



COMPUTE



STORE



ANALYZE

# **Cray Earth Sciences Update**

**Phil Brown**

**[philipb@cray.com](mailto:philipb@cray.com)**

**Earth Sciences Segment Leader**

***iCAS, 17<sup>th</sup> September 2015, Annecy***



- **Cray's Presence in Earth System Modelling Community**
- **Emerging Trends in Weather & Climate Data Analysis**
- **Cray's Vision for Converged Modelling & Analytics Platforms**

# Cray Solutions for Earth System Modelling

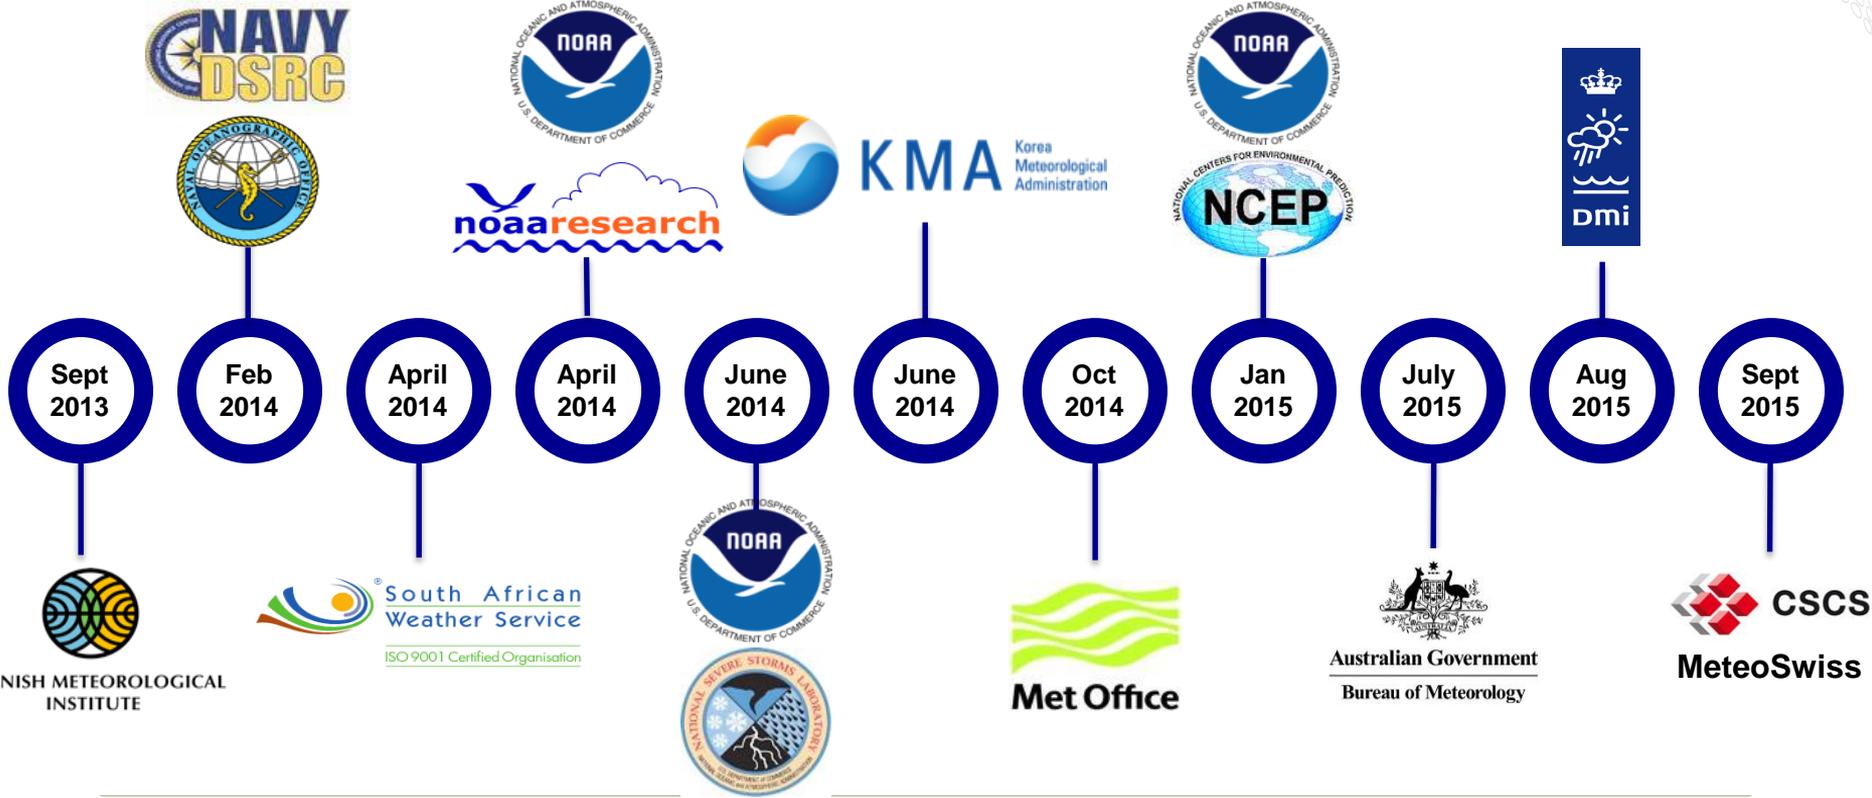


- **Cray offers a complete, single-vendor solution for ESM:**
  - Balanced integration of compute, storage and analytics
- **Cray's industry leading technology enables customers to address a broader & more challenging set of problems**
- **Cray has huge experience delivering and operating the world's largest and most complex systems, both in operational & research environments**
  - Emphasis on total cost of ownership – from power and upgradability, to efficiency.
- **Cray is committed to long-term partnerships that provide significant ongoing value to our customers.**

---

COMPUTE | STORE | ANALYZE

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



COMPUTE | STORE | ANALYZE

# Bureau of Meteorology, Australia



- In the second half of 2015 Cray will deploy a new Cray supercomputer at the Bureau of Meteorology. The new Cray facility will allow BoM to increase forecasting capabilities by providing high performance computing to support:
  - Increased resolution and model enhancements
  - Introduce ensemble runs for global & local regions



System Overview

- Dual Cray XC40 systems for operations and research/backup
- Systems expanded in 2018 with latest Intel Xeon processors
- Cray Sonexion storage
- Applications include global, regional, city-scale & tropical cyclone and climate research.



Testimonial

"At the moment we're at about a 4km grid degree of resolution on the city models. By 2020 we'll be down to a 1km grid for the city model. We'll also be running more regular and frequent updates - 24 times a day by 2020, where at the moment we run it four times daily. We'll also be able to run high, on-demand models for tropical cycles and run two extreme weather events side by side. The benefits to the community, the economy, emergency services, national security and various parts of the industry will be quite marked."

Lesley Seebeck, BoM CIO. *As told to iTnews.com.au*

COMPUTE | STORE | ANALYZE

# Danish Meteorological Institute



- Cray was recently selected by DMI to provide their next two generations of supercomputing systems, succeeding their current Cray XT5 platforms. The supercomputers will be hosted in Iceland, taking advantage of the cool climate and abundant geo-thermal and hydro electricity.
- Initially the Cray systems will deliver three times the performance of the current systems, rising to a 10x performance after a mid-life upgrade in 2018

## System Overview

- Dual Cray XC40 systems for operations and research/backup
- Systems upgraded in 2018 with latest Intel Xeon processors
- Cray Sonexion storage
- Applications include high resolution regional forecasts for Denmark and the North Atlantic including Iceland & Greenland



## Testimonial

It was paramount to us to find a business partner with the organization and skills to install a supercomputing solution in Iceland that could operate remotely 1,300 miles away from Denmark, and Cray filled all the criteria

Mrs. Marianne Thyrring, Director General of the Danish Meteorological Institute.

# Trends in Weather & Climate Data Analysis



- **Data volumes generated by supercomputers growing rapidly**
  - Driven by increasing resolutions, ensembles members & number of experiments
- **Cumulative archives also expanding rapidly**
  - ~40-50% CAGR not unusual
  - DKRZ expect to grow at ~75PB p/a with new supercomputer
  - ECMWF archive growing at about 50PB/year (>60% CAGR)
- **Even simple data analysis efforts can be I/O bound**

# Weather/Climate Informatics



- **Emerging analysis approaches adopted from data analytics & machine learning space**
- **Examples includes:**
  - Complex network & graph based approaches
  - Scalable optimization methods
  - Supervised/unsupervised learning
- **Some use-cases:**
  - Detecting links between elements in climate system
  - Automated forecaster guidance/decision support
  - Optimizing integration of multi-model climate ensembles
- **Computation & communication intensity likely to grow, with IO remaining very important**
- **Analysis will also become more tightly coupled within overall workflow**

---

COMPUTE | STORE | ANALYZE

# Weather/Climate Informatics



5<sup>th</sup> In  
Septer

2015

Hosted

- supervised/uns

ICCS 2015: “Computational Science at the Gates of Nature”

## Sixth Workshop on Data Mining in Earth System Science (DMESS 2015)

Co-conveners: [Forrest M. Hoffman](#), [Jitendra Kumar](#), and [J. Walter Larson](#)

Reykjavík, Iceland | June 1–3, 2015

UNIVERSITY OF MINNESOTA  
Driven to Discover™

NSF

Fifth Annual Workshop on  
**Understanding Climate Change from Data**  
August 4-5, 2015 | University of Minnesota Minneapolis, MN

- Analysis will also become more tightly coupled within overall workflow

COMPUTE | STORE | ANALYZE

# *Cray Today*



COMPUTE | STORE | ANALYZE

# Urika-XA Extreme Analytics Platform



## Urika-XA



**Turnkey Advanced Analytics Platform**

**Next-Generation System Architecture**

**Engineered for Performance**

**Advanced Analytics at Lower TCO**

- Open platform for both pre-configured and user-installed tools
- Hadoop and Spark ecosystem
- Emerging high performance analytic workloads
- Unified system management interface
- High performance 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 performance networking
- Cray Adaptive Runtime for Hadoop
- Scale out to multi-rack configurations

COMPUTE

STORE

ANALYZE

# Cray DataWarp I/O Acceleration for Cray XC40



- **Pure performance**
  - 70 thousand to 40 million IOPS per system
  - Quality of Service to applications
- **Breakthrough efficiencies**
  - 5x the bandwidth of disk at the same cost
- **Flexible Usage Models**
  - Local and Shared I/O models
  - No application changes required

**CRAY**  
**DATAWARP™**



*DataWarp overcomes the performance gap between compute and disk storage*

COMPUTE | STORE | ANALYZE

# Containers on Cray XC

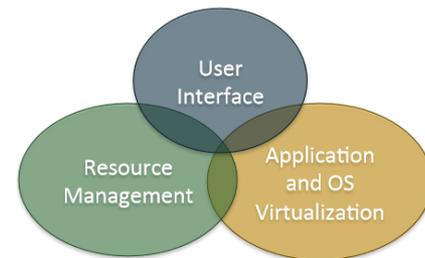


- **Allow user-defined environment**
  - Either customized, or pre-defined for particular workflow
- **NERSC are experimenting with this – likely to form a key element of the 2016 Cori XC40/KNL system**

## Prototype Implementation: “Shifter”



- **Supports**
  - Docker Images
  - CHOS Images
  - Will be able to support other image types (e.g., qcow2, vmware, etc)
- **Basic Idea**
  - Convert native image format to common format
  - Construct chroot tree on compute nodes using common format image
  - Directly use linux VFS namespaces to support

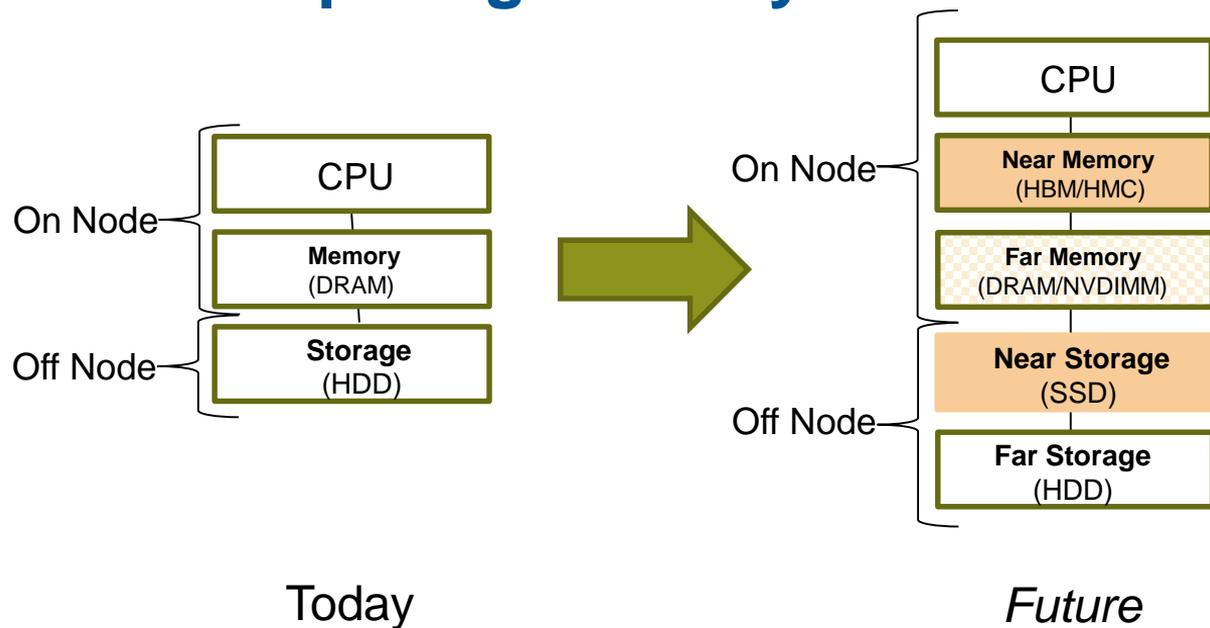


# Our Vision...

A wide-angle photograph of a two-lane asphalt road stretching into the distance towards a bright sunset. The sun is a large, glowing orb on the horizon, casting a warm orange and yellow light across the sky and the road. The sky is filled with soft, wispy clouds. To the left of the road, there is a green field and a body of water reflecting the sunset. To the right, there are rolling hills and a fence line. The overall mood is one of hope and forward-looking vision.

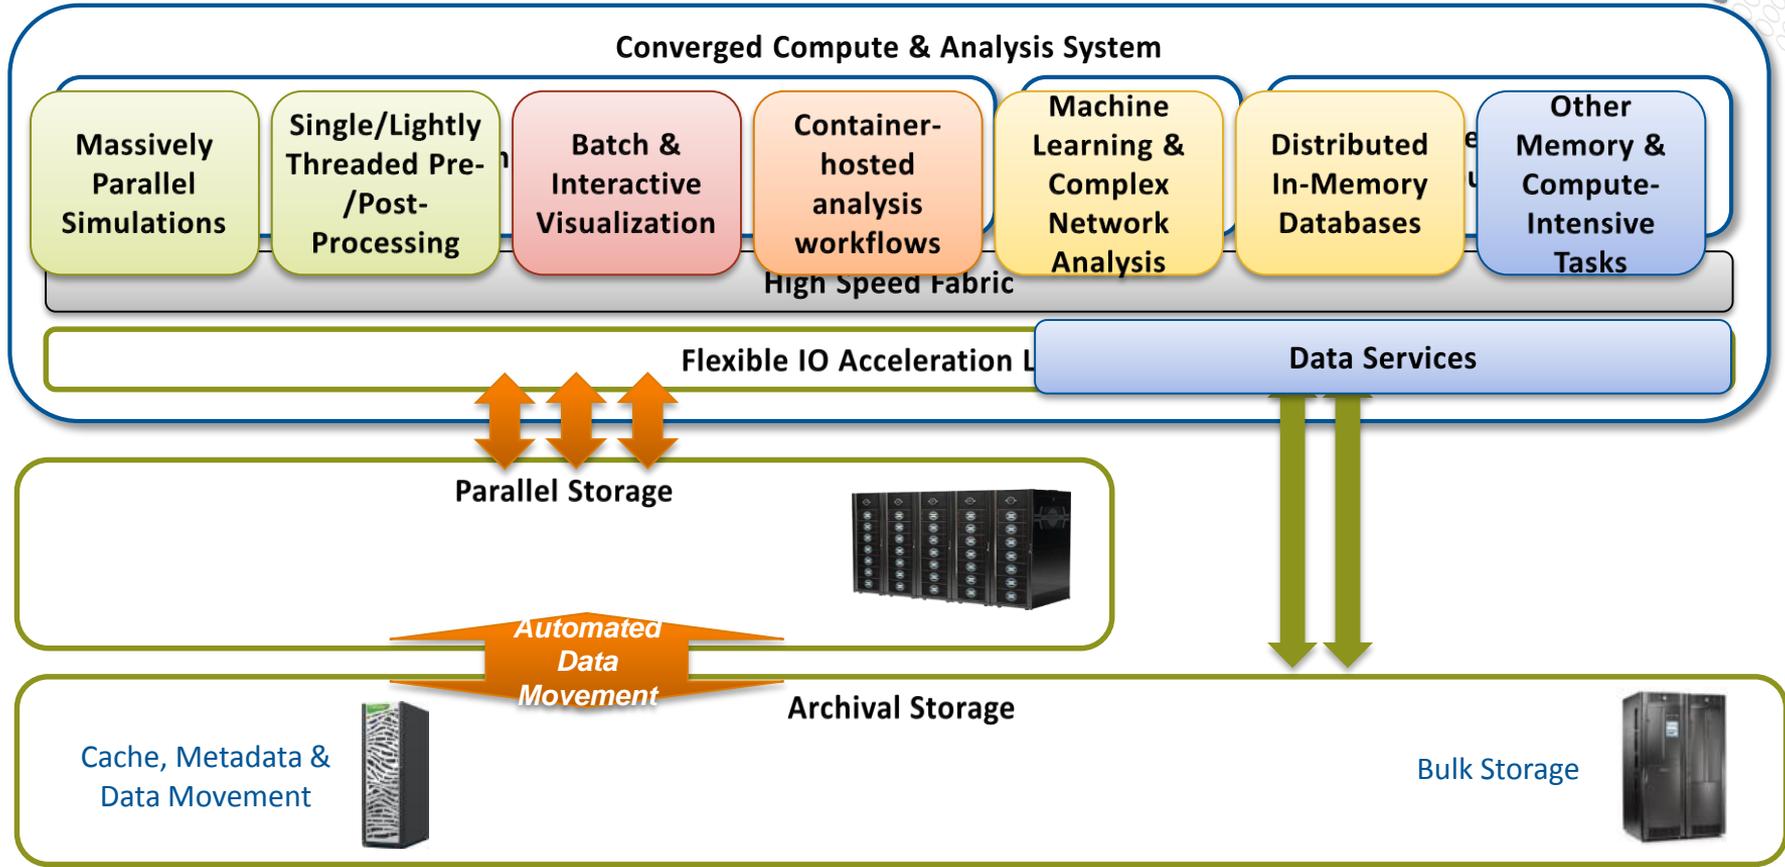
**Build a world-class integrated supercomputing environment that enables transformational computing across a broad set of science, engineering and advanced analytics (big data) applications**

# Exascale Computing Memory Trends



- **Good:** Helps reduce/mitigate cost of moving data
- **Bad:** Even more complexity in programming models

# Future Converged Architecture



COMPUTE | STORE | ANALYZE

# Integrated HPC Environments are the capability that will turn data in to insight and discovery

## Compute



## Store



## Analyze





COMPUTE



STORE



ANALYZE

The background of the slide is a dramatic, high-contrast image of a stormy sky. Dark, heavy clouds are illuminated from below, creating a bright, ethereal glow. In the lower-left quadrant, a bright, jagged lightning bolt strikes downwards. The overall color palette is dominated by deep blues, greys, and a hint of yellowish-orange light from the bottom right, suggesting a sunset or sunrise breaking through the storm.

**Thank you!**