HPC Trends and Directions in the Earth Sciences

Per Nyberg
nyberg@cray.com
Sr. Director,
Business Development

CRAY®
THE SUPERCOMPUTER COMPANY

Cray Solutions for the Earth Sciences

Why Cray ?

- Cray solutions support:
  - Range of modeling capabilities
  - Research and operational environments
  - Shortened research to operations
- Experience in delivering and operating the world's largest and most complex systems
- Commitment to long-term partnerships

Market Presence

- Cray systems are used in a wide range of areas and configurations:
  - NWP and Climate
  - Used as model development platforms for extreme scale architectures
  - Compute and storage sizes range from Terascale to Petascale

What's Next ?

- Continued emphasis on system-wide approach to scalability – hardware and software
- Impact of latency sensitive big and fast data
- Impact of analytics workloads on HPC architectures
- Ultimately integrated HPC environments are the capability that will turn data in to insight and discovery.

Cray Announcements in Weather, Climate and Oceanography over the Last Two Years

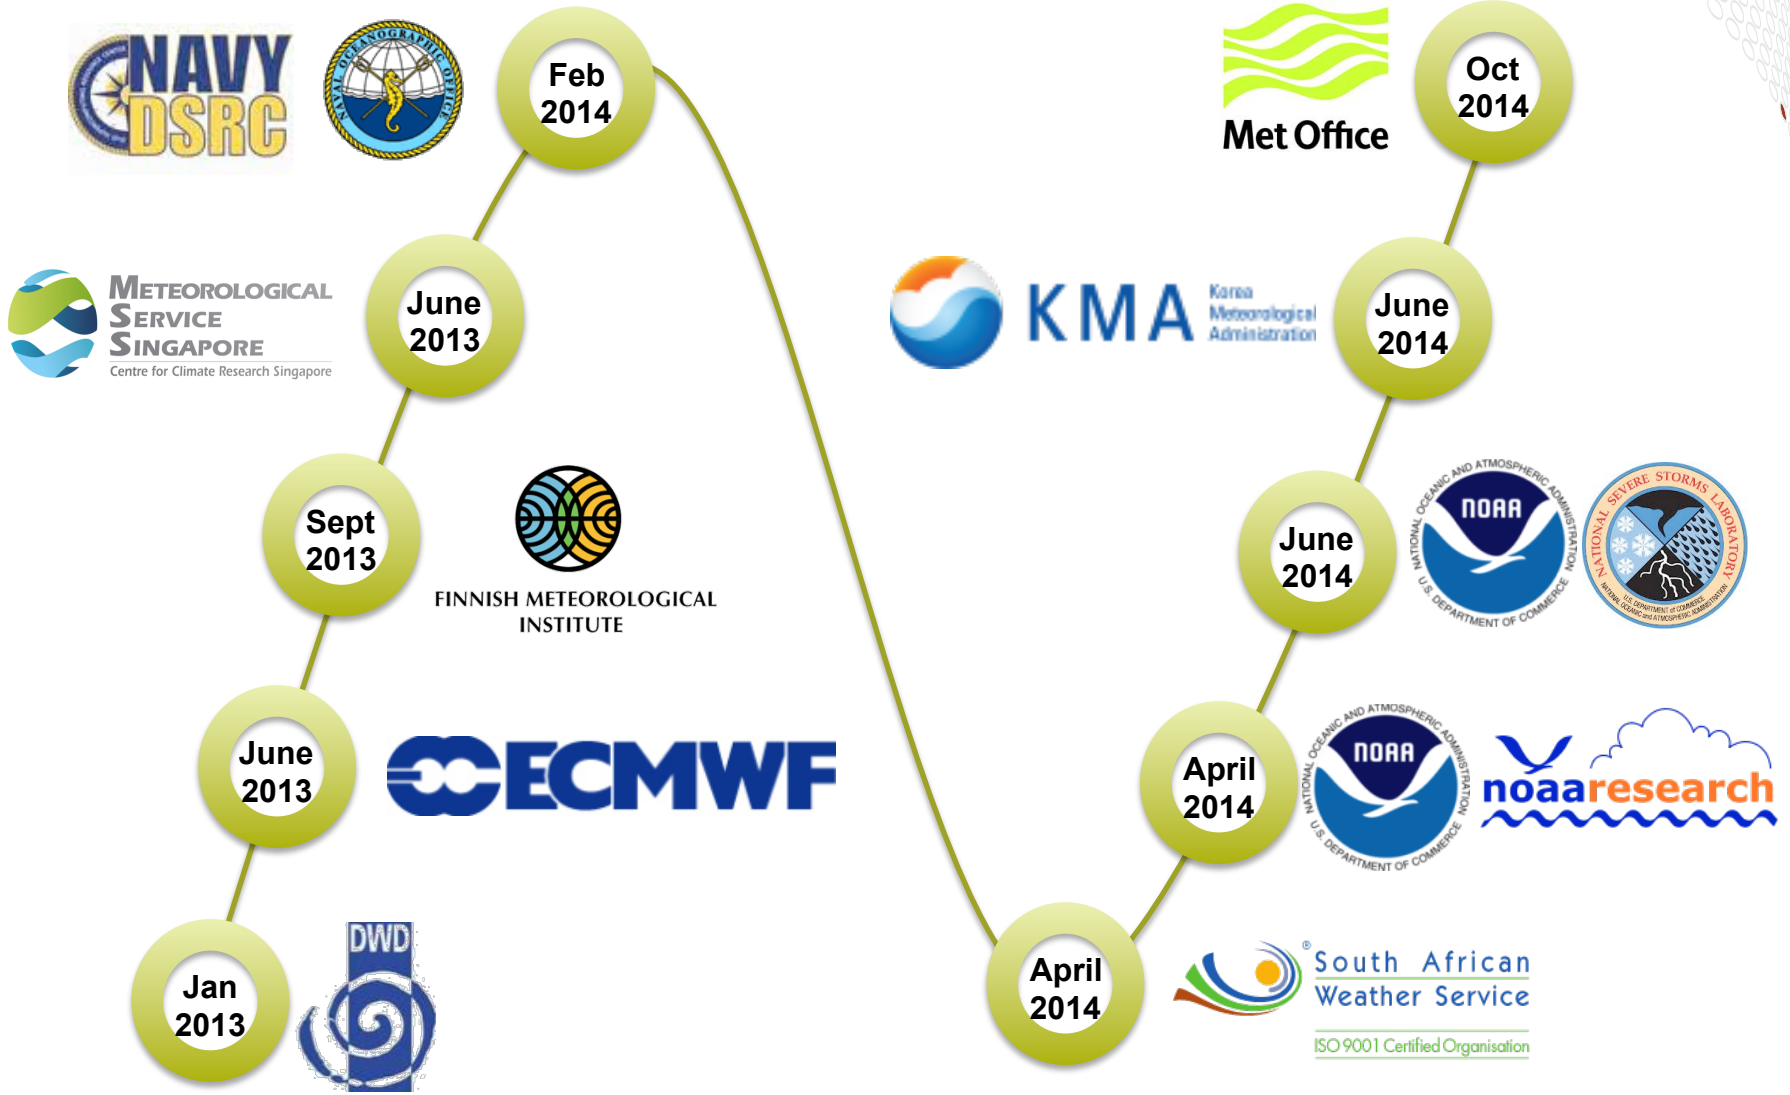

Cray XC Design Areas for Weather and Climate

- **Tightly integrated hardware and software stack designed for…**
- **Sustained Application Performance and Scalability:**
  - Uniquely capable of both large scale high resolution deterministic and ensemble workloads
  - Highly tuned software stack maximizes the cycles to the application
- **Operational workloads and environments**
- **Investment Protection – Upgradable by Design:**
  - Key system elements (CPUs, memory) can be easily upgraded as faster and more capable components become available
  - Systems can be easily expanded with additional blades / cabinets
  - Programming Environment designed for application life and portability
- **User Productivity**

Trends and Market Drivers…

Data and Computational Drivers

- **Today's science is:**
  - Data-intensive
  - Data-driven
  - Compute-intensive
  - Multi-disciplinary
  - Multi-scale, multi-physics
- **Driven by multiple dimensions.**
- **Data Tsunami is defying standard approaches to interpretation**
  - Volume and complexity of data are too much for either humans or current technologies for effective analysis
- **Creating necessity to re-examine current applications and develop new applications domains**

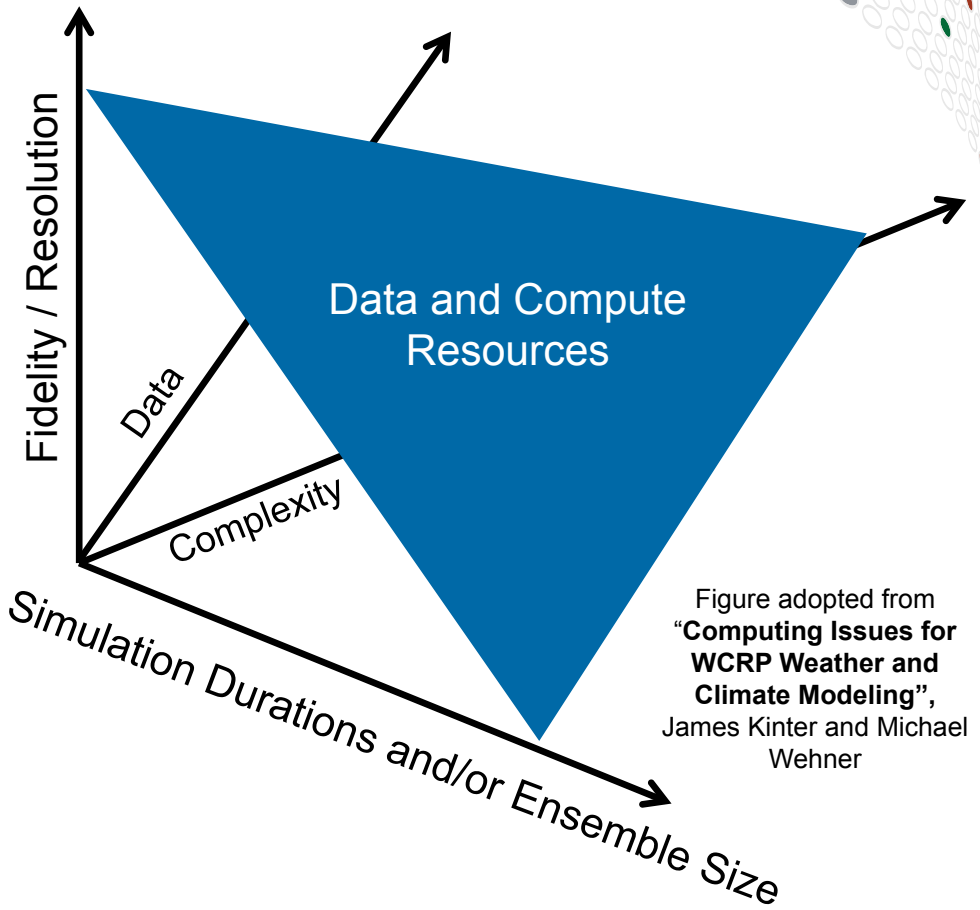

Figure adopted from "**Computing Issues for WCRP Weather and Climate Modeling**", James Kinter and Michael Wehner

Requires multi-dimensional technology stack – hardware, system software and applications

The Current Wave of Big Data….

Complex, unstructured, created by sensors,….

Data driven discovery and advanced analytics are rapidly becoming a competitive differentiators providing insight and predictive capabilities…

COMPUTE | STORE | ANALYZE

“Simulation was the first big data market” -- IDC

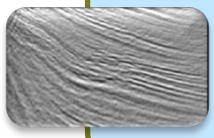

Terascale	Petascale	Exascale
Isotropic/anisotropic FWI elastic modeling/RTM	*Elastic FWI visco elastic modeling*	*Visco elastic FWI petro-elastic inversion*

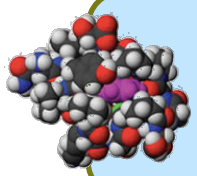

Terascale	Petascale	Exascale
Identify all protein sequences using public resources and metagenomics data, and systematic modelling of proteins belonging to the family	*Improving the prediction of protein structure by coupling new bio-informatics algorithm and massive molecular dynamics simulation approaches*	*Systematic identification of biological partners of proteins.*

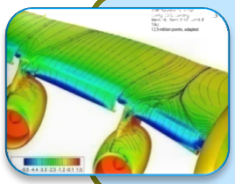

Terascale	Petascale	Exascale
Aero Optimisation & CFD-CSM	*Full multidisciplinary design optimization*	*CFD-based noise simulation*

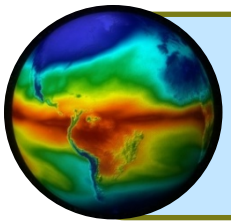

Terascale	Petascale	Exascale
Coupled AGCMs at 50km atmosphere and 1 deg ocean	*Earth System Models at 10 km atmosphere and 1/10 deg ocean*	*Full earth system models with carbon feedback cycle at 1km atmosphere and 1/100 deg ocean*

Major changes to applications and algorithms to address complexity and extend scalability

COMPUTE | STORE | ANALYZE

The Evolving HPC Workload: Simulation, data driven discovery and advanced analytics

Largest ever hurricane simulation at an urban scale resolution of 500m

Large-scale engine bay airflow optimization to improve vehicle fuel consumption

Highest-definition sub-surface images of oil fields improving development and production efficiencies

Healthcare fraud detection across disparate databases

Cyberattack detection through the discovery of relationships from multiple sources in billion point graphs

Viewer sentiment analysis to optimize television guest appearances

Predictive models to determine optimal athlete performance and likelihood of injury

Large, dynamic and complex analytics for portfolio risk analysis

The Evolving HPC Workload: *Economics of emerging use-cases will take part in driving architectures across all areas*

The Demand Side

Healthcare Fraud Detection (Source: IDC)

- 5 separate databases for the big USG health care programs under Centers for Medicare and Medicaid Services (CMS)
- Estimated fraud: $150B ~ $450B
- <$5B caught today)
- ORNL, SDSC have evaluation contracts to unify the databases and perform fraud detection on various architectures.

Sports Analytics (Source: Vince Genarro)

- Player, field of play and consumer data has exploded
- The stakes have grown dramatically
- $50—$100 million decisions are commonplace
- Winning Drives Profitability

Largest ever hurricane simulation at an urban scale resolution of 500m

Large-scale engine bay airflow optimization to improve vehicle fuel consumption

Highest-definition sub-surface images of oil fields improving development and production efficiencies

Health fraud detection across disparate databases

Cyberattack detection through the discovery of relationships from multiple sources in billion point graphs

Viewer sentiment analysis to optimize television guest appearances

Predictive models to determine optimal athlete performance and likelihood of injury

Large, dynamic and complex analytics for portfolio risk analysis

| STORE | ANALYZE

Growth in Data Science: Monsanto Example

Background

- In Oct 2013 Monsanto acquired The Climate Corporation for ~$1B.
- The Climate Corporation's expertise is in data science.
- Use a wide range of information to provde insights and recommendations for farmers, such as planting and irrigation schedules.
- "…offering <farmers> novel options in the way they manage risk on farm – including weather, which is the single biggest risk farmers face on an annual basis."

What do they do ?

- Based on a Hadoop implementation that creates weather projections for the next two years at every 2.5 by 2.5 kilometer grid across the U.S.
- Mapped out the most likely 10,000 outcomes per location using different variations of likely patterns to create a probabilistic view of weather.

Goal

- Goal is to help farmers understand the risks of their practices and to help them reduce those risks by changing their practices and by helping underwrite weather insurance against adverse effects.

Sources:
http://www.monsanto.com/features/pages/monsanto-acquires-the-climate-corporation.aspx
http://www.datamation.com/applications/hadoop-makes-a-big-data-splash-1.html

Emergence of Latency-Sensitive Analytics

Low-latency applications require performance optimizations
- *Memory-storage hierarchies*
- *Fast interconnects*

Response time frames

<30ms

30ms

10min

>10min

- *streaming data*
- *tweets*
- *event logs*
- *IoT*
- *SQL/ad hoc queries*
- *BI*
- *visualization*
- *exploration*

Low-Latency

- *summarization*
- *aggregation*
- *indexing*
- *ETL*

Batch

Hadoop tomorrow

Hadoop today

Streaming data / Data in motion

Stationary data

Enabling More Complexity & Capability …Big Data → Fast Data

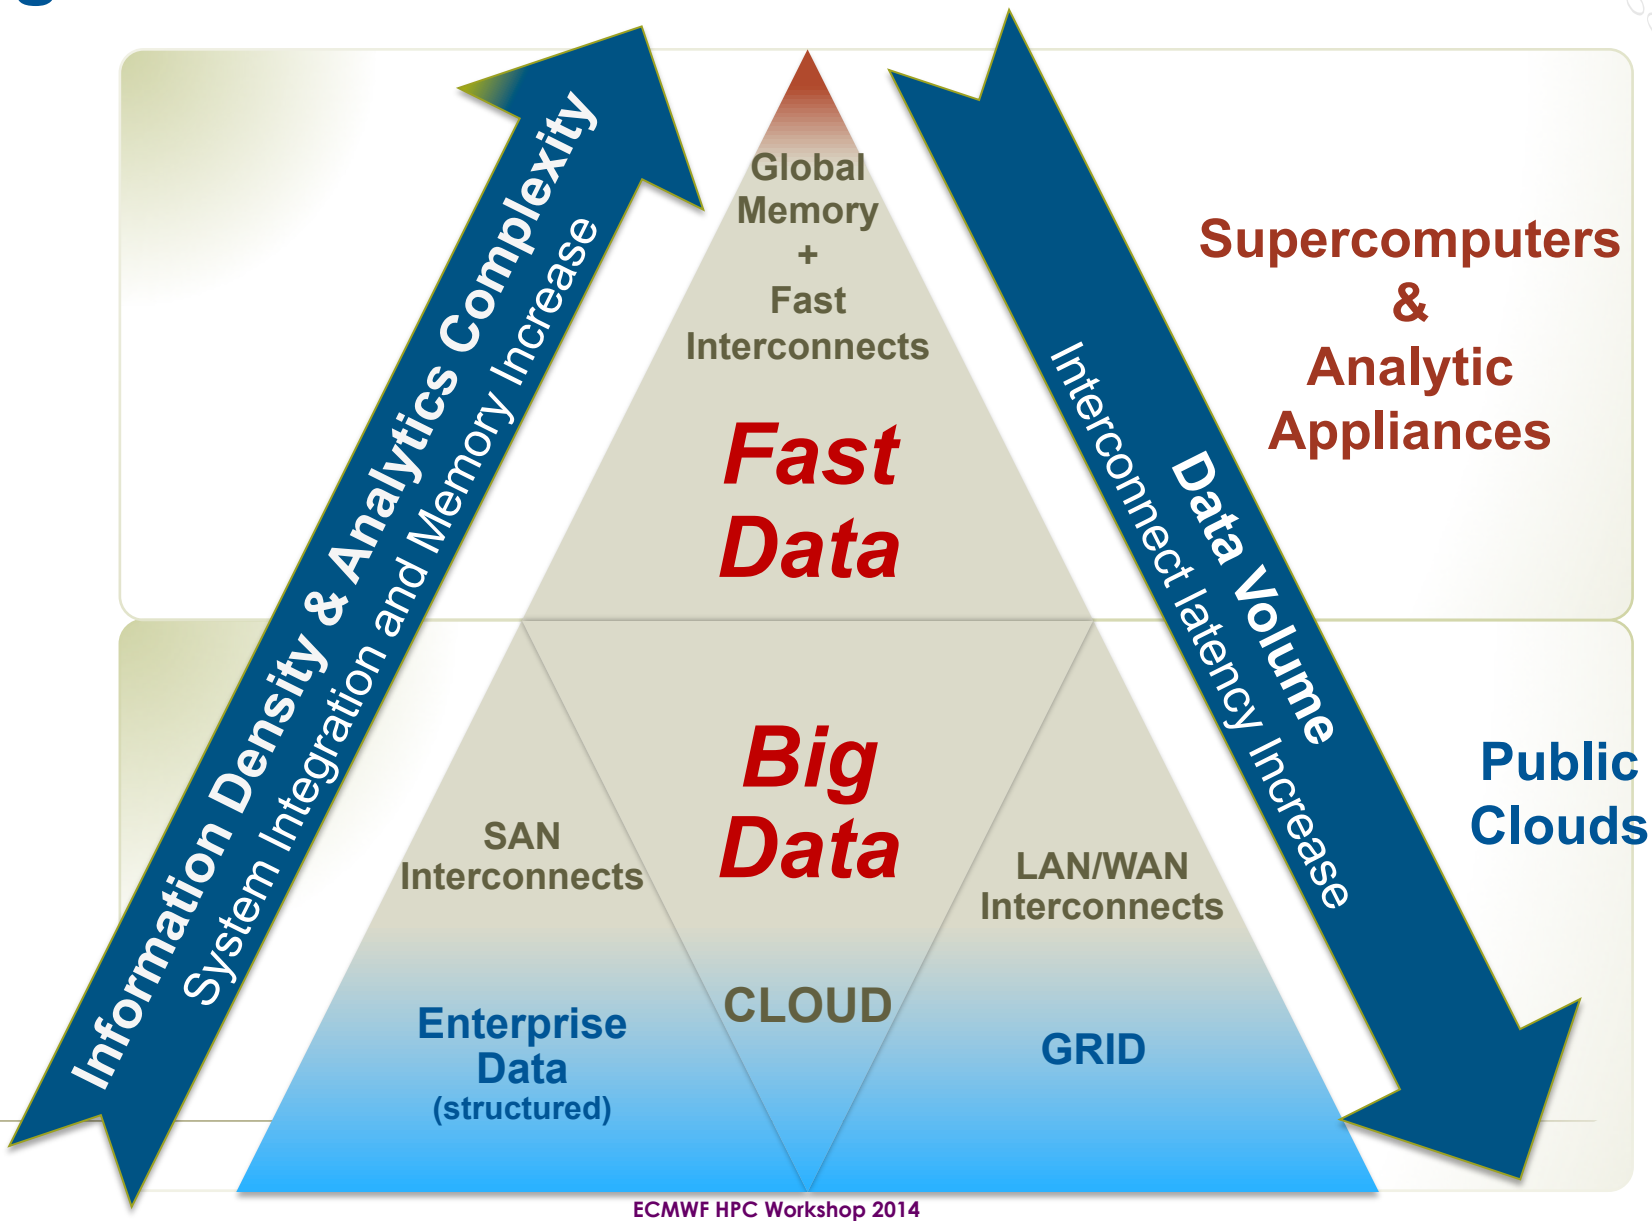

System Architecture Differences…

Supercomputing

- Scalable computing w/high BW, low-latency, global memory architectures
- Tightly integrated processor-memory-interconnect & network storage
- Minimize data movement – load the “mesh” into memory
- Move data for loading, check-pointing or archiving
- “Basketball court sized” systems

Large-Scale Data Analytics

- Distributed computing at largest scale
- Divide-and-conquer approaches on Service Orientated Architectures
- Maximize data movement-- scan/sort/stream all the data all the time
- Lowest cost processor-memory-interconnect & local storage
- “Warehouse sized” private and public clouds

The Fusion of Supercomputing with Large Scale Data Analytics

Key Trends

- The cost per FLOP is dropping dramatically
- Cost of moving data dominates performance and energy consumption
- New levels of memory and storage hierarchy are emerging
- Effective Parallelism of O(Million ~ Billion)
- Analytic application needs for memory size and IO performance are insatiable!
- Future systems are being influenced by the converged infrastructure of compute & storage
- Converged infrastructure of Hyperscale demands management capability of O(10000) nodes

Changing Needs
The Fusion of Supercomputing and Big & Fast Data

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

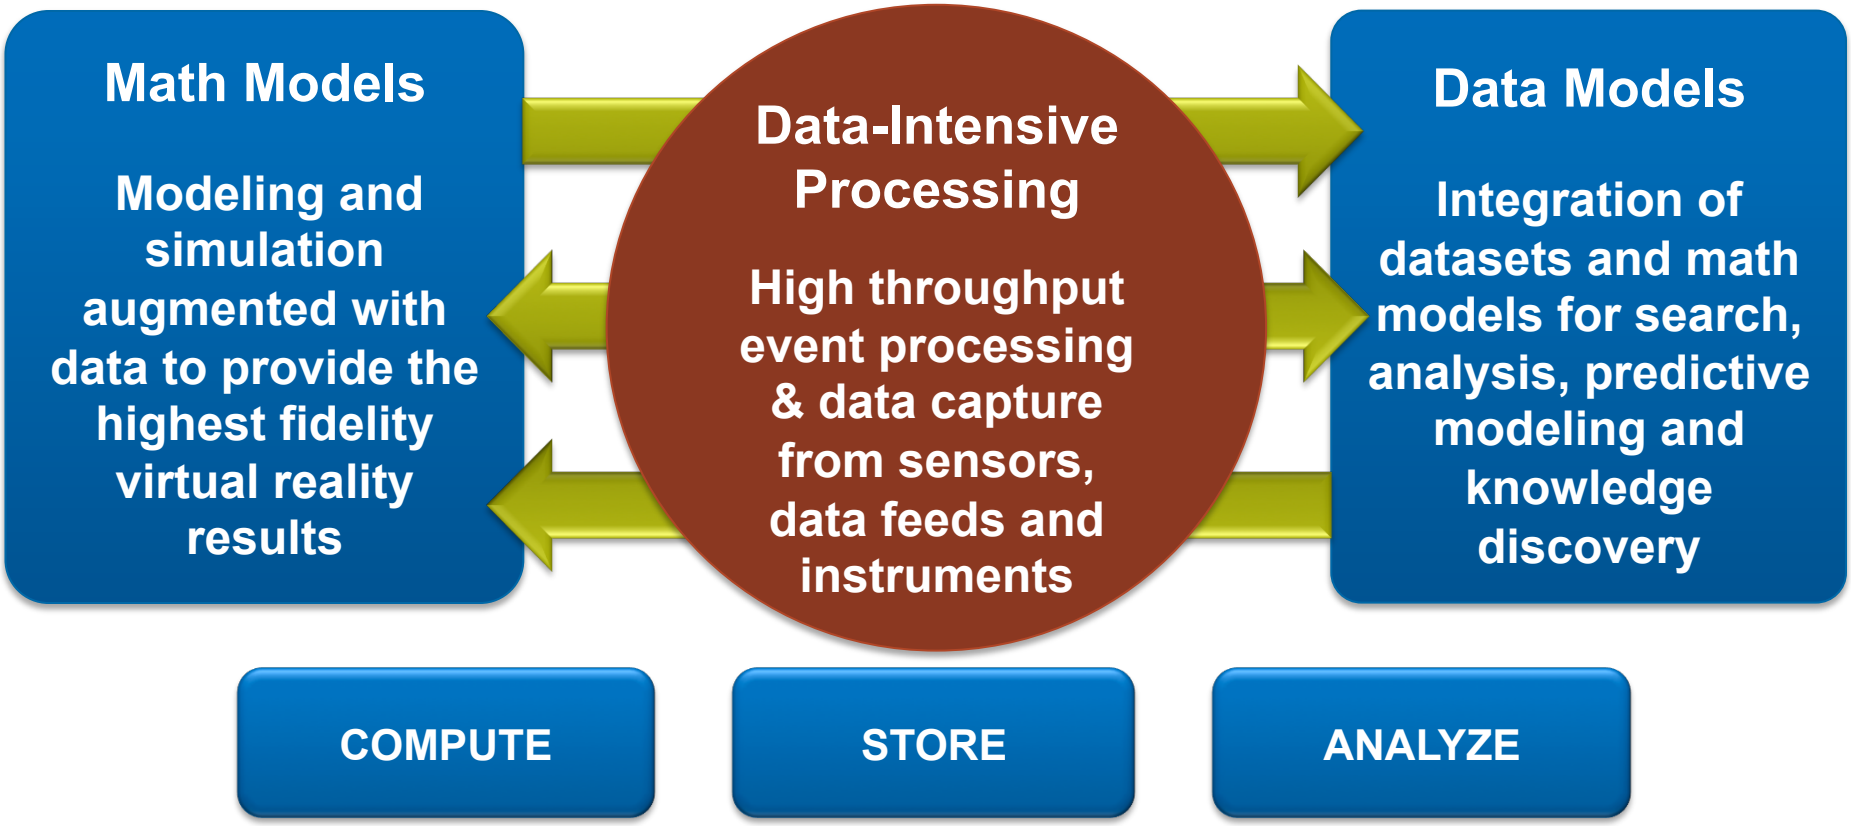

Cray Integrated Environments for Simulation and Discovery

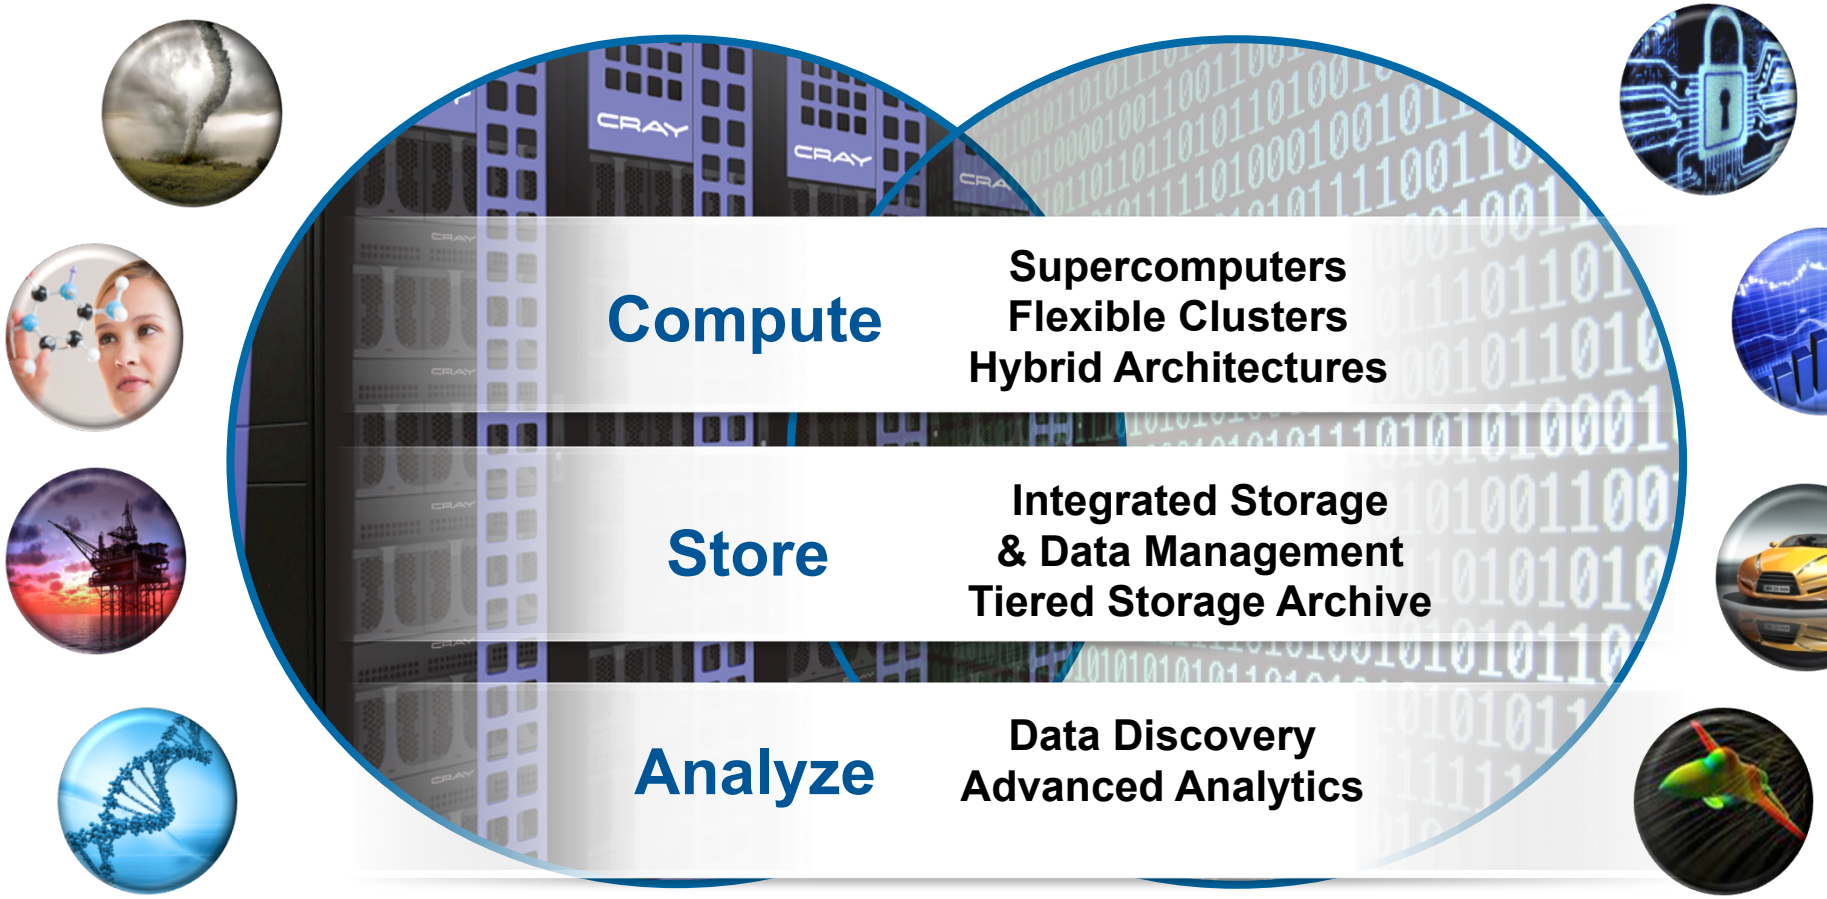

Recently Announced: Urika-XA Platform

Turnkey Advanced Analytics Platform

Next-Generation System Architecture

Engineered for Performance

- *Hadoop and Spark ecosystem*
- *Emerging analytic workloads*
- *Open platform for current and future frameworks*
- *Single pane of glass for system management*
- *Innovative use of storage technologies*
- *Battle-tested on cutting-edge government/scientific analytic applications*
- *Ready for the enterprise*
- *Dense footprint: over 1,500 cores, 6TB memory*
- *38TB SSD and 120TB POSIX-compliant high-performance storage*
- *InfiniBand*
- *Cray Adaptive Runtime for Hadoop*
- *Scale out to multi-rack configurations*

Recent Cray XC40 Announcement

What is it?

- Follow-on to highly successful Cray XC30
- Already received orders for over 70,000 sockets
- Much of which is already installed

What's New ?

- DataWarp applications I/O accelerator
  - New tier of high performance flash SSD directly connected to the Cray XC40 compute nodes.
  - Delivers a balanced and cohesive end-to-end system architecture
- Support for Intel® Xeon® processor E5-2600 v3 product family, formerly code named "Haswell,"

Comes in small packages too – the Cray XC40-AC

- Cray's core technologies… in a smaller package
- Same Powerful and Productive Software Ecosystem
- Same Aries Infrastructure
- Same DataWarp I/O Options
- …Only the cooling is different.

CRAY XC40-AC

COMPUTE | STORE | ANALYZE

Trends in the Memory/Storage Subsystem

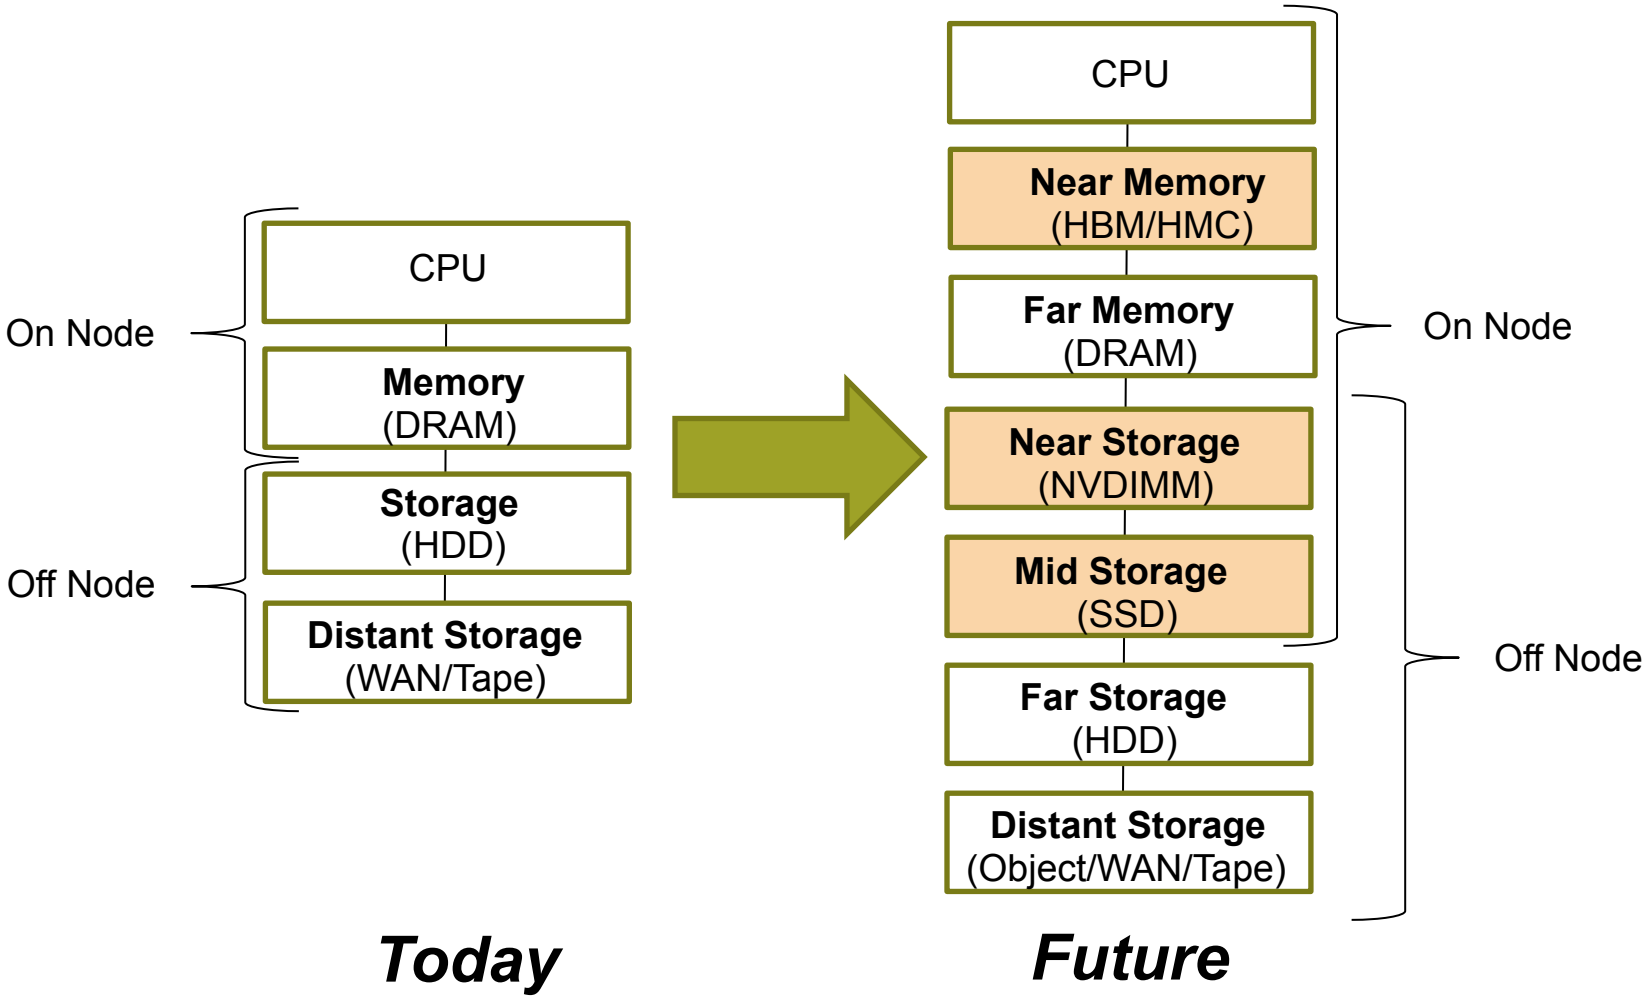

Two Recent Large XC Orders that will Impact Future Technologies and Applications Throughout the Community

- **Los Alamos / Sandia – Trinity**
  - 42 Pflop system
  - Heterogeneous System – Haswell + KNL
  - 3PB SSD DataWarp Capability

- **NERSC8**
  - Installation Date: 2Q 2016
  - Self-hosted KNL System
  - ~28 Pflops Peak
  - Transitioning user base to “many-core” processing
  - NERSC Exascale Science Applications Program (NESAP)

NERSC Exascale Science Applications Program (NESAP)

- **Collaboration with Cray and Intel to prepare for "Cori", the Cray XC to be deployed at NERSC in 2016.**
- **NESAP was launched to ensure that the highly diverse workloads of the DOE science community continue to be supported as over 5,000 users make the transition to Cori.**
- **Application areas:**
  - Fusion Energy Sciences, High Energy Physics, Nuclear Physics,…
  - Biological and Environmental Research:
    - ESM Global Climate Modeling, John Dennis (NCAR)
    - High-Resolution Global Coupled Climate Simulation Using The Accelerated Climate Model for Energy (ACME), Hans Johansen (LBNL)
    - Multi-Scale Ocean Simulation for Studying Global to Regional Climate Change, Todd Ringler (LANL)
- **Large number of community models are involved.**

Moving Forwards: Adapting to Simulation and Data-Intensive Workloads by Adding Value at the Edge of the Network

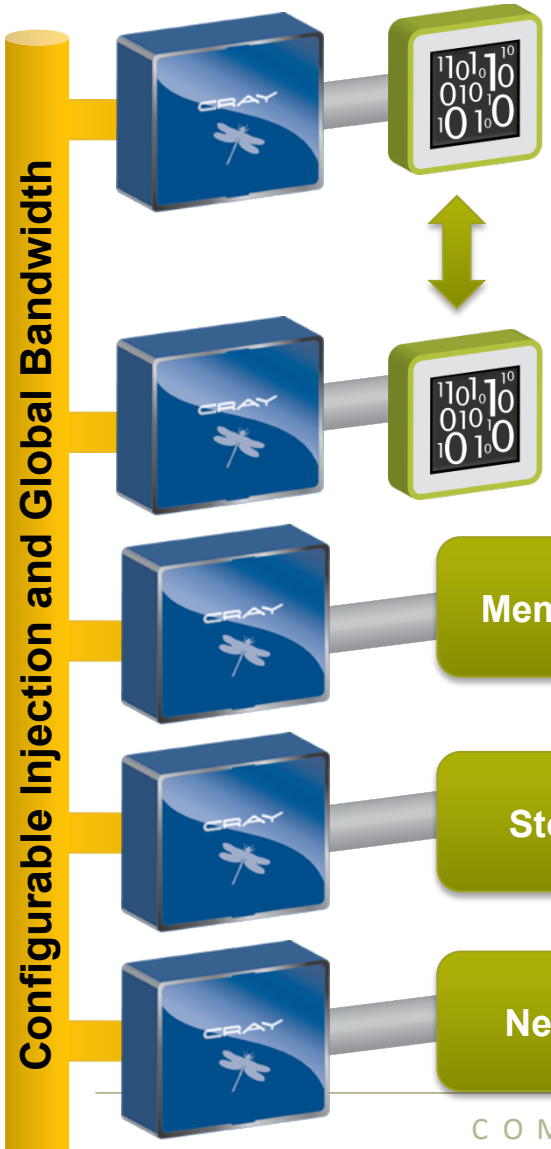

Data Processors
- *x86*
- *Many-core / GPU*
- *Multi-threaded*
- *Flexible compute for specialized functions*

System-wide approach to workflow performance

System Software
- *High throughput scheduling*
- *Adaptive runtimes / Adaptive routing*
- *Compilers, Auto-tuning, Programming models*
- *Middleware stack to support simulation and analytics workflows*

Memory Pools
- *DRAM*
- *Active NVRAM/NVDIMM*
- *HMC/HBM*

Storage I/O
- *Tiered Storage from SSD's to Disk to Tape*
- *Tightly coupled Software Stack for RAID*
- *Very Low Latency*

Network I/O
- *Intelligent I/O & Protocol Load Balancing*
- *Specialized Data Ingest*
- *Policy Based Data Movement/Software Defined Network*

COMPUTE | STORE | ANALYZE

Summary

- Weather, climate and ocean modeling is a key community within Cray's customer base.
- Data driven discovery and advanced analytics are rapidly becoming a competitive differentiators in both traditional and non-traditional HPC areas.
- The HPC workload is evolving and is driving the need for advanced architectures.
- Cray's vision is that the fusion of Supercomputing and Big & Fast Data is the capability that will turn data into insight and discovery.

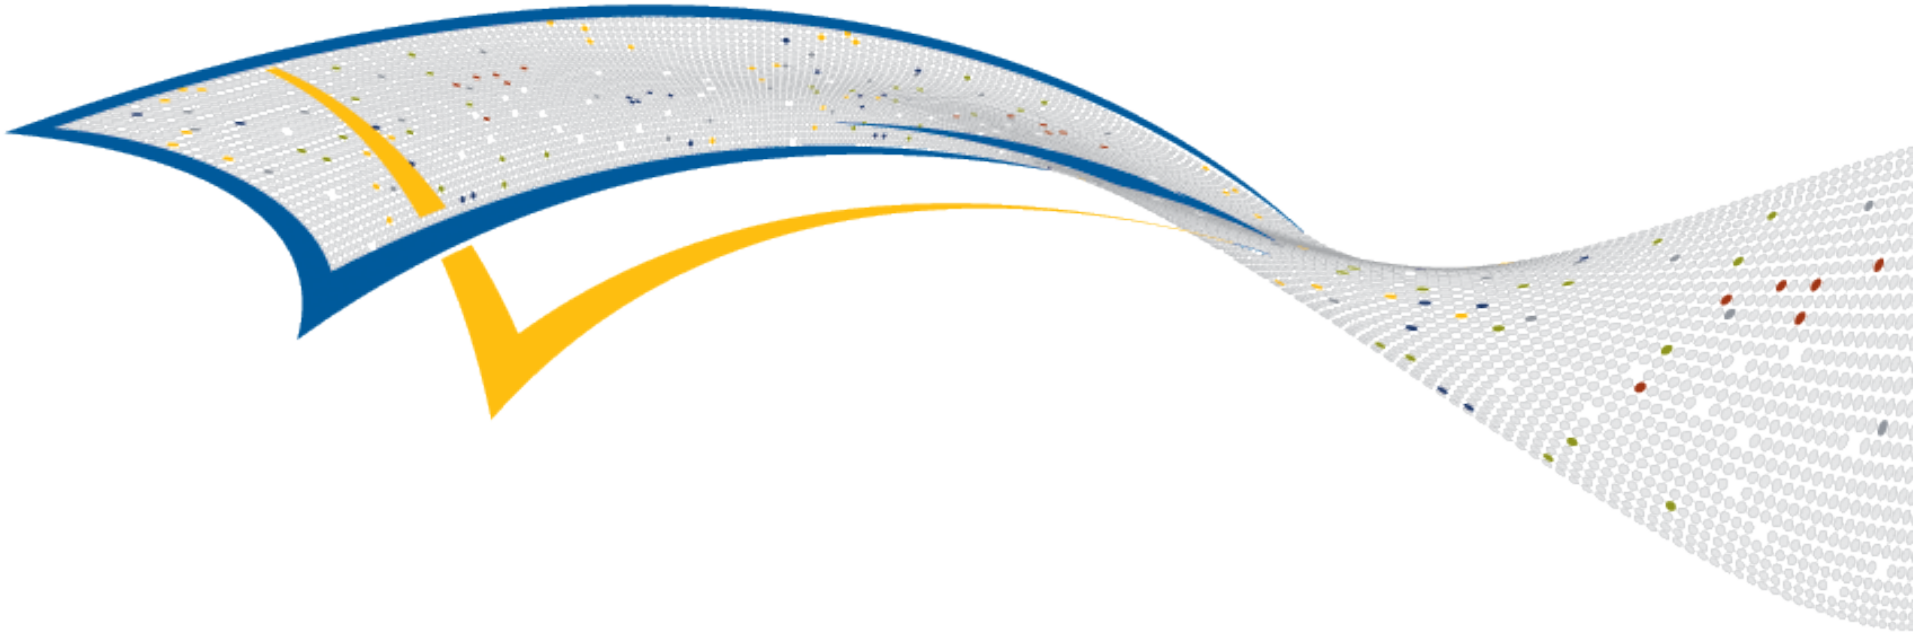

Thank you for your attention

CRAY®

THE SUPERCOMPUTER COMPANY

COMPUTE | STORE | ANALYZE