “the larger the data the simpler the algorithm”**

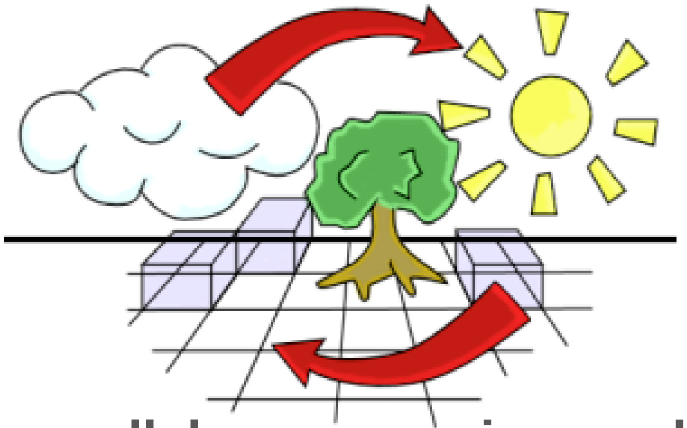

HPC provides new parallel programming models that utilize highly scalable Global Memory Architectures supported by highest BW, lowest latency interconnects, with powerful algorithms for high fidelity modeling and simulation using highly iterative processing of both capability and capacity workloads that increasingly support data assimilation (from sensors)

Image courtesy of University of Michigan –Atmospherics Dynamics Modeling Group

The "Big Data" Challenge

Supercomputing minimizes data movement – "data movement" is highly restricted for defensive or resiliency such as loading, check pointing or archiving.
Programming model is imperative (C++/Fortran + MPI) with focus on the details of **how** parallel programming is done

Data-intensive computing is all about data movement - scanning, sorting, streaming and aggregating *all the data all the time* to get the answer or discover new knowledge from unstructured or structured data sources.
Programming model is declarative (query) or functional with emphasis on ***what*** is being computed versus ***how*** it is computed

Cloud Computing is all about virtualization -- Application access to converged infrastructure (Compute/Network/Storage) via IP Stack
Programming Model is Platform as a Service with APIs for ***what*** is being computed rather then ***where*** the computing is done

Multiple Aspects of Big Data

Business Reporting	• Transaction Analytics (OLAP): • **ad hoc SQL queries on structured data** in relational databases by ***Analysts*** producing Business Intelligence Reports • Looking at all the data, O(100) TB, all the time
Search & Correlation	• Textual Analytics (Hadoop Ecosystem) • **API for analysis of unstructured data** in massive data sets by ***Programmers*** seeking “long tail” insights • Looking at all the data, O(1000) TB, once at a time
Discovery	• Graph Analytics (Semantic Web → Warehouse) • **ad hoc SPARQL queries on graph data using new standards for open linked data RDF, OWL, etc** • By ***Analysts*** seeking discovery via hypothesis • Looking at all the data, O(100) TB, and relationships

YarcData: Purpose-built for the Unique Challenges of Data Discovery

Do not know the relationships In the data → Shared Memory Model — Access all data with uniform, low latency regardless of partitioning, layout or access pattern

Do not know the desired insight or the right question to ask → Memory Accelerator — Investigate multiple, changing hypotheses in parallel without prefetching/caching

Do not know the paths/linkages to explore diverse data sets → In Memory, Graph Analytical Database — Explore diverse data fused without upfront modeling and independent of linkage/traversal path

Sandia National Laboratories

	# PROCESSORS	TIME
Traditional Approaches after months of optimization	48	10.8 Hours
YarcData	32	30 sec

1,944 Times Faster !

Institute for Systems Biology

"In the amount of time it takes to validate one hypothesis, we can now validate 1000 hypotheses – increasing our success rate significantly." – Dr. Ilya Shmulevich

It is Really About Decision Making through Fact Finding and Equation Solving

Analyst Query

Key Function	Language	Data Approach	"Airline" Example
OLTP	Declarative (SQL)	Structured (relational)	ATM transactions Buying a seat on an airplane
OLAP Ad Hoc	Declarative (SQL+UDF) or NoSQL	Structured (relational)	Business Intelligence analysis of bookings for new ad placements or discounting policy

The "Old World" of Fortran and SQL

Key Function	Language	Data Approach	"Airline" Example
Optimize Models	Procedural (Solver Libs)	Optimization <-> Simulation	Complex Scheduling Estimating empty seats
Simulate Models	Procedural (Fortran, C++)	Matrix Math (Systems of Eq's)	Mathematical Modeling and simulation (design airplane)

Languages & Tools for Programmers

Big Data → Fast Data

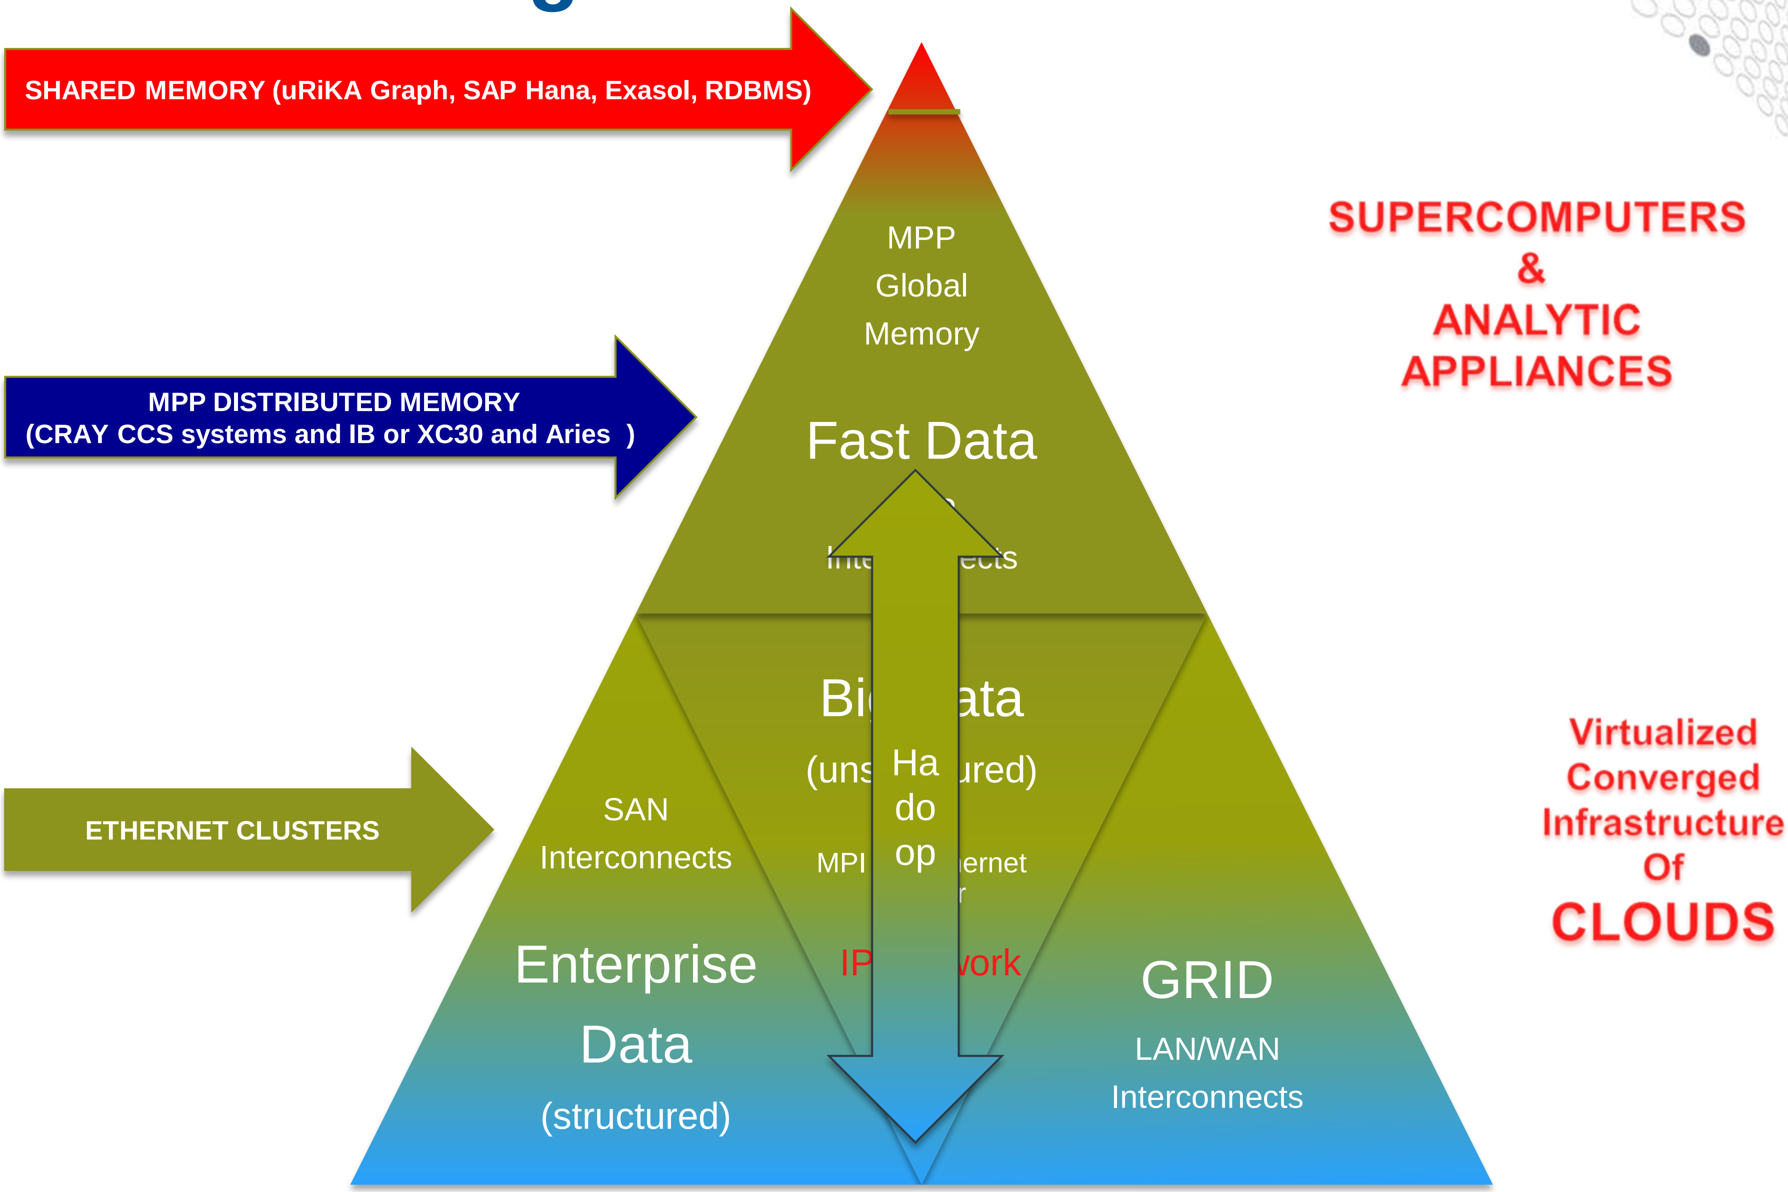

Cray Brings Supercomputing to Analytics

Highest Performance

Best Scalability

Tightest System Integration

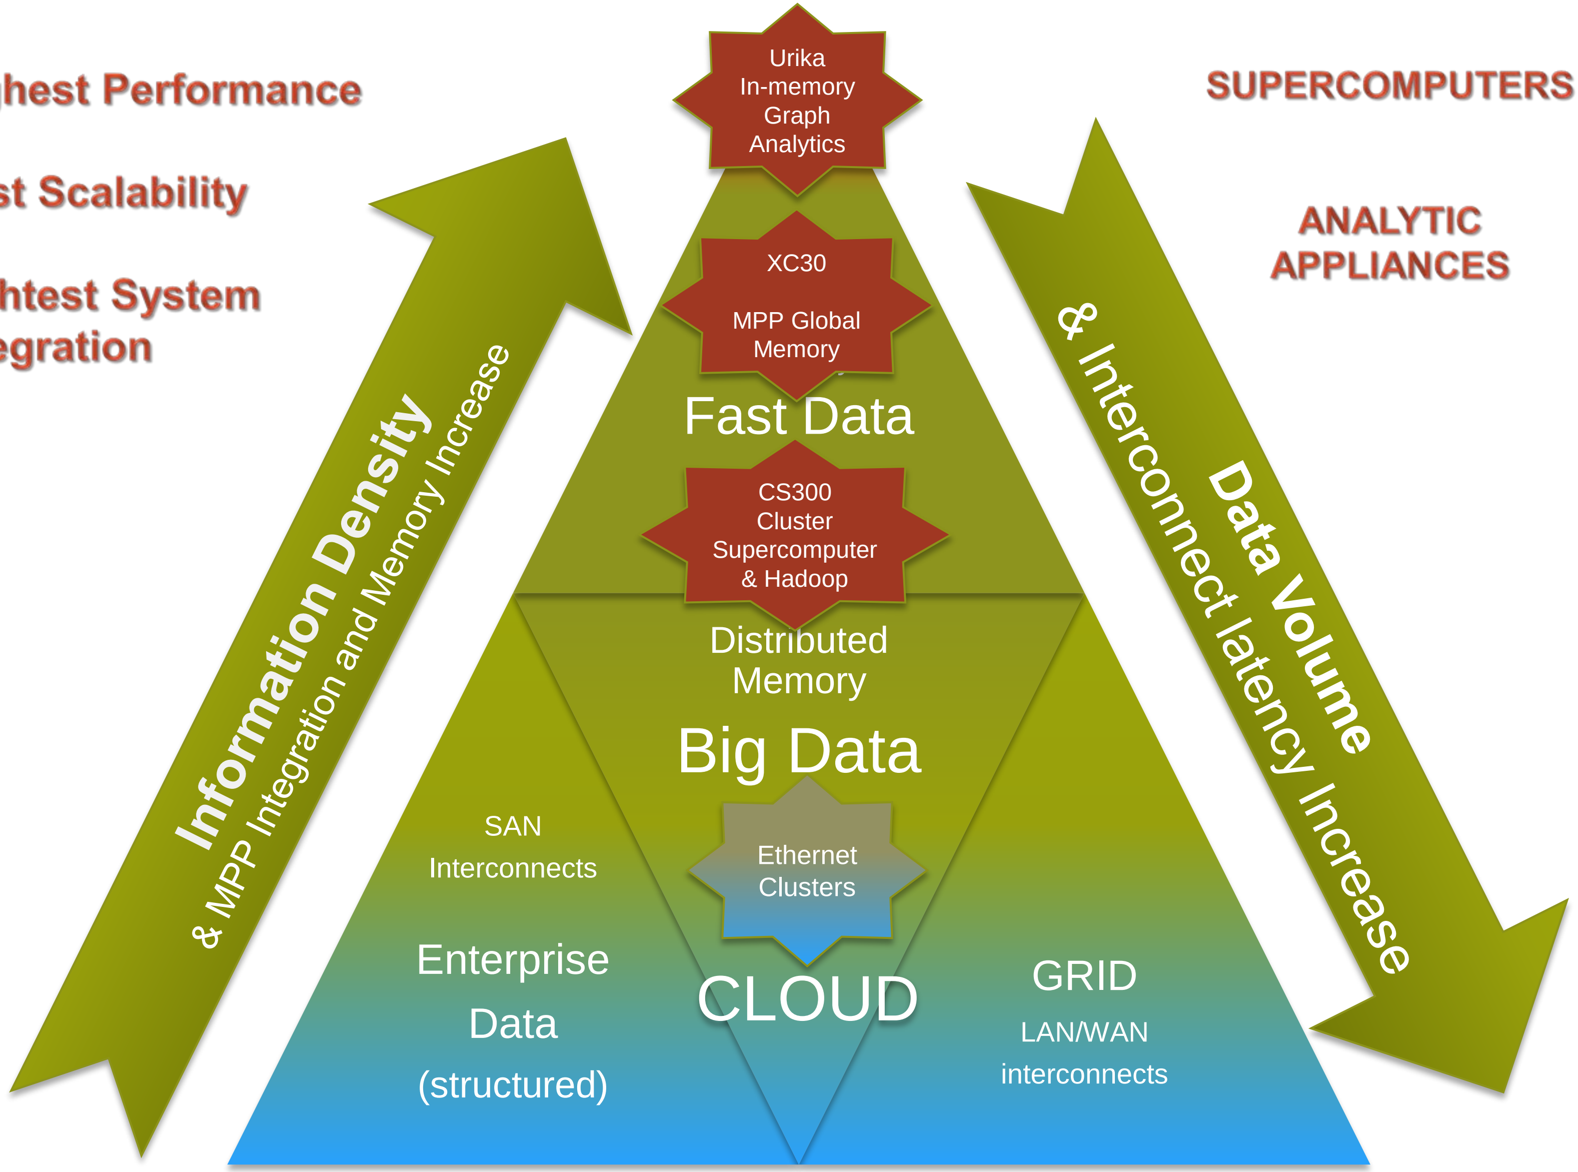

Programming Emphasis

"what" plus "how"

Highly optimized C++/Fortran
Thread/Task + Msg Parallelism
Auto-tuned libraries and
Tools that expose low level

"what" without "how"

Functional and Declarative Languages

"what" without "how" or "where"

PaaS

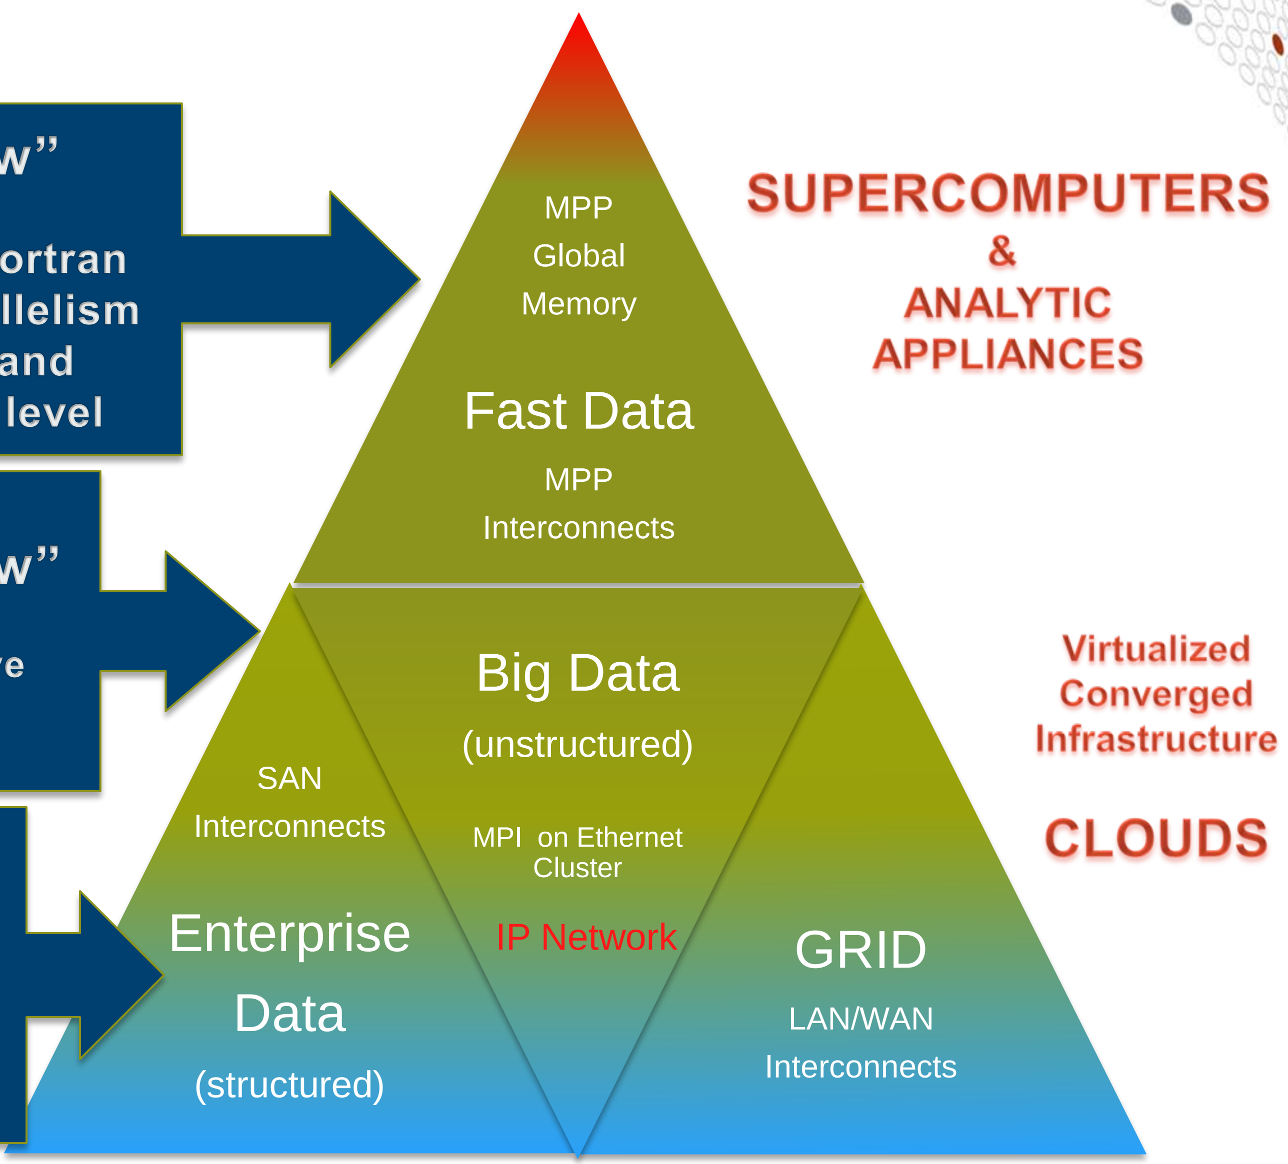

The Future: Global Memory and Latency Hiding

Data Reuse Near Previous Data Access

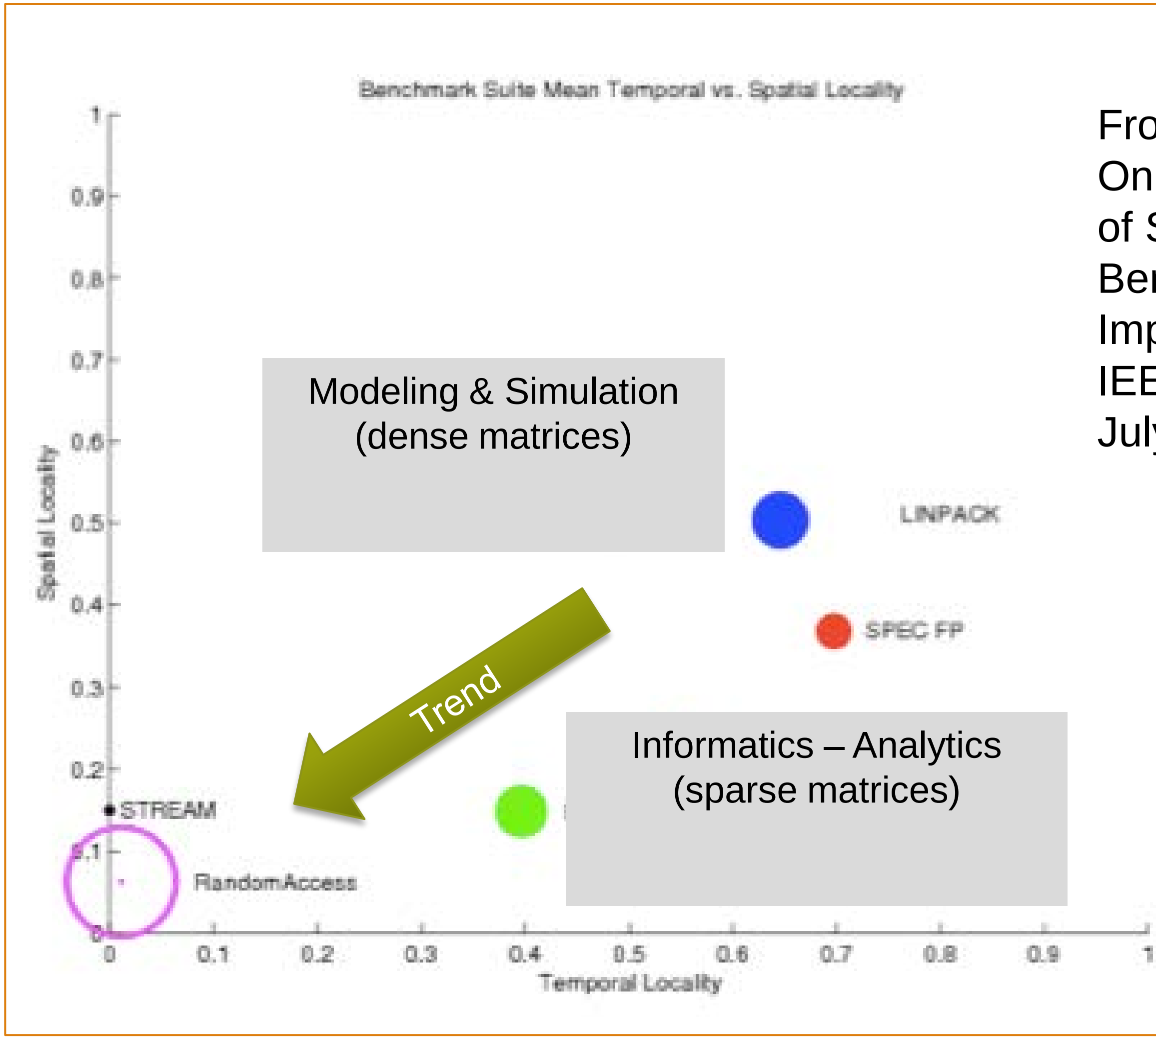

Data Reuse over Time

From: Murphy and Kogge,
On the Memory Access Patterns of Supercomputer Applications: Benchmark Selection and Its Implications,
IEEE Trans. On Computers,
July 2007

The Challenge of Expressing Analytics?
Queries and Program Code Do Not Mix Well

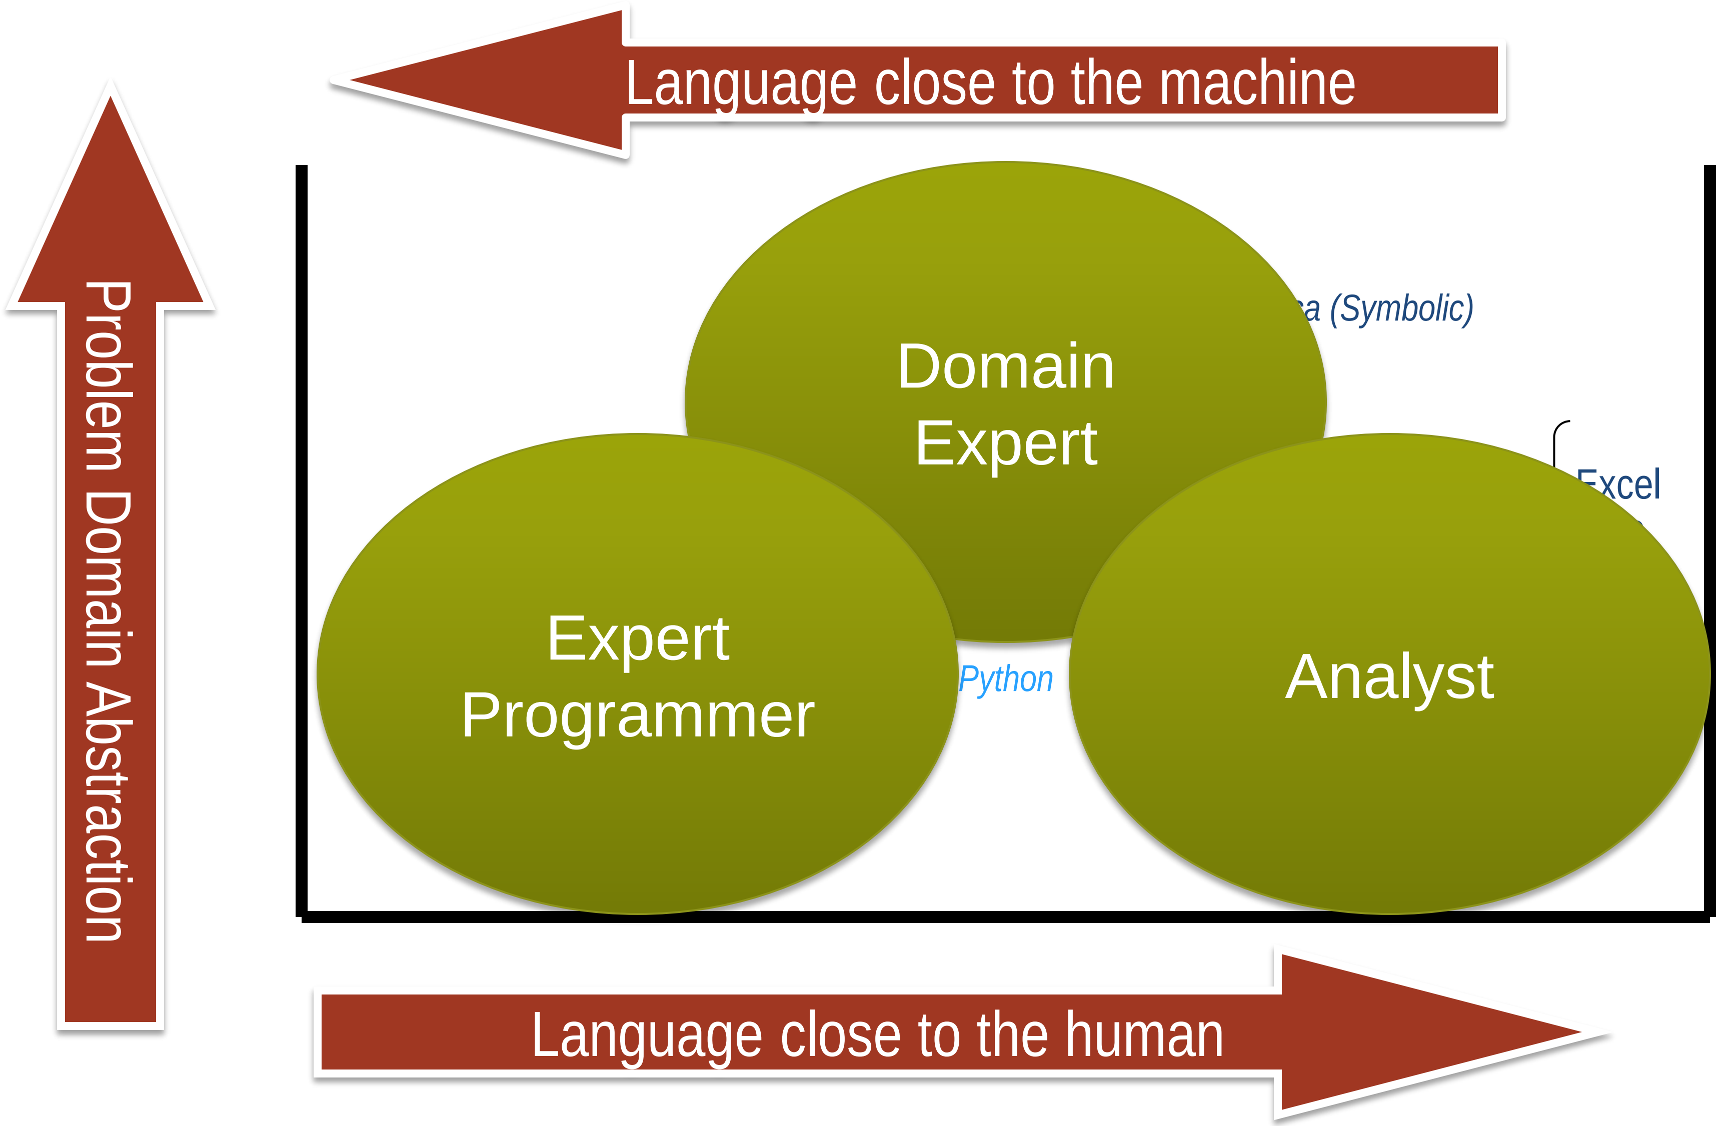

Adapting to Data-Intensive Computing: Adding Value at the Edge of the Network

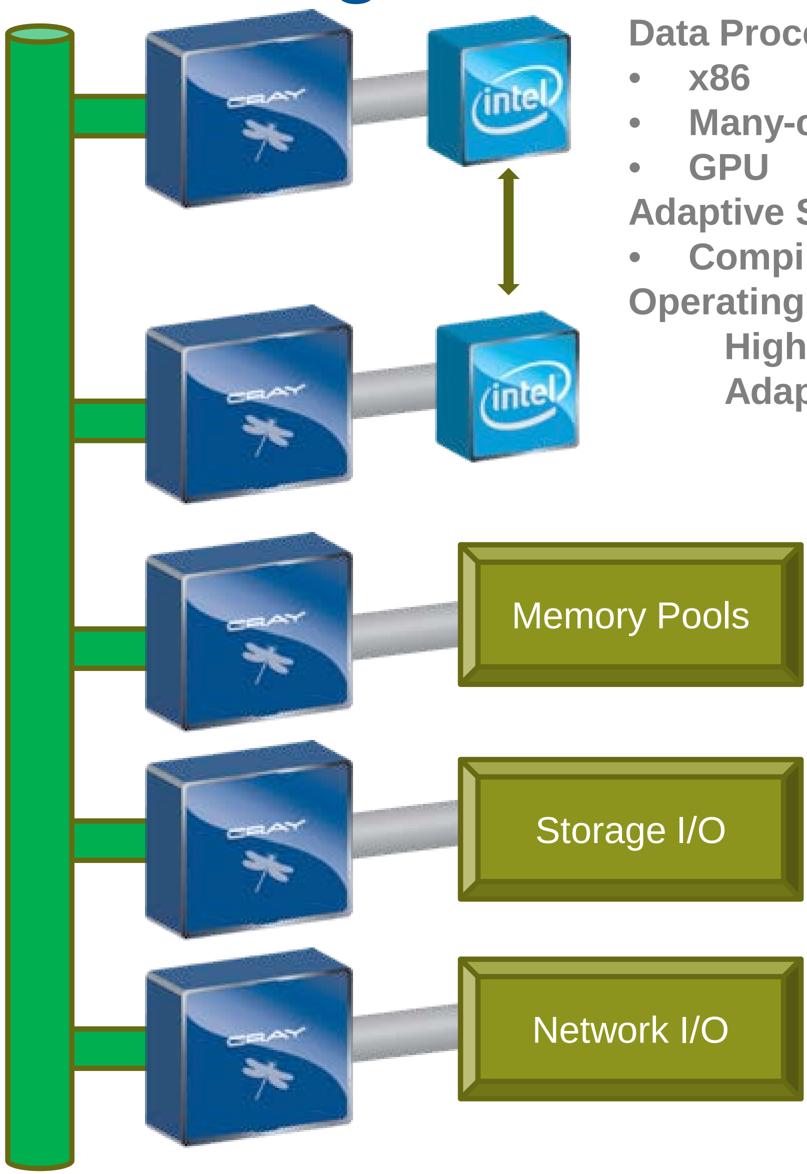

Data Processors
- x86
- Many-core
- GPU

Adaptive Supercomputing
- Compilers, Auto-tuning

Operating System

High throughput scheduling,
Adaptive Runtime, Network

High Performance Memory
- DRAM
- Active NVRAM

Storage
- SAS to PCIe
- Tightly coupled Software Stack for RAID
- Very Low Latency

Intelligent I/O
- Protocol Load Balancing
- Specialized Data Ingest
- Policy Based Data Movement

  Software Defined Network

The most capable interconnects will be key to analytic workloads

Future Architectural Possibilities

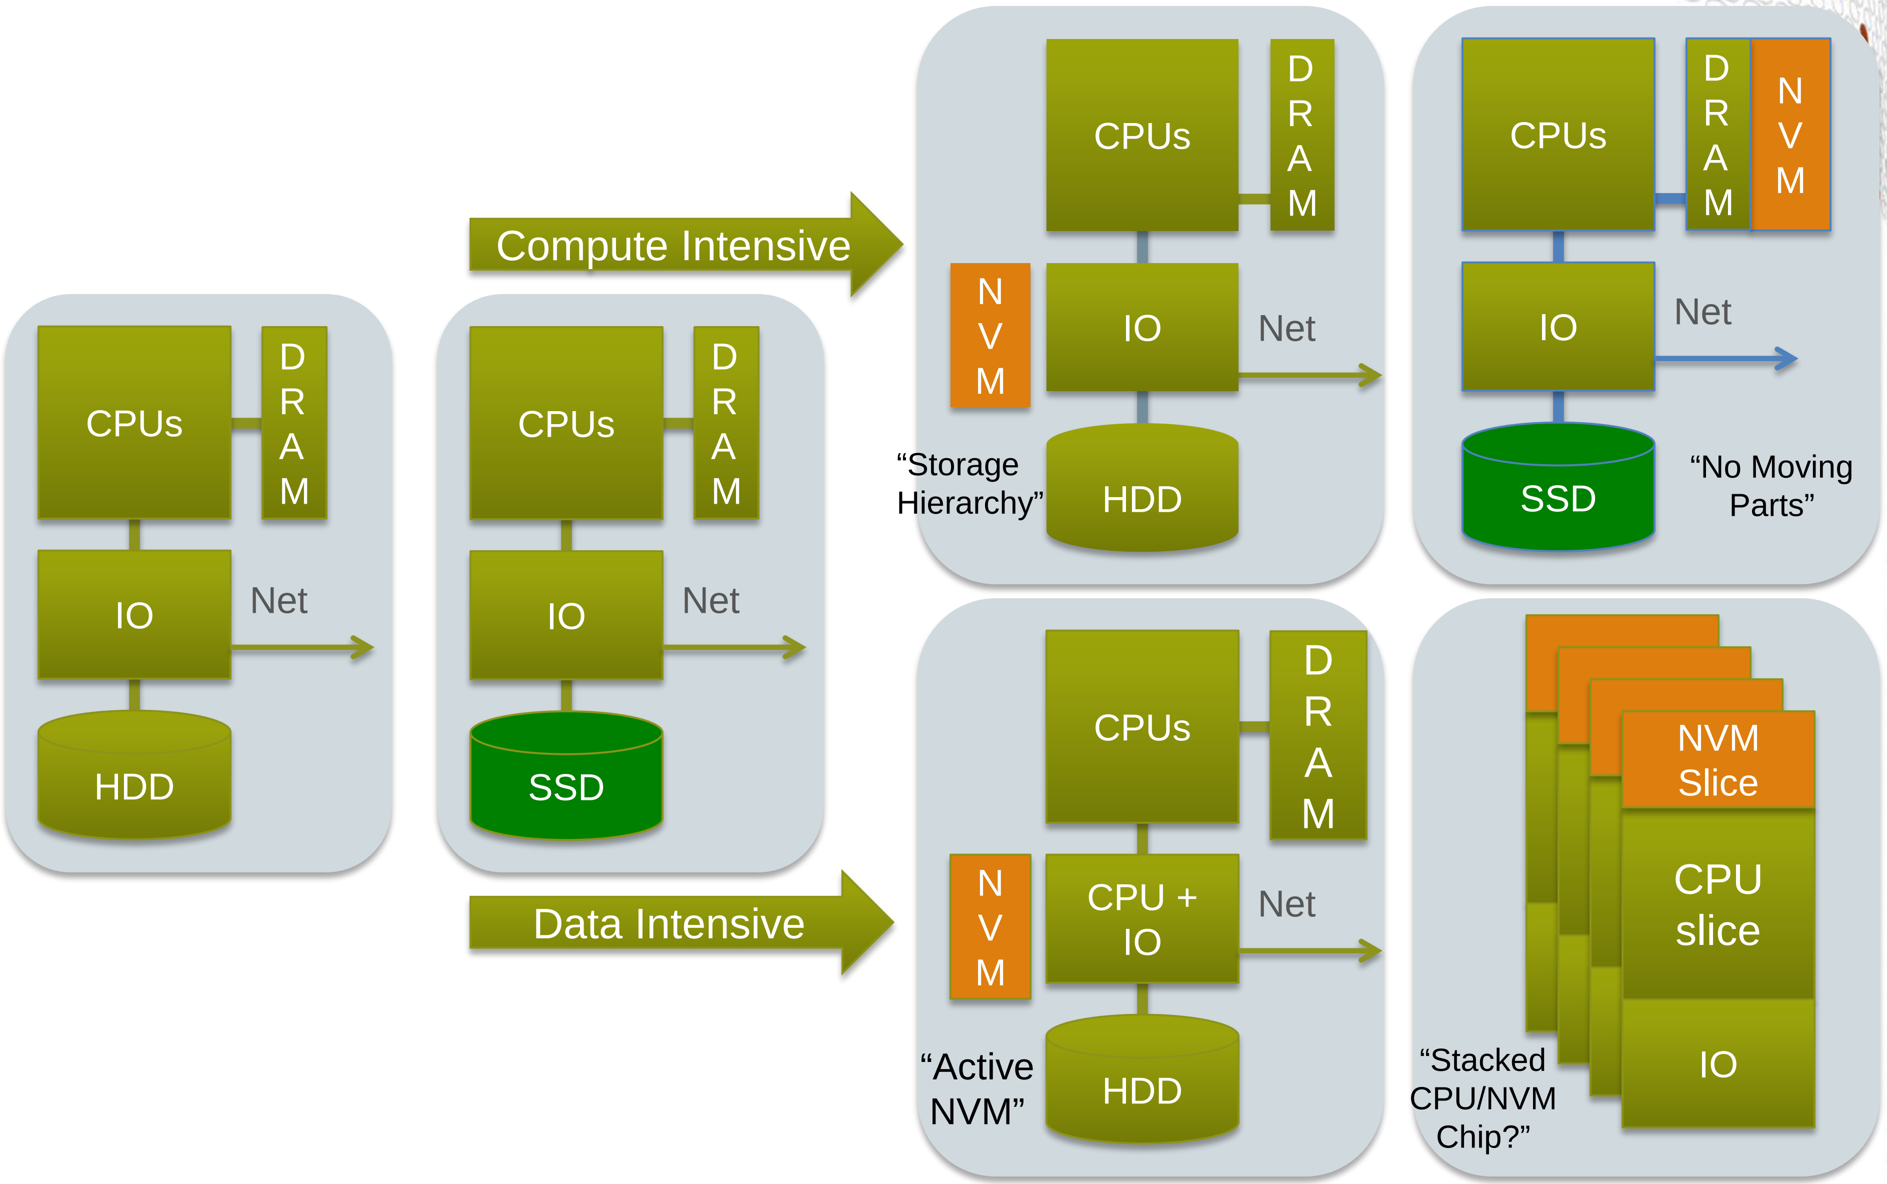

The Next Generation of Analytics Will Need Very High Performance Global Memory Operations

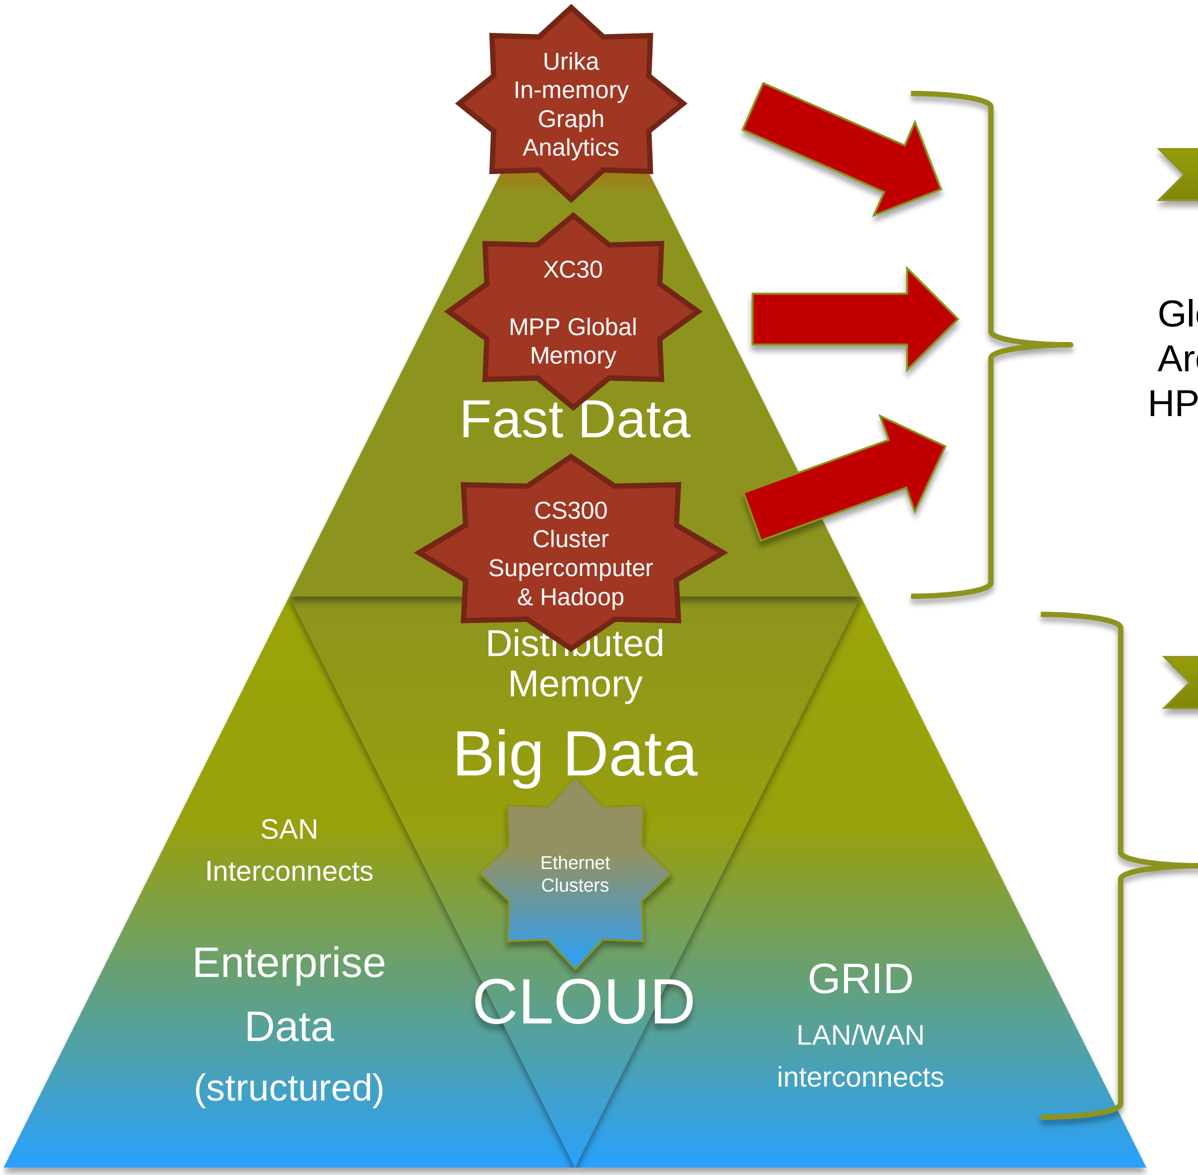

+ 5 Years	+ 10 Years

Capability Workloads

MPP Global Memory Architecture for HPC & Analytics	MPP Global Memory Architecture for Analytics with Exascale Technology and NVM

Capacity Workloads

Ingest and Archive

High Throughput,
Multi-tenancy
Cloud Computing

"The SW Defined Datacenter"

Cray's Roadmap "Fusion"

Modeling & Simulation "Compute"

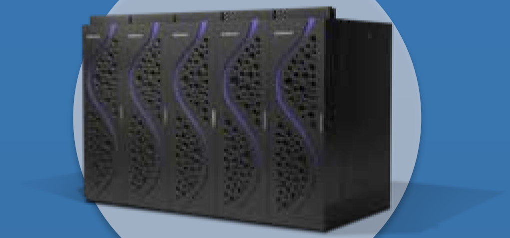

Data Ingest "Store"

Appliances "Analyze"

Multiple Roadmap Steps:

Combine workflows & data into a single system
Aggressive local and global memory capabilities
Tightly integrated software stack (& runtimes) across all 3 capabilities

Thank You!

bill.blake@cray.com

Cray Products Update

Barry Bolding
Daniel Kim
Peg Williams
Ramesh Menon

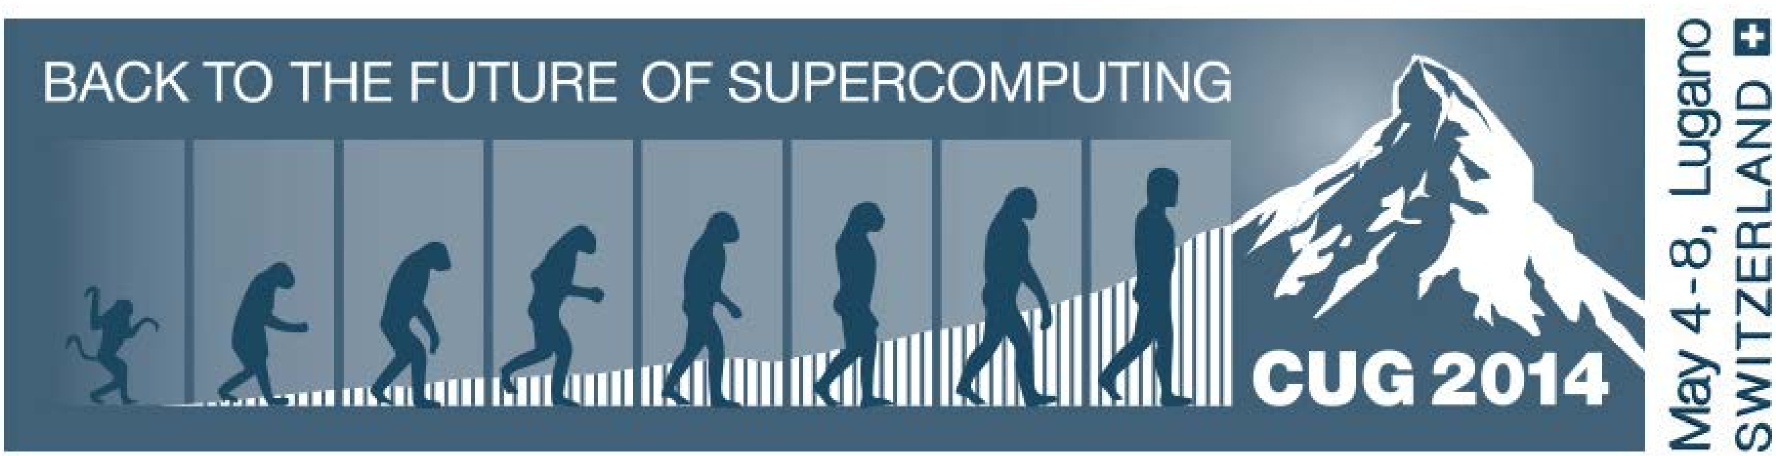

Big Data Workflows Driving Research

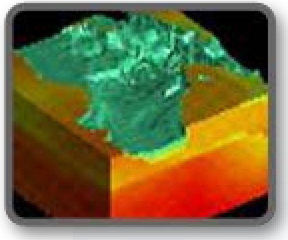

Earth Sciences
CLIMATE CHANGE
& WEATHER PREDICTION.
REMOTE SENSING.

Life Sciences
DRUG DISCOVERY, GENOMIC
RESEARCH, COMPLEX
MODELING

Defense & National Security
WARFIGHTER SUPPORT,
THREAT PREDICTION &
STOCKPILE STEWARDSHIP

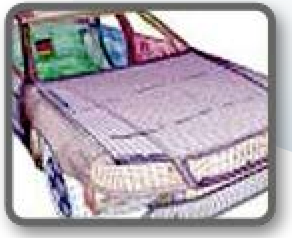

Manufacturing
AIRCRAFT DESIGN,
CRASH SIMULATION &
FLUID DYNAMICS

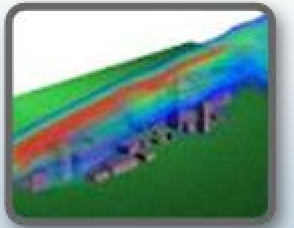

Energy
SEISMIC IMAGING &
RESERVOIR SIMULATION

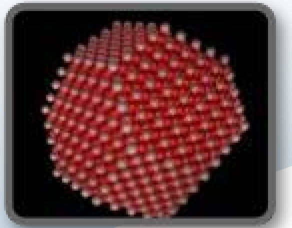

Higher Education
UNIVERSITY-DRIVEN SCIENCE,
NEW ENERGY SOURCES &
EFFICIENT COMBUSTION

CRAY SONEXION

Cray Datacenter Solutions

CRAY TIERED ADAPTIVE STORAGE

Cray XC30

Yarcdata Urika®
Hadoop Solutions

Cray CS300

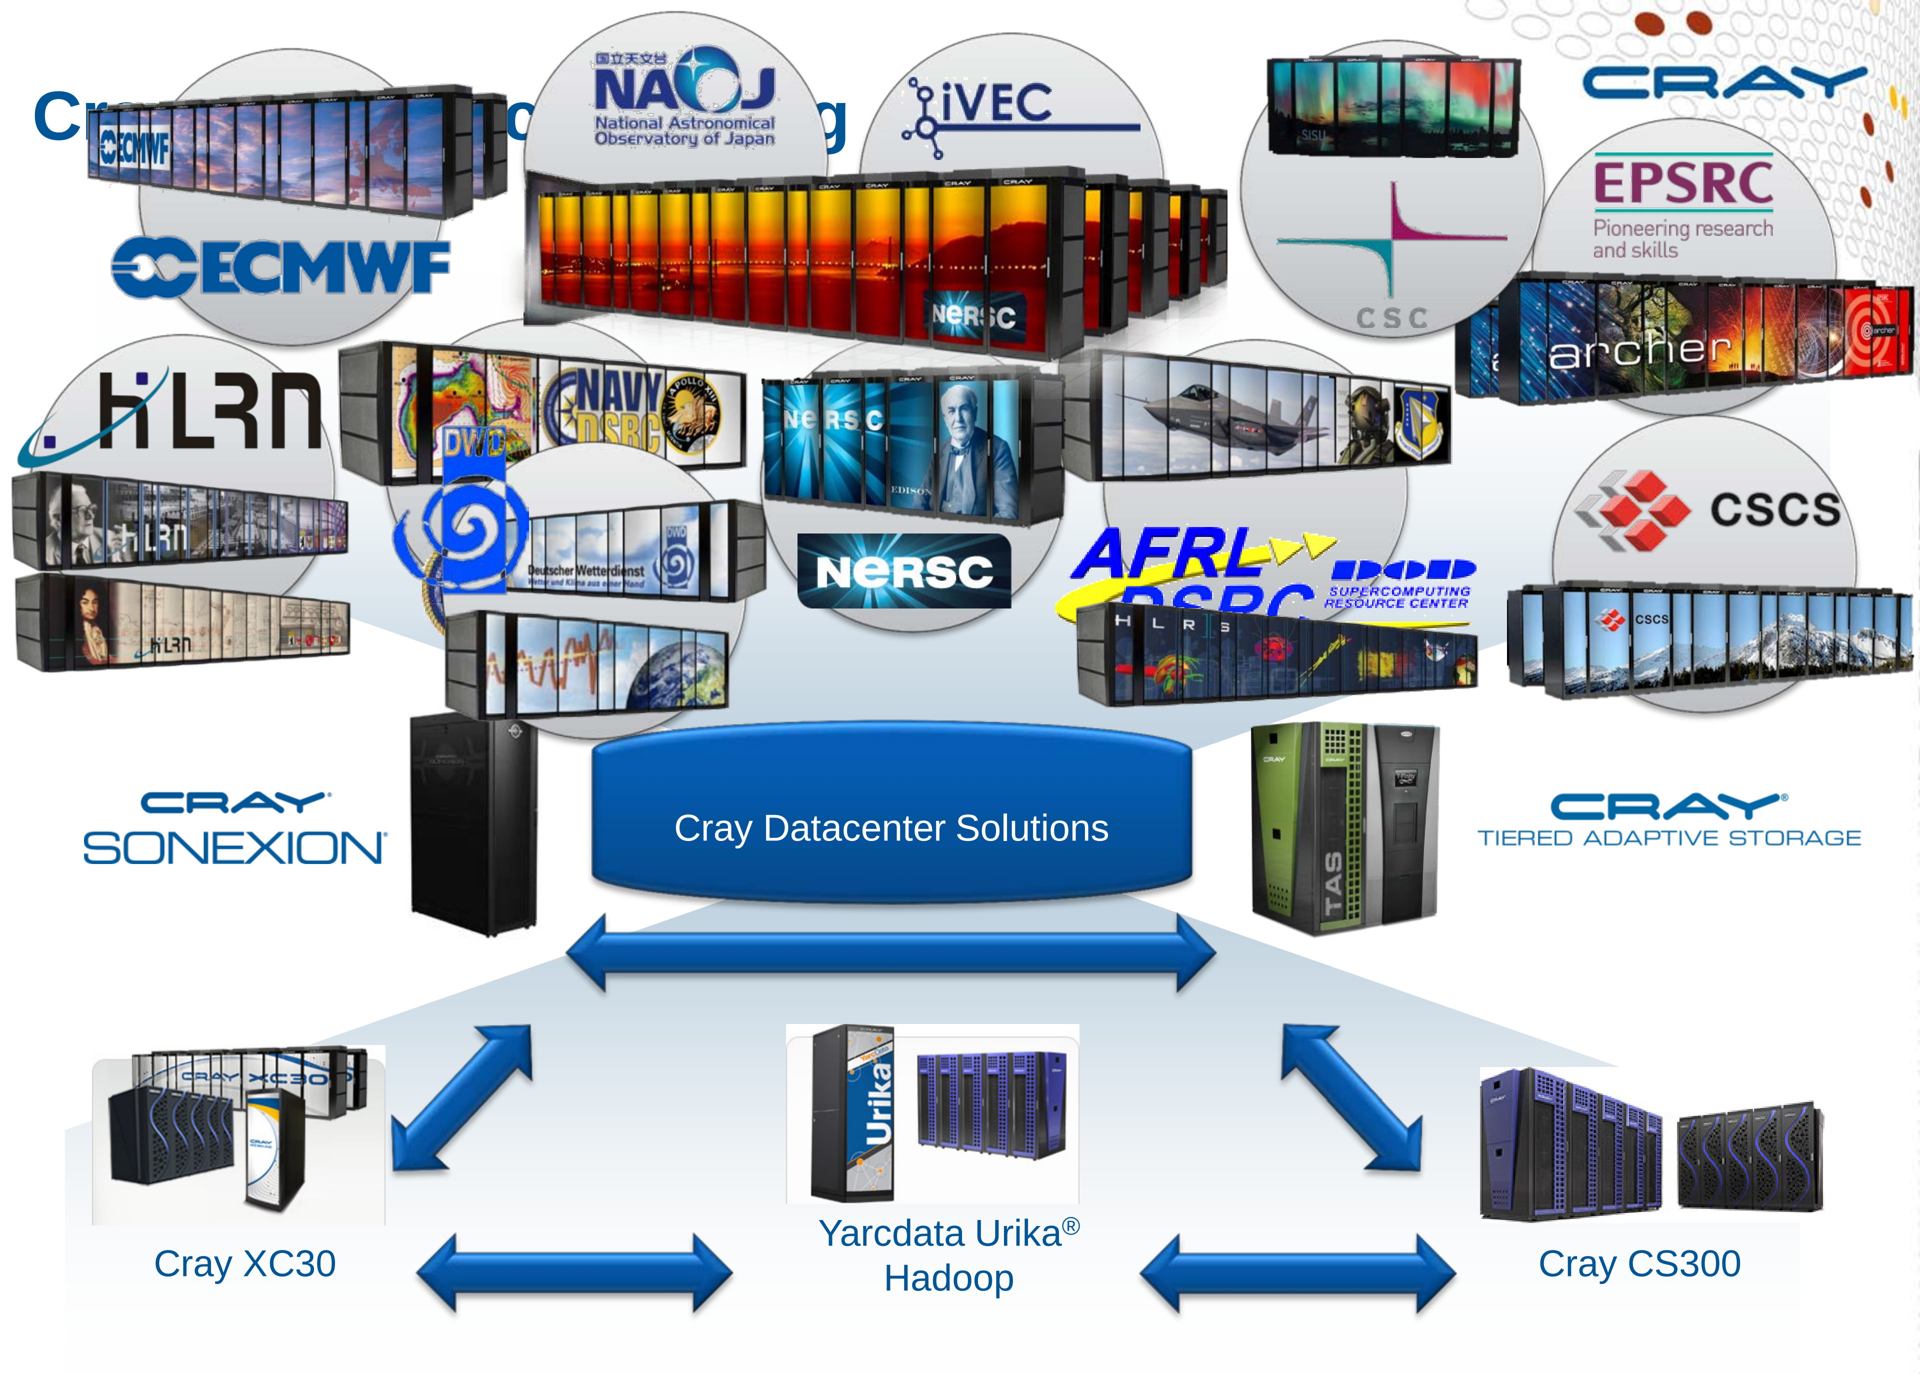
CRAY
ECMWF
NAOJ
National Astronomical Observatory of Japan
iVEC
CSC
EPSRC
Pioneering research and skills
archer
HLRN
NERSC
Deutscher Wetterdienst
CSCS
AFRL
SUPERCOMPUTING RESOURCE CENTER
CRAY SONEXION
Cray Datacenter Solutions
CRAY TIERED ADAPTIVE STORAGE
TAS
Urika
Cray XC30
Yarcdata Urika®
Hadoop
Cray CS300


CRAY
CRAY
Supercomputing Leadership
Lawrence Livermore National Laboratory
筑波大学
University of Tsukuba
KYOTO UNIVERSITY FOUNDED 1897
NOAA
HOPPER
NERSC
KRAKEN
Los Alamos NATIONAL LABORATORY EST. 1943
Sandia National Laboratories
CHUGACH
HLRS
OAK RIDGE National Laboratory
BLUE WATERS
NSF
NCSA
ILLINOIS
INDIANA UNIVERSITY

The Cray XC30 Family

Leadership Class Supercomputing

Cray's Leadership Class Supercomputing

Sustained Application Performance and Scalability	Production Supercomputing	Investment Protection – Upgradable by Design	User Productivity

Years of experience deploying extremely large and successful systems…and delivering breakthrough science & engineering

Scientific Discovery on NCSA's "Blue Waters"

- **HIV Capsid – NAMD**
  - Enable virus-busting therapies
- **Space Weather – VPIC**
  - Magnetic reconnection of high temperature plasmas
- **Stellar Process Study – PPM**
  - Inertial confinement fusion
- **Numerous 1PF Applications in production**

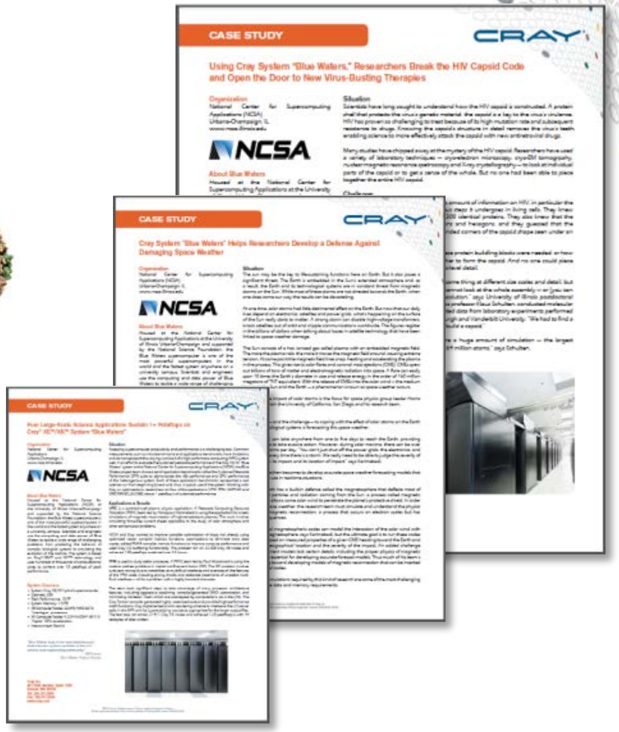

Cray is all about the Science on the Products, more than the products themselves!

Scientific Discovery on NERSC's – "Edison"

- Big bang simulation
- Combustion research
- Dark matter investigations
- High resolution Xray imaging

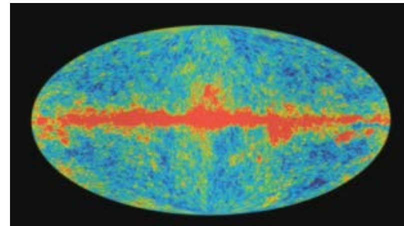

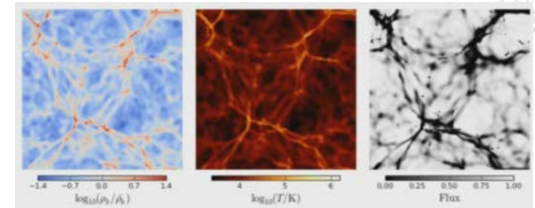

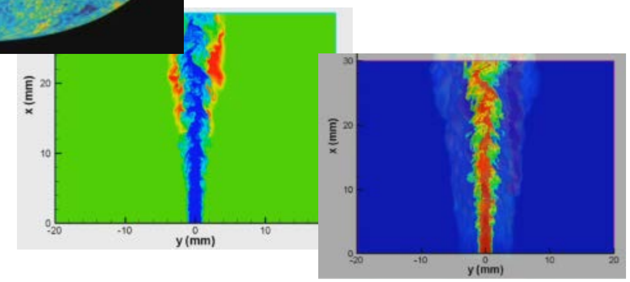

Edison XC30 Supercomputer

Hopper XE6 Supercomputer

Sonexion Storage System

Mendel CS300 Cluster

The Cray CS300 Family

Leadership Cluster Supercomputing

Cray CS300 Series
Cluster Supercomputers

Air-Cooled

Liquid-Cooled

- **Flexible and Scalable**
- **Reliable**

- **Manageable**
- **End-To-End Solutions**

Only from Cray

Advanced Cluster Engine (ACE) Management Software

- Cray's Scalable Cluster Management Software
- Framework for Cray's Cluster HPC Software Stack
- Customizable design to work with open source and commercial applications and development tools
- Deployed in some of the largest production clusters in the world

Cray Programming Environment (CPE)

- Cray's Optimized C, C++ & Fortran Compilers
- Cray's Scientific & Math Libraries
    - LibSci, Libsci_acc
- Cray's Perftools

Unique offerings to Solve Challenging Problems

Cray CS300 Cluster Deployments

Tsukuba U, “COMA” 1PF Phi

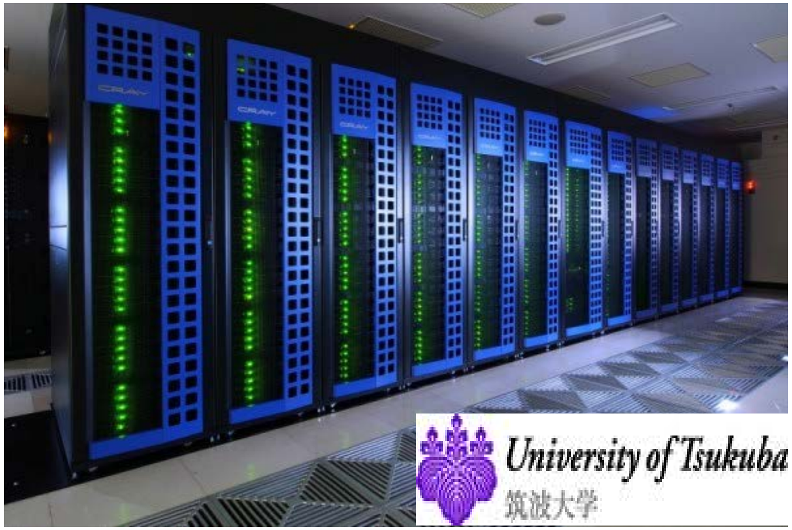

Kyoto U, CS300 - LC

Mississippi State University

The new High Performance Computing *Collaboratory* (HPC²) liquid-cooled CS300-LC™ cluster supercomputer, nicknamed “Shadow”.

"This investment is the latest example of Mississippi State's commitment to providing powerful, energy-efficient and technologically-advanced HPC system for scientific research”.

Trey Breckenridge, director of high performance computing

Analytics in Production

YarcData
Getting to *Eureka!* faster

Discovery through fast hypothesis validation

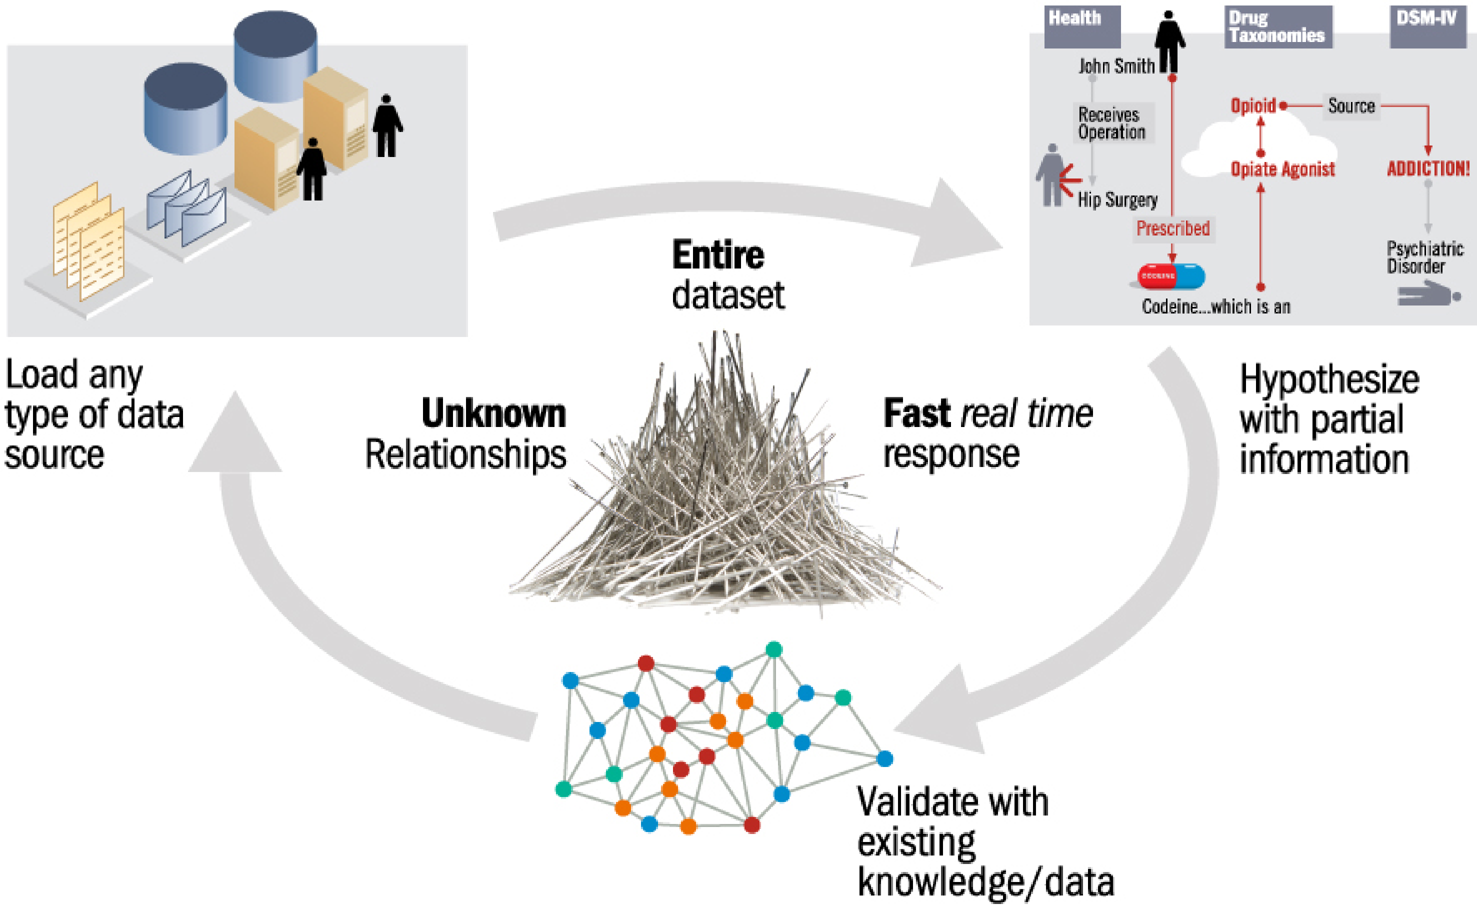

> "In the amount of time it takes to validate one hypothesis, we can now validate 1,000 hypotheses – massively improving our success rate and systematizing serendipity."
>
> \- YarcData Customer

Discovering New Drugs

- **Usage model:** Identify the complex relationships between the disease state and patients for the development of new biomarkers and drug pathways and identify new disease targets or companion diagnostics

- **Data sets:** Medline, PubMed, TCGA, Uniprot, Pfam, CRO, Clinical Trials…

- **Technical Challenges:** Size and Diversity of data; Complex inter-relationships; Entity resolution; Probabilistic relationships

- **Users:** Life sciences researchers in drug discovery and drug development

Institute for Systems Biology
Revolutionizing Science. Enhancing Life.

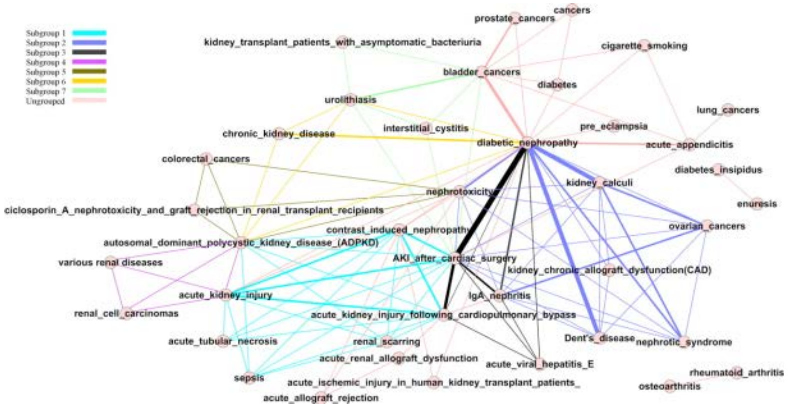

YarcData
Getting to Eureka! faster

Discovering New Matchup Advantages

- **Usage model:** Identify similar pitchers and explore corresponding linkages and evaluate batter performance against these groups to optimize lineup

- **Data sets:** Pitch data including release point, velocity, pitch type, pitch sequence, movement; Hit data including launch angle and batted ball velocity, Fatigue and batter performance; Venue and stage of game data

- **Technical Challenges:** Recent explosion of new data sources and game-by-game additions

- **Users:** Managers, Baseball Operations

MAJOR LEAGUE BASEBALL

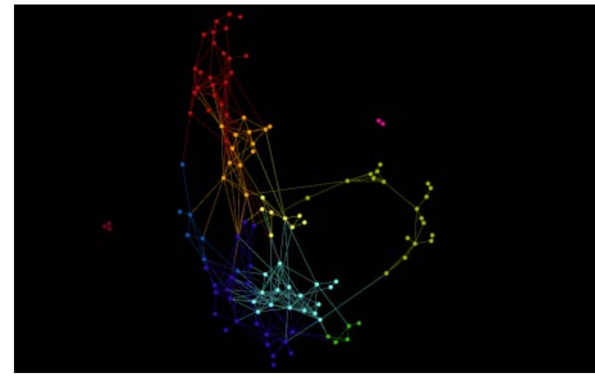

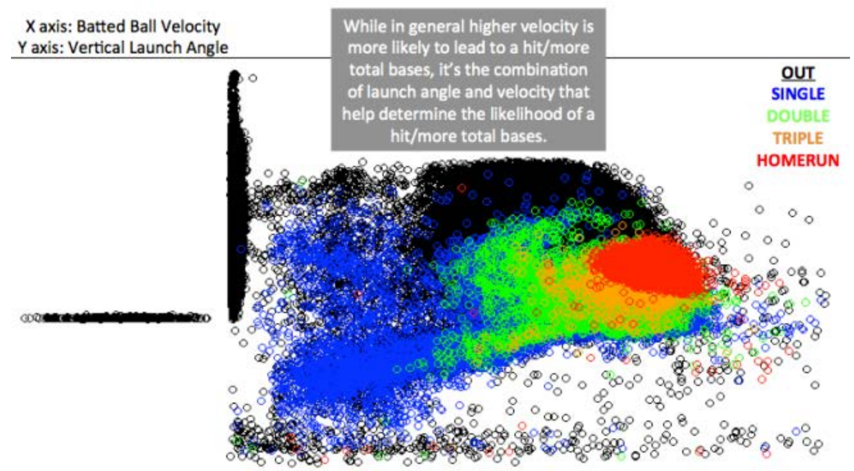

The Cray Storage Family

Leadership in Tiered Data Storage

Tiered Storage Workflow

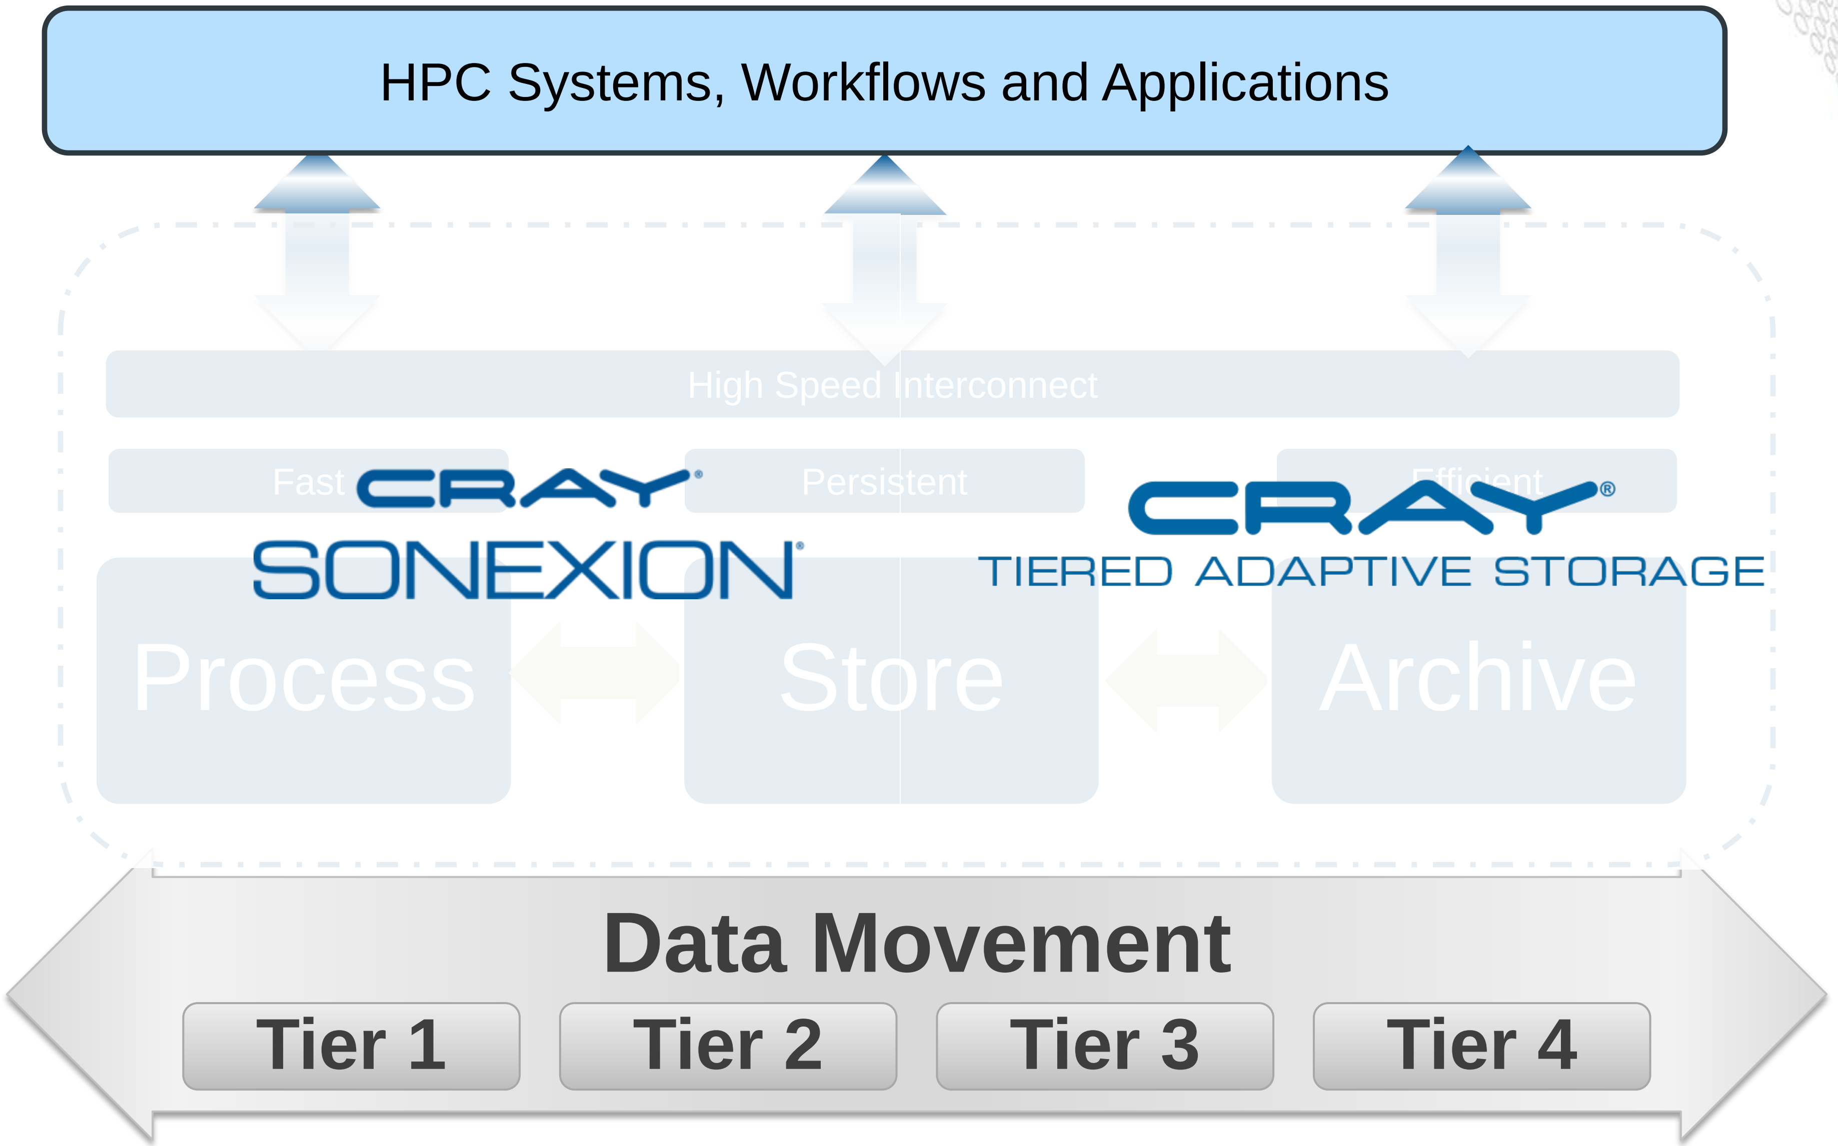

Cray Storage Sytems

CRAY SONEXION

CRAY TIERED ADAPTIVE STORAGE

True Lustre <-> HSM connectivity is coming

Powered By Versity

Scalable building blocks

- Best-of-breed storage technologies
- Open systems and software

Scale optimally – small to large systems

- Gigabytes to terabytes of performance
- Terabytes to exabytes of capacity

But different: Comparing TAS to SAM-QFS

Category	Oracle SAM-QFS	Cray TAS
Support M		W
Flexib		
Scal Oper		
Prod and		
Oper know		tools
Plan Integ		e oyed
Best		ces ble
Time t		
Migrate Data in Place	• Oracle supports in place migration from Solaris to Solaris	• Cray supports in place migration from Solaris to Linux

Things you can get with TAS that you don't get with SAM/QFS

1) A partnership with Cray instead of Oracle…
2) A Roadmap for HPC
3) TAS integration with Lustre
And much more
4) Support for Linux extended attributes
6) Enhanced failover
7) Large Inode support
8)Automated Failover
9) Native 4KB sector
10) …

Thank You!

Legal Disclaimer

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA, and YARCDATA. The following are trademarks of Cray Inc.: ACE, APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, THREADSTORM. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

THE SUPERCOMPUTER COMPANY