Cray Shasta for Modeling, Simulation, Analytics and AI for Weather and Climate

Ilene Carpenter

September 2019

icarpenter@cray.com

Weather, Water and Climate: Why CRAY?

Most of the world's operational weather forecast centers use Cray systems

- Reliability
  - Operationally proven, unrivalled experience
- Performance
  - Balance performance & throughput across workflow
  - Software development environment, performance tools & application support experts
- Long-term customer partnerships

CRAY Growth in Earth Sciences

CRAY

May 2016 — Canada

Oct 2016 — Met Office

Feb 2017 — Global NWP Centre

Mar 2017 — Ministry of Earth Sciences, Government of India

Jun 2017 — NIWA Taihoro Nukurangi

May 2018 — Japan Meteorological Agency; ZAMG Zentralanstalt für Meteorologie und Geodynamik

Sept 2018 — IBS Center for Climate Physics

March 2019 — Royal Thai Navy

July 2019 — Oak Ridge National Laboratory; Schweizerische Eidgenossenschaft, Confédération suisse, Confederazione Svizzera, Confederaziun svizra

IBS Center for Climate Physics

NCEP

ANI

Australian Government Bureau of Meteorology

Met Office

DWD

Japan Meteorological Agency

CSCS

South African Weather Service
ISO 9001 Certified Organisation

Korea Meteorological Administration

noaaresearch

Meteorological Service Singapore
Centre for Climate Research Singapore

NAVY DSRC

NCMRWF

FINNISH METEOROLOGICAL INSTITUTE

NIWA Taihoro Nukurangi

ECMWF

DMI

PRESS RELEASE

View printer-friendly version << Back

SWISS FEDERAL OFFICE OF METEOROLOGY AND CLIMATOLOGY ADVANCES WEATHER FORECASTING WITH NEW CRAY SUPERCOMPUTER AND STORAGE

Swiss National Weather Center to Deploy its Third Cray CS-Storm System with Cray ClusterStor Storage to Enhance the Quality and Accuracy of the Country's Weather Predictions

SEATTLE and FRANKFURT, Germany, June 18, 2019 (GLOBE NEWSWIRE) -- Today at the 2019 International Supercomputing Conference in Frankfurt, Germany, global supercomputer leader Cray Inc. (Nasdaq:CRAY) announced that the Swiss National Supercomputing Centre (CSCS) is adding a third Cray® CS-Storm™ supercomputer with Cray ClusterStor® storage to support the development of cutting-edge weather service products at the Swiss Federal Office of Meteorology and Climatology (MeteoSwiss). The new CS-Storm system will join two Cray supercomputers already on site at CSCS, used by MeteoSwiss to run production weather forecasting. The job of the newly-accepted CS-Storm is to allow MeteoSwiss to run larger ensemble predictions and higher resolution weather models to improve forecast accuracy.

"Forecasting weather is like a 24-hour marathon–it never stops and it requires technology to keep pace," said Prof. Dr. Thomas Schulthess, director of CSCS. "And the complexity is growing in weather and climate forecasting and modeling as we put advanced workloads to the test. Our partnership with Cray continues to provide us the compute resources necessary so customers like MeteoSwiss can tackle the biggest challenges in meteorological research. This new system will make it possible for them to more accurately predict the weather."

MeteoSwiss continually invests in innovation and development for the benefit of the economy and public safety. With CSCS' purchase of its first two Cray CS-Storm systems in 2015, MeteoSwiss became the first major national weather service to use a GPU-accelerated supercomputer to run production numerical weather models. And with its latest CS-Storm – designed for organizations that require the performance and efficiency of today's latest GPU accelerators – MeteoSwiss will continue with the GPU approach.

"Complete, extremely detailed forecasts is our vision for the future," said Philippe Steiner, Head of Numerical Prediction at MeteoSwiss. "Switzerland's unique orography adds a level of complexity to forecasting that requires us to develop advanced, data-intensive modeling methods and techniques. With this latest technology from Cray, MeteoSwiss will be able to run increasingly complex, high-resolution models for higher-accuracy forecasts and nowcasts, as well as to develop additional weather forecast products."

MeteoSwiss found success with its existing Cray supercomputers and selected this new CS-Storm to provide the additional computational power required to process increasing volumes of weather observations and produce higher fidelity forecasts. The CS-Storm system was also selected for its ability to run numerical weather forecasts within a reduced energy footprint (as compared to competing solutions), and for the reliability the platform provides MeteoSwiss when running critical workloads.

"The new Cray CS-Storm at CSCS equips MeteoSwiss with the computational performance required to power next-generation weather forecasts. Cray is proud to play a part in helping support the amazing work MeteoSwiss does to keep Swiss citizens safe, informed and protected through accurate and timely weather forecasts," said Ilene Carpenter, Earth Sciences Segment Director at Cray.

CSCS' new CS-Storm® is configured with 18 compute nodes, each with 8 NVIDIA® V100 GPUs and 2 Intel® Xeon® Gold 6134 CPUs, and includes two Cray ClusterStor® L300 storage systems. The CS-Storm was accepted in April 2019 and will become fully operational in 2020.

MeteoSwiss acquires 2nd GPU-based system for production weather forecasting.

New Cray CS-Storm has 18 nodes, each with 8 NVIDIA V100 GPUs and 2 Intel Xeon processors.

Challenges in Weather, Water and Climate

- Diversity and specialization in processors and the need for performance portability
- Integration of AI, analytics and traditional simulation workloads
- Rapidly increasing data volumes
- Integration with cloud to handle big data and maximize its value
- Maintaining and improving runtime reproducibility

Increasing Diversity in Technology

- End of Dennard scaling, architectural specialization to get performance, different processors optimized for different workloads

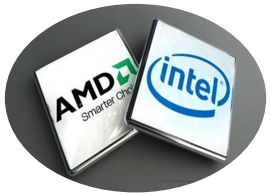

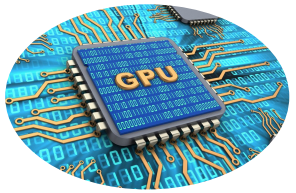

- New memory and storage technologies - Flash, non-volatile memory, HBM

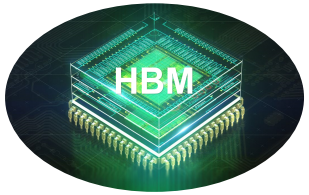

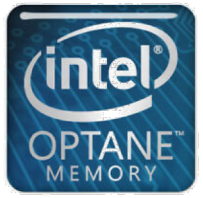

Increasingly heterogeneous, data-centric workloads

Diverse software stacks and usage styles

- Interactive computing
- High productivity interpreted languages
- New analytics and AI frameworks
- Containerized software
- On-demand cloud computing
- Object storage


CRAY SHASTA
CRAY

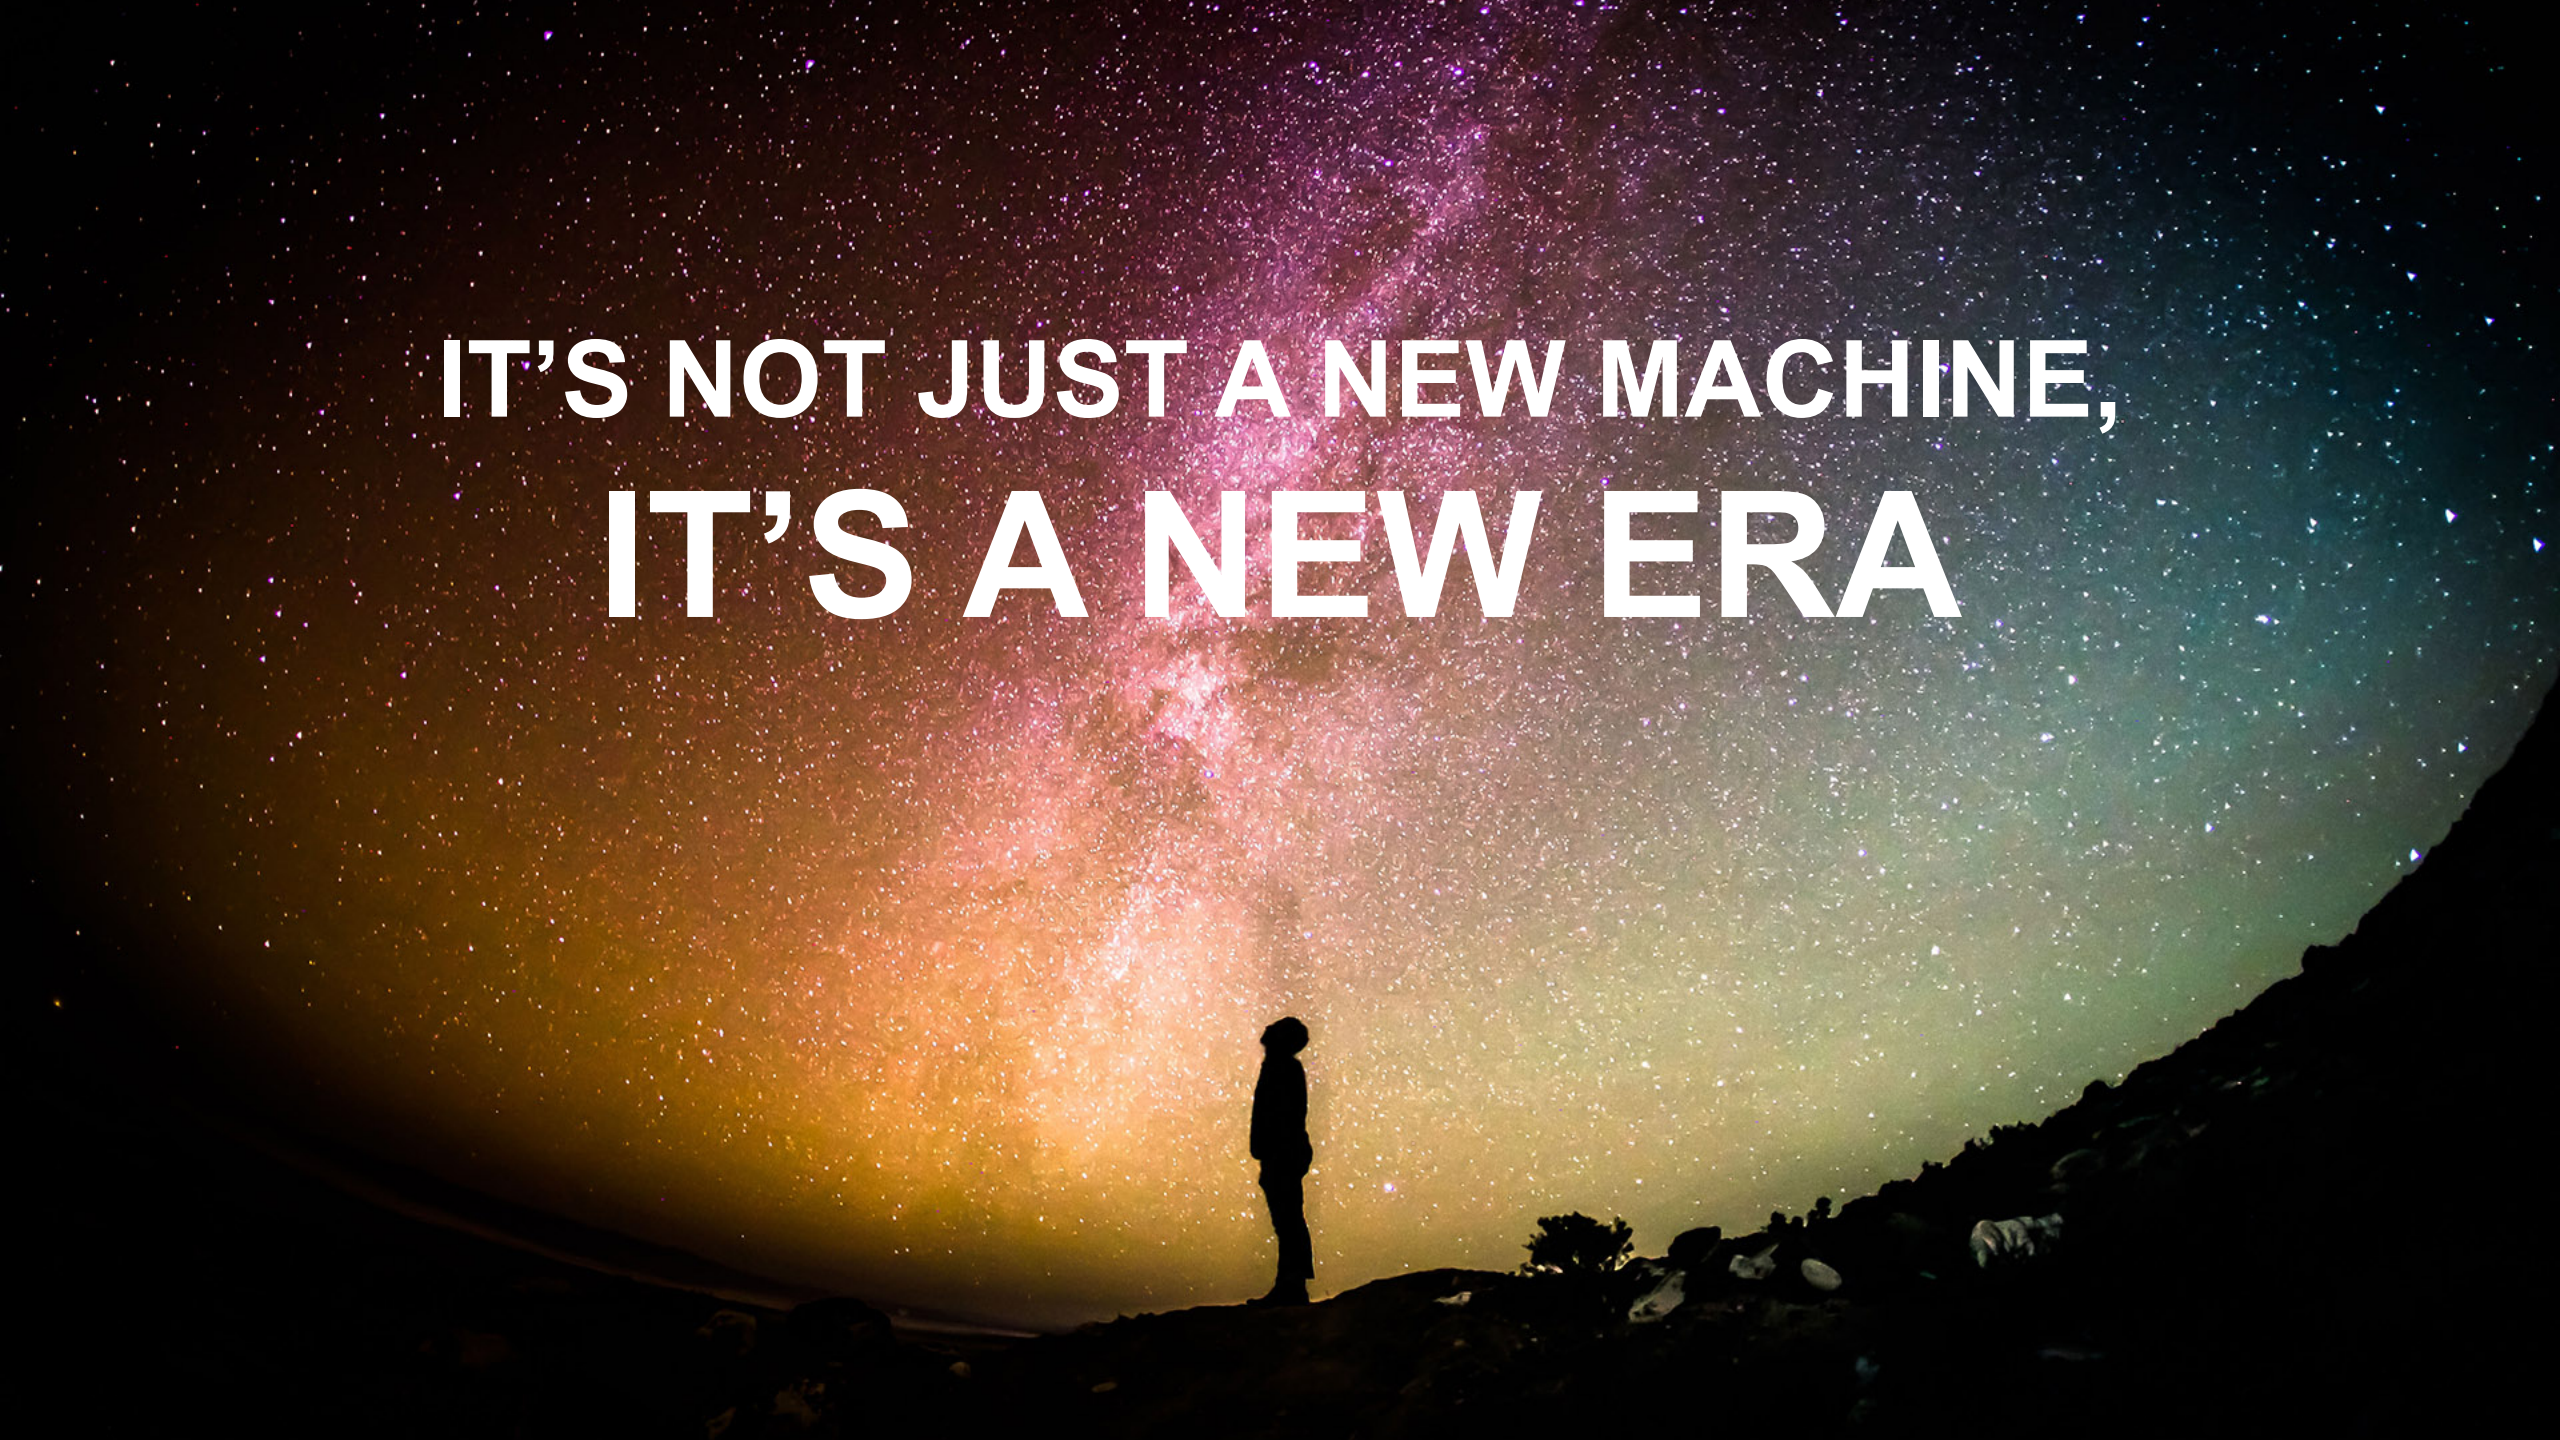
IT’S NOT JUST A NEW MACHINE,
IT’S A NEW ERA


MODELING &
SIMULATION
BIG DATA
ANALYTICS
ARTIFICIAL
INTELLIGENCE

Supercomputers are Critical to Simulation

Largest ever storm prediction model

- Over 4 billion points used to simulate the landfall of Hurricane Sandy
- Urban scale grid resolution of 500m (compared to standard 3km)
- Enables the research to understand fine grained properties of hurricanes

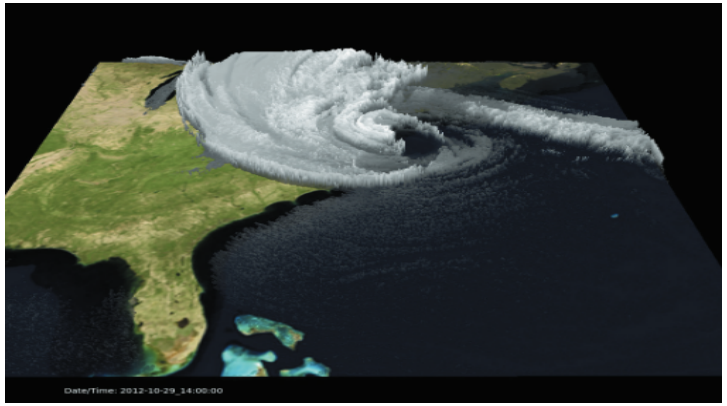

Key to the development of new antiretroviral drugs

- Determined the precise chemical structure of the HIV capsid – the protein shell that protects the virus's genetic material and is a key to its virulence.
- Requires the assembly of more than 1,300 identical proteins – in atomic-level detail.

Studying crop devastation by whiteflies to address a major cause of hunger in East Africa

- Understanding the DNA of the species by generating phylogenetic trees
- With only 500 whiteflies in a genetic dataset, the possible relationships between these flies run into the octillions (10^25)

...and to Machine and Deep Learning

Crop data is key to decision makers

- Applying Deep Learning on satellite data, the two major crops can be distinguished with 95% accuracy just a few months after planting and well before harvest.
- More timely estimates could be used for a variety applications, including supply-chain logistics, commodity market future projections, and more.

NCSA

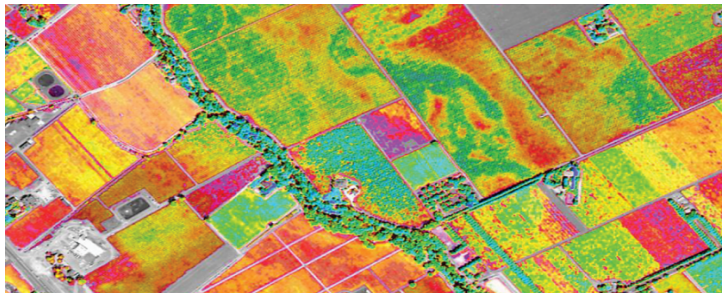

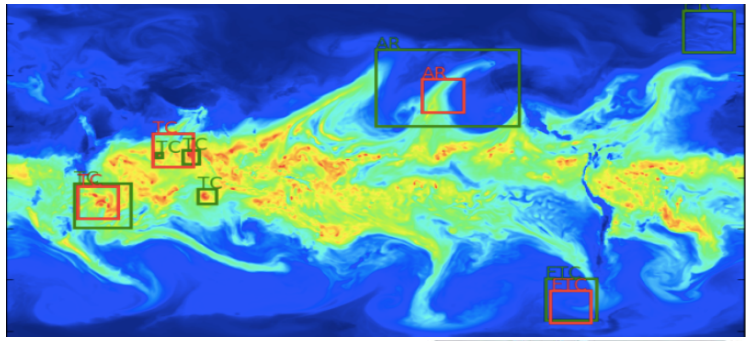

NERSC

Quantitatively assess how extreme weather will change in the future

- A single climate simulation can produce over 100TB of data with archives reaching over 5PB.
- Deep learning techniques are ideal for pattern recognition over large data sets

Development of systems for connected cars and autonomous technologies

- These advances could not have been realized without the application of deep learning to object detection in image and full motion video

SAMSUNG RESEARCH AMERICA

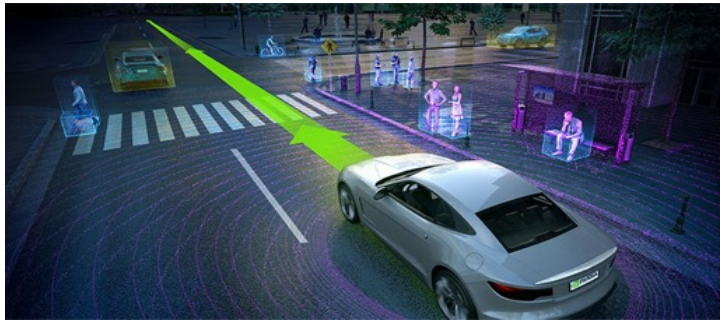

Blue Waters – world's largest geospatial system

Blue Waters will be the most powerful dedicated, non-classified geospatial system in the world.

The EarthDEM project between NCSA and NGA follows the successful ArcticDEM project (Polar Geospatial Center at Univ. of Minnesota and Ohio State Univ.).

- Goal: produce global, very high-resolution DEM data sets

NCSA to help build world's most powerful geospatial system

August 2, 2019 by staff Leave a Comment

NCSA has announced a new collaboration between the Blue Waters Project, the National Geospatial-Intelligence Agency (NGA), the University of Minnesota, and The Ohio State University to produce digital elevation models (DEM) of the entire Earth, among other geospatial research projects.

> *"We're very excited to announce this relationship with NGA," said NCSA Director William "Bill" Gropp. "This is the beginning of building a long-term strategic relationship between NGA and the University of Illinois, centered on high-performance computing and data analysis. The Blue Waters collaboration is the first of what we anticipate will be many years of research collaborations between NGA and Illinois faculty as well as NCSA."*

The collaboration, which is proceeding with an initial $11.1 million in funding from the NGA through the National Science Foundation (NSF) for one year, will make Blue Waters the most powerful dedicated, non-classified geospatial system in the world, capable of bringing unprecedented speed and efficiency to global mapping, and fundamentally changing the way humans view the Earth. This collaboration will be a continuation of the existing Blue Waters machine and will leverage the Blue Waters Project's domain expertise in optimizing workflows, applications, and complex data analysis, The Ohio State University's programming

Dept. of Energy next-generation systems

All Flash Storage System

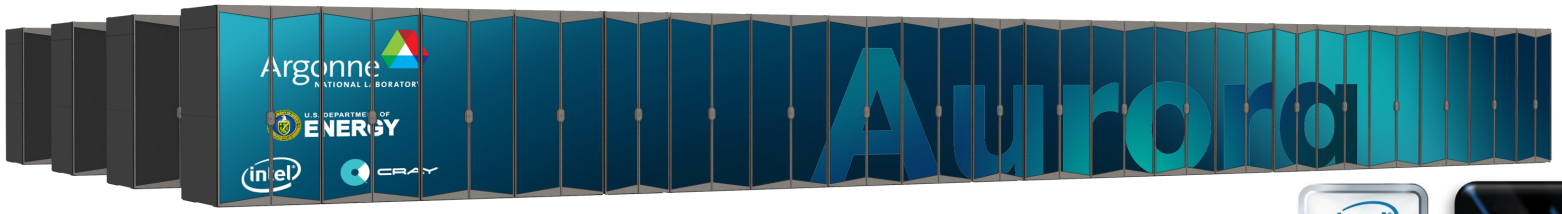

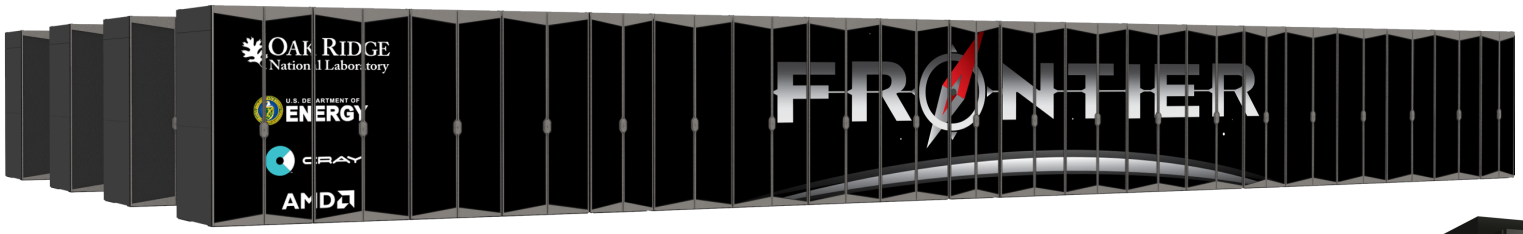

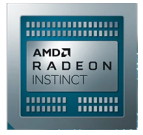

? ?

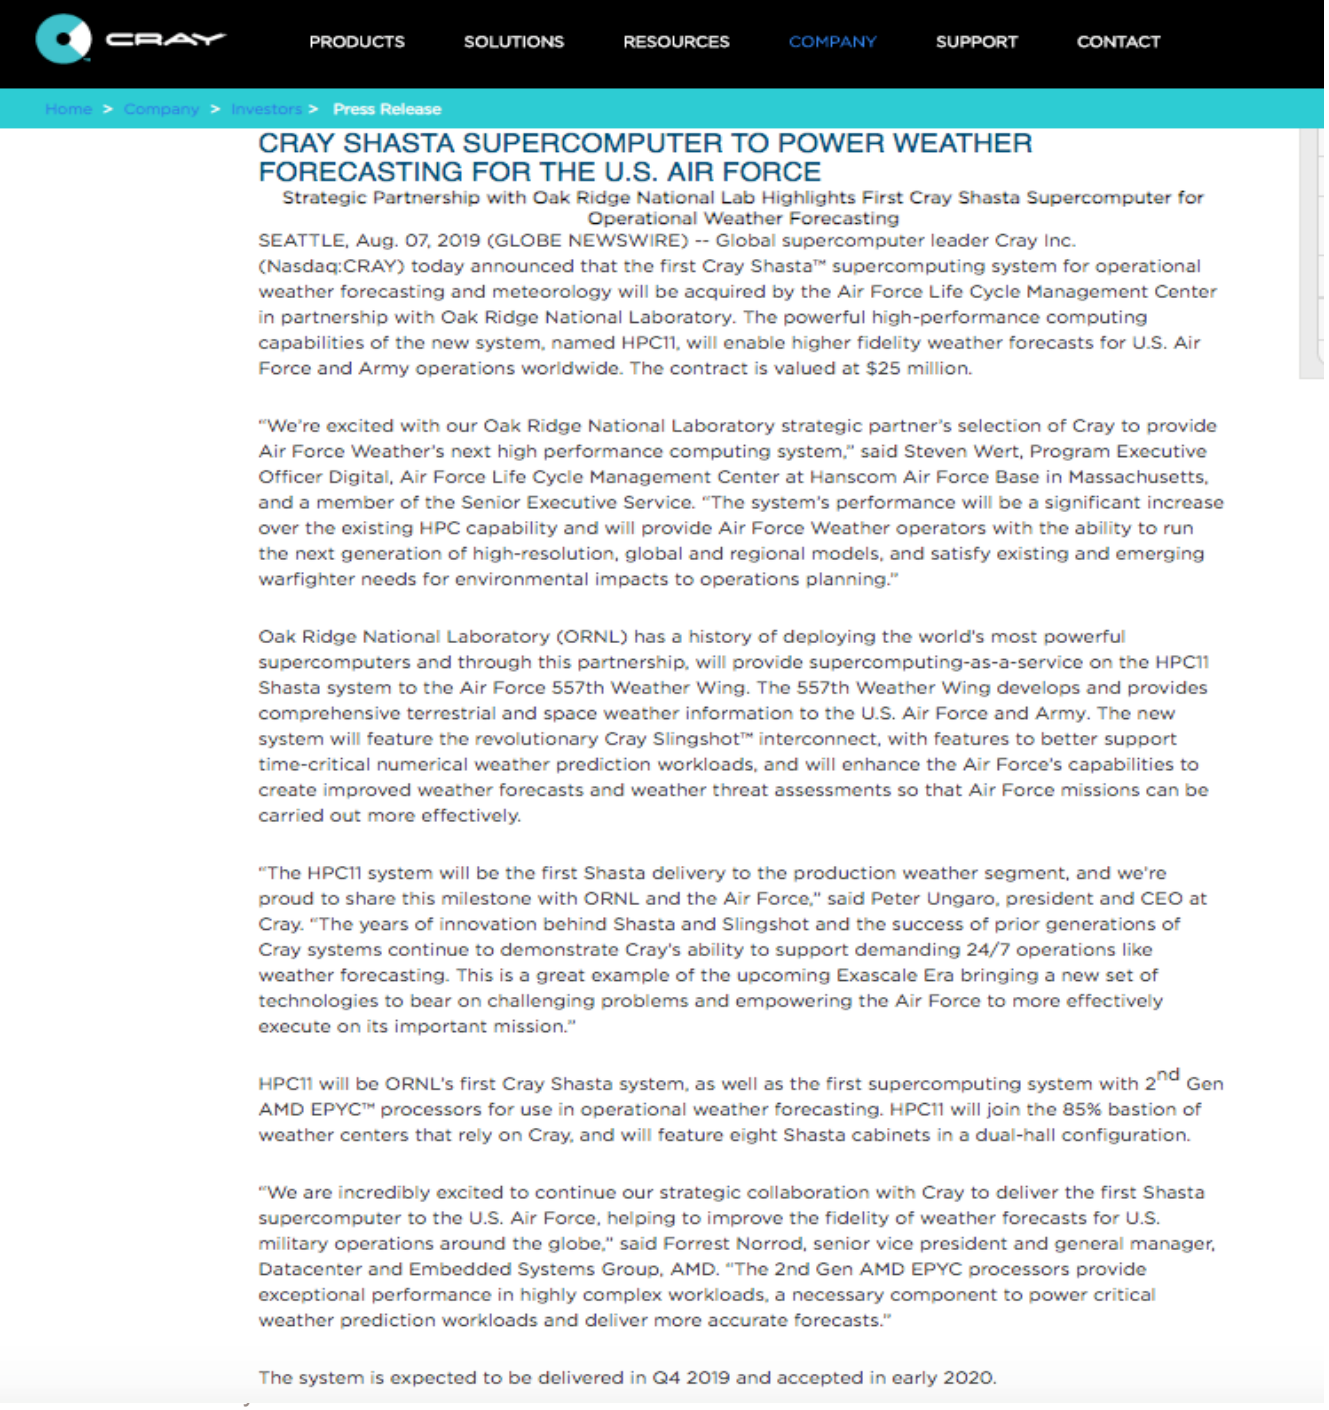

CRAY PRODUCTS SOLUTIONS RESOURCES COMPANY SUPPORT CONTACT

Home > Company > Investors > Press Release

CRAY SHASTA SUPERCOMPUTER TO POWER WEATHER FORECASTING FOR THE U.S. AIR FORCE

Strategic Partnership with Oak Ridge National Lab Highlights First Cray Shasta Supercomputer for Operational Weather Forecasting

SEATTLE, Aug. 07, 2019 (GLOBE NEWSWIRE) -- Global supercomputer leader Cray Inc. (Nasdaq:CRAY) today announced that the first Cray Shasta™ supercomputing system for operational weather forecasting and meteorology will be acquired by the Air Force Life Cycle Management Center in partnership with Oak Ridge National Laboratory. The powerful high-performance computing capabilities of the new system, named HPC11, will enable higher fidelity weather forecasts for U.S. Air Force and Army operations worldwide. The contract is valued at $25 million.

"We're excited with our Oak Ridge National Laboratory strategic partner's selection of Cray to provide Air Force Weather's next high performance computing system," said Steven Wert, Program Executive Officer Digital, Air Force Life Cycle Management Center at Hanscom Air Force Base in Massachusetts, and a member of the Senior Executive Service. "The system's performance will be a significant increase over the existing HPC capability and will provide Air Force Weather operators with the ability to run the next generation of high-resolution, global and regional models, and satisfy existing and emerging warfighter needs for environmental impacts to operations planning."

Oak Ridge National Laboratory (ORNL) has a history of deploying the world's most powerful supercomputers and through this partnership, will provide supercomputing-as-a-service on the HPC11 Shasta system to the Air Force 557th Weather Wing. The 557th Weather Wing develops and provides comprehensive terrestrial and space weather information to the U.S. Air Force and Army. The new system will feature the revolutionary Cray Slingshot™ interconnect, with features to better support time-critical numerical weather prediction workloads, and will enhance the Air Force's capabilities to create improved weather forecasts and weather threat assessments so that Air Force missions can be carried out more effectively.

"The HPC11 system will be the first Shasta delivery to the production weather segment, and we're proud to share this milestone with ORNL and the Air Force," said Peter Ungaro, president and CEO at Cray. "The years of innovation behind Shasta and Slingshot and the success of prior generations of Cray systems continue to demonstrate Cray's ability to support demanding 24/7 operations like weather forecasting. This is a great example of the upcoming Exascale Era bringing a new set of technologies to bear on challenging problems and empowering the Air Force to more effectively execute on its important mission."

HPC11 will be ORNL's first Cray Shasta system, as well as the first supercomputing system with 2nd Gen AMD EPYC™ processors for use in operational weather forecasting. HPC11 will join the 85% bastion of weather centers that rely on Cray, and will feature eight Shasta cabinets in a dual-hall configuration.

"We are incredibly excited to continue our strategic collaboration with Cray to deliver the first Shasta supercomputer to the U.S. Air Force, helping to improve the fidelity of weather forecasts for U.S. military operations around the globe," said Forrest Norrod, senior vice president and general manager, Datacenter and Embedded Systems Group, AMD. "The 2nd Gen AMD EPYC processors provide exceptional performance in highly complex workloads, a necessary component to power critical weather prediction workloads and deliver more accurate forecasts."

The system is expected to be delivered in Q4 2019 and accepted in early 2020.

CRAY

OAK RIDGE National Laboratory

ORNL will provide supercomputing-as-a-service on the HPC11 system to the Air Force 557th Weather Wing.

First Shasta system for production weather forecasting.

- Slingshot interconnect
- 2nd gen AMD EPYC processors
- 8 cabinets in a hall A/hall B configuration

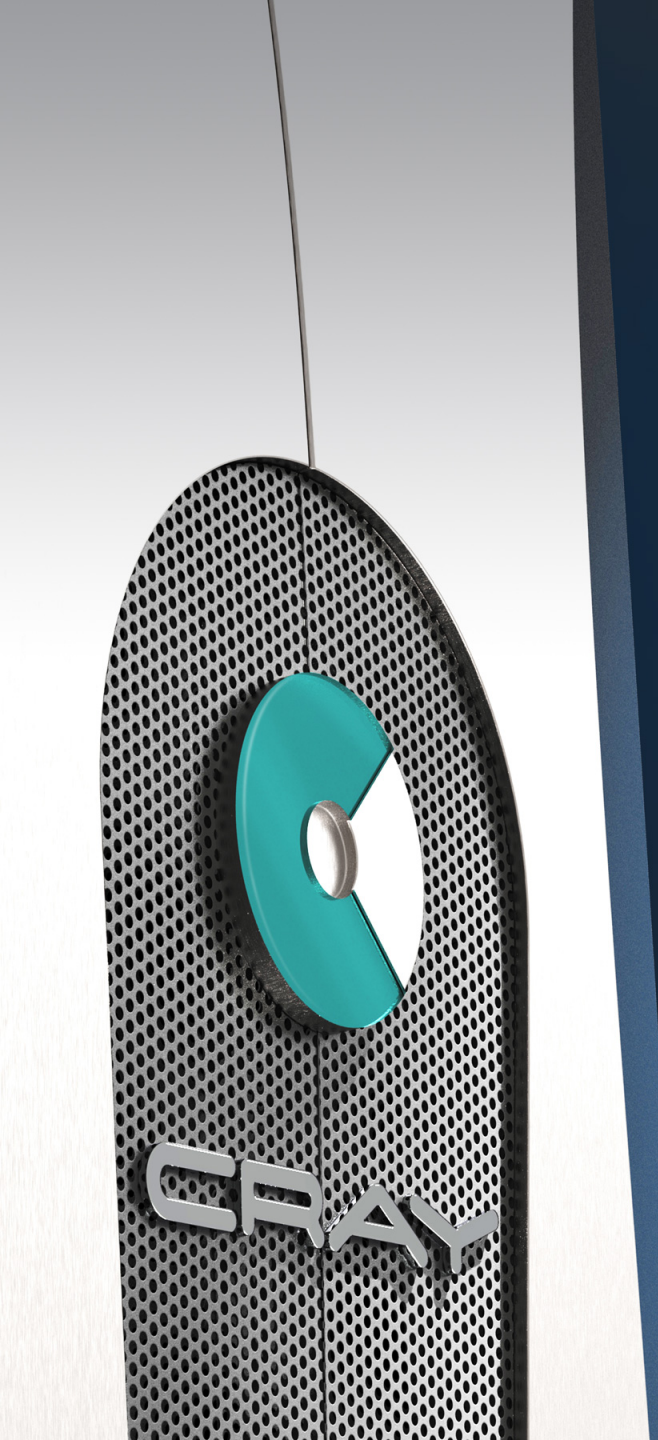

SHASTA HARDWARE

DIVERSITY OF PROCESSORS

DATA-CENTRIC APPROACH

RETHINKING INTERCONNECT

Slingshot Interconnect

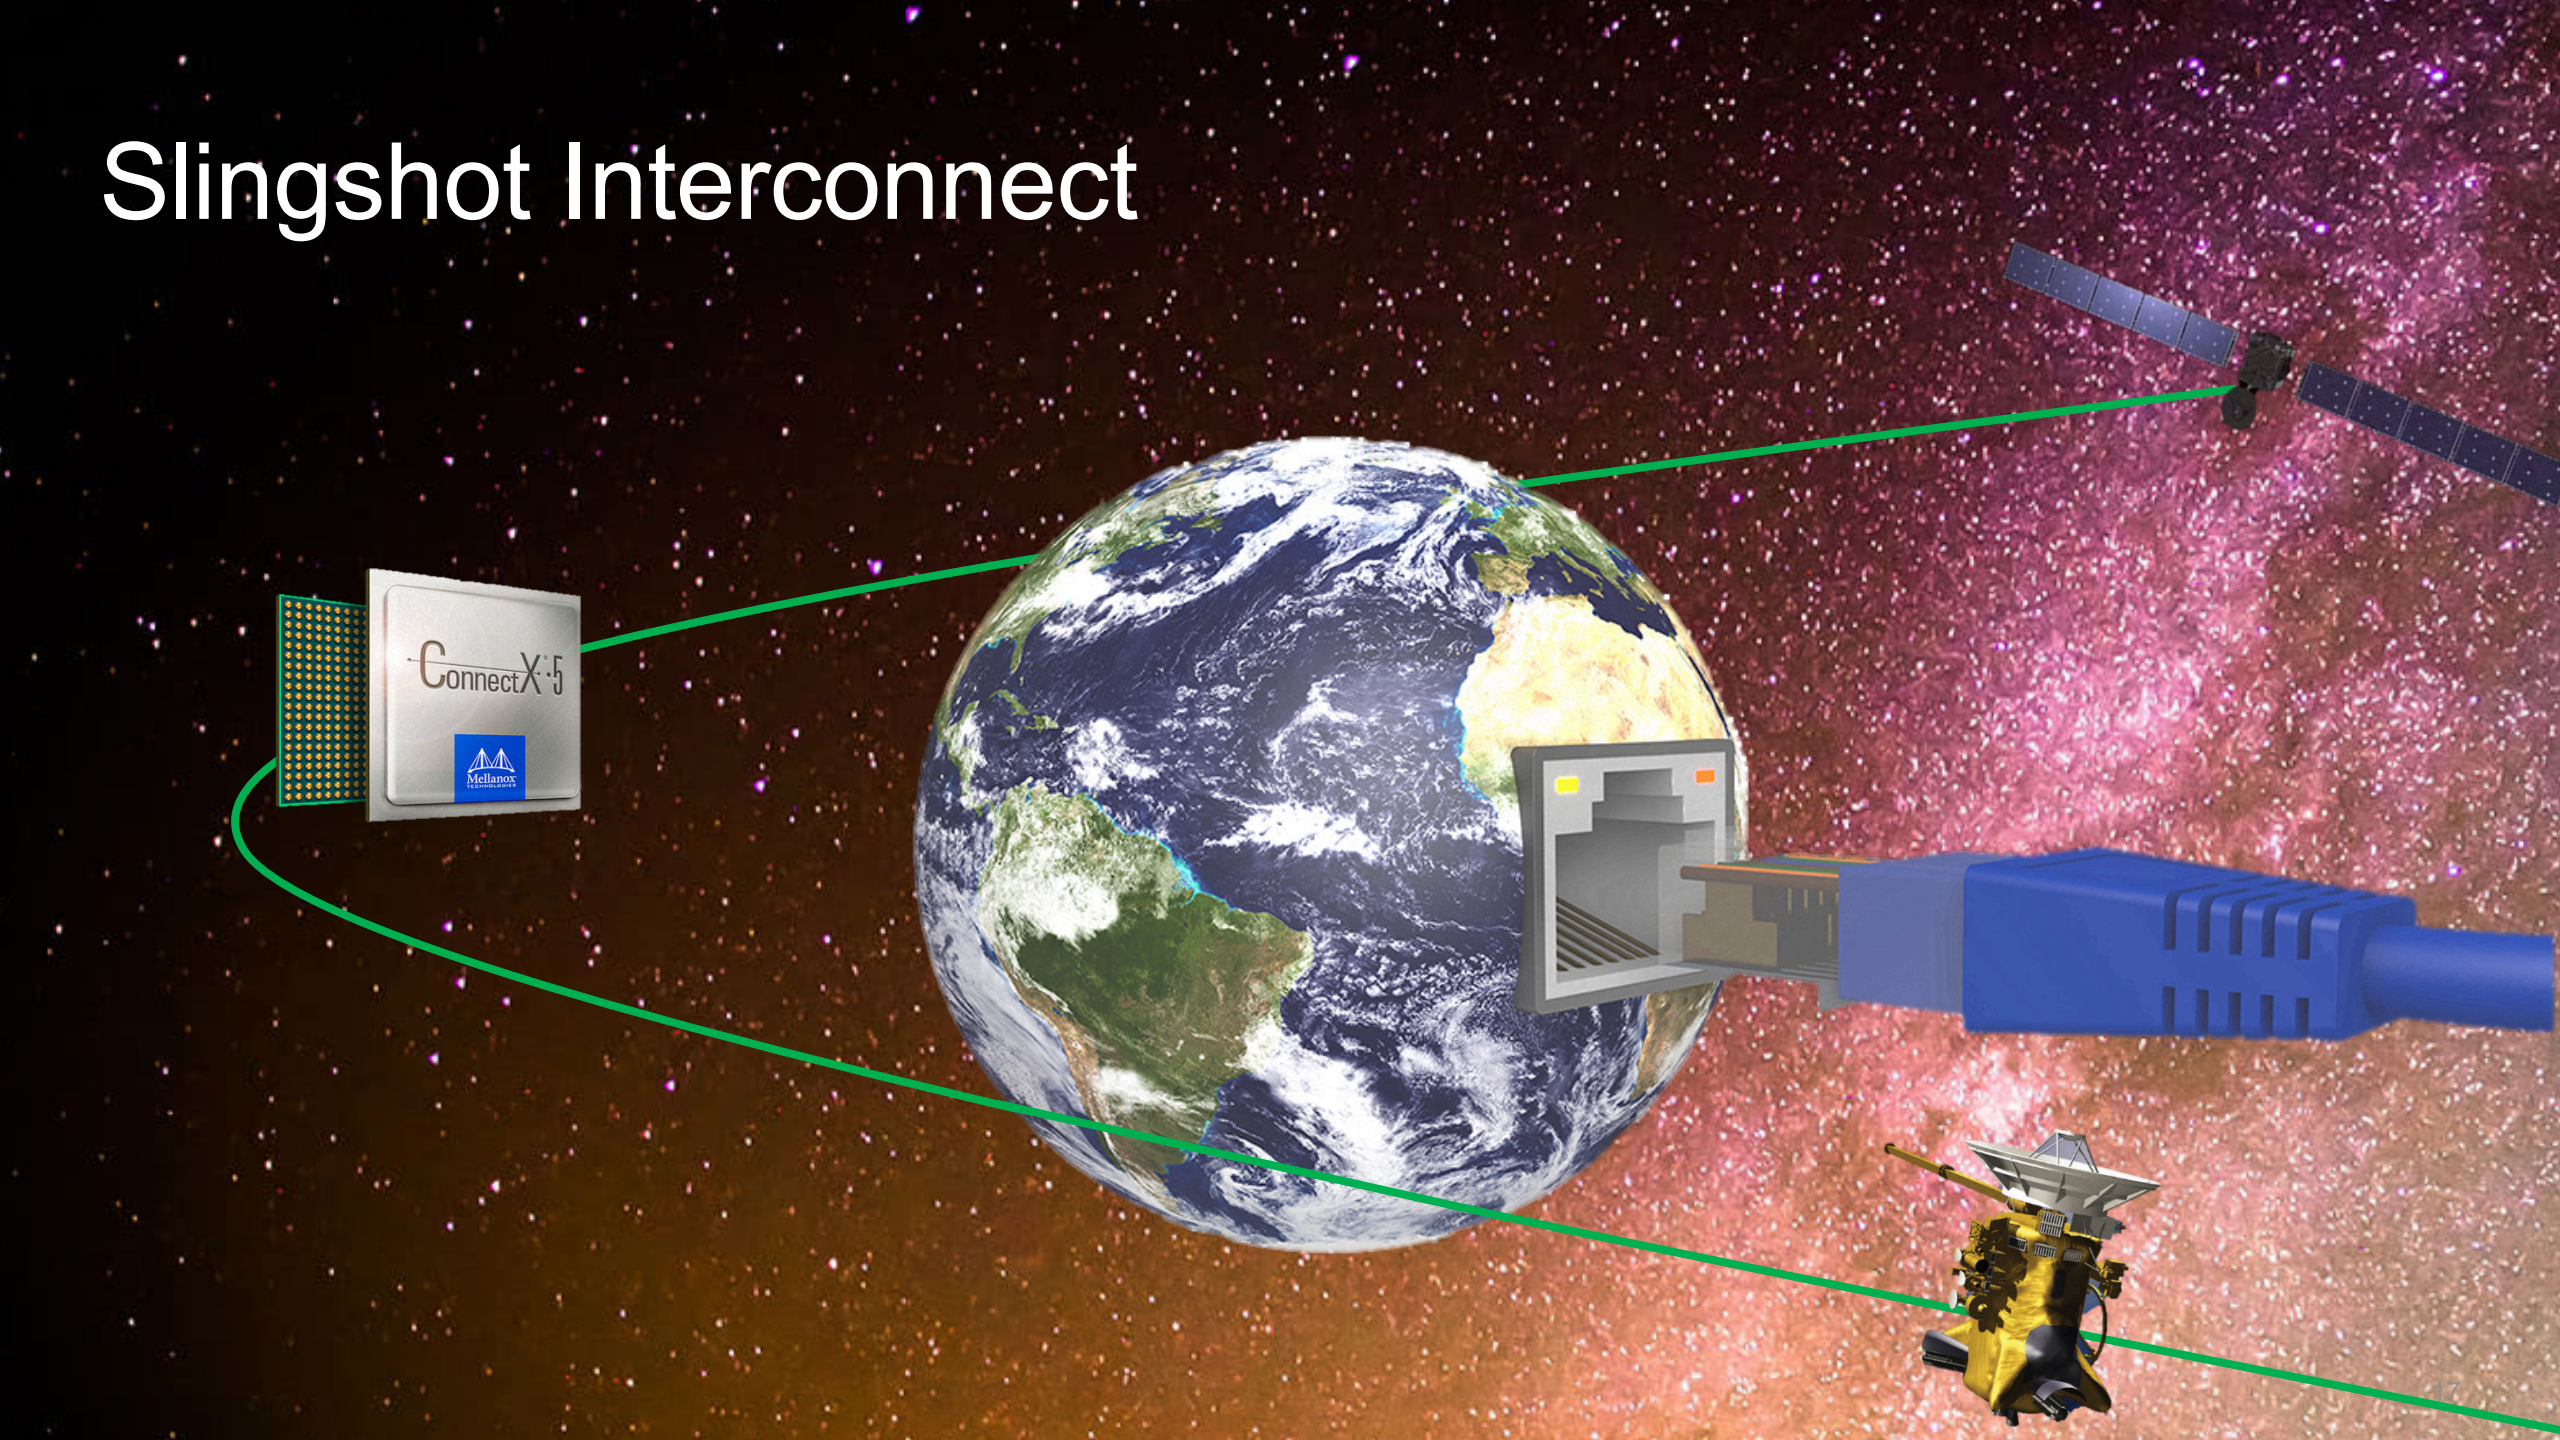

Slingshot – 6 Major Enhancements

Slingshot will be a **great interconnect** for our customers

This is Cray's **8th** supercomputing interconnect

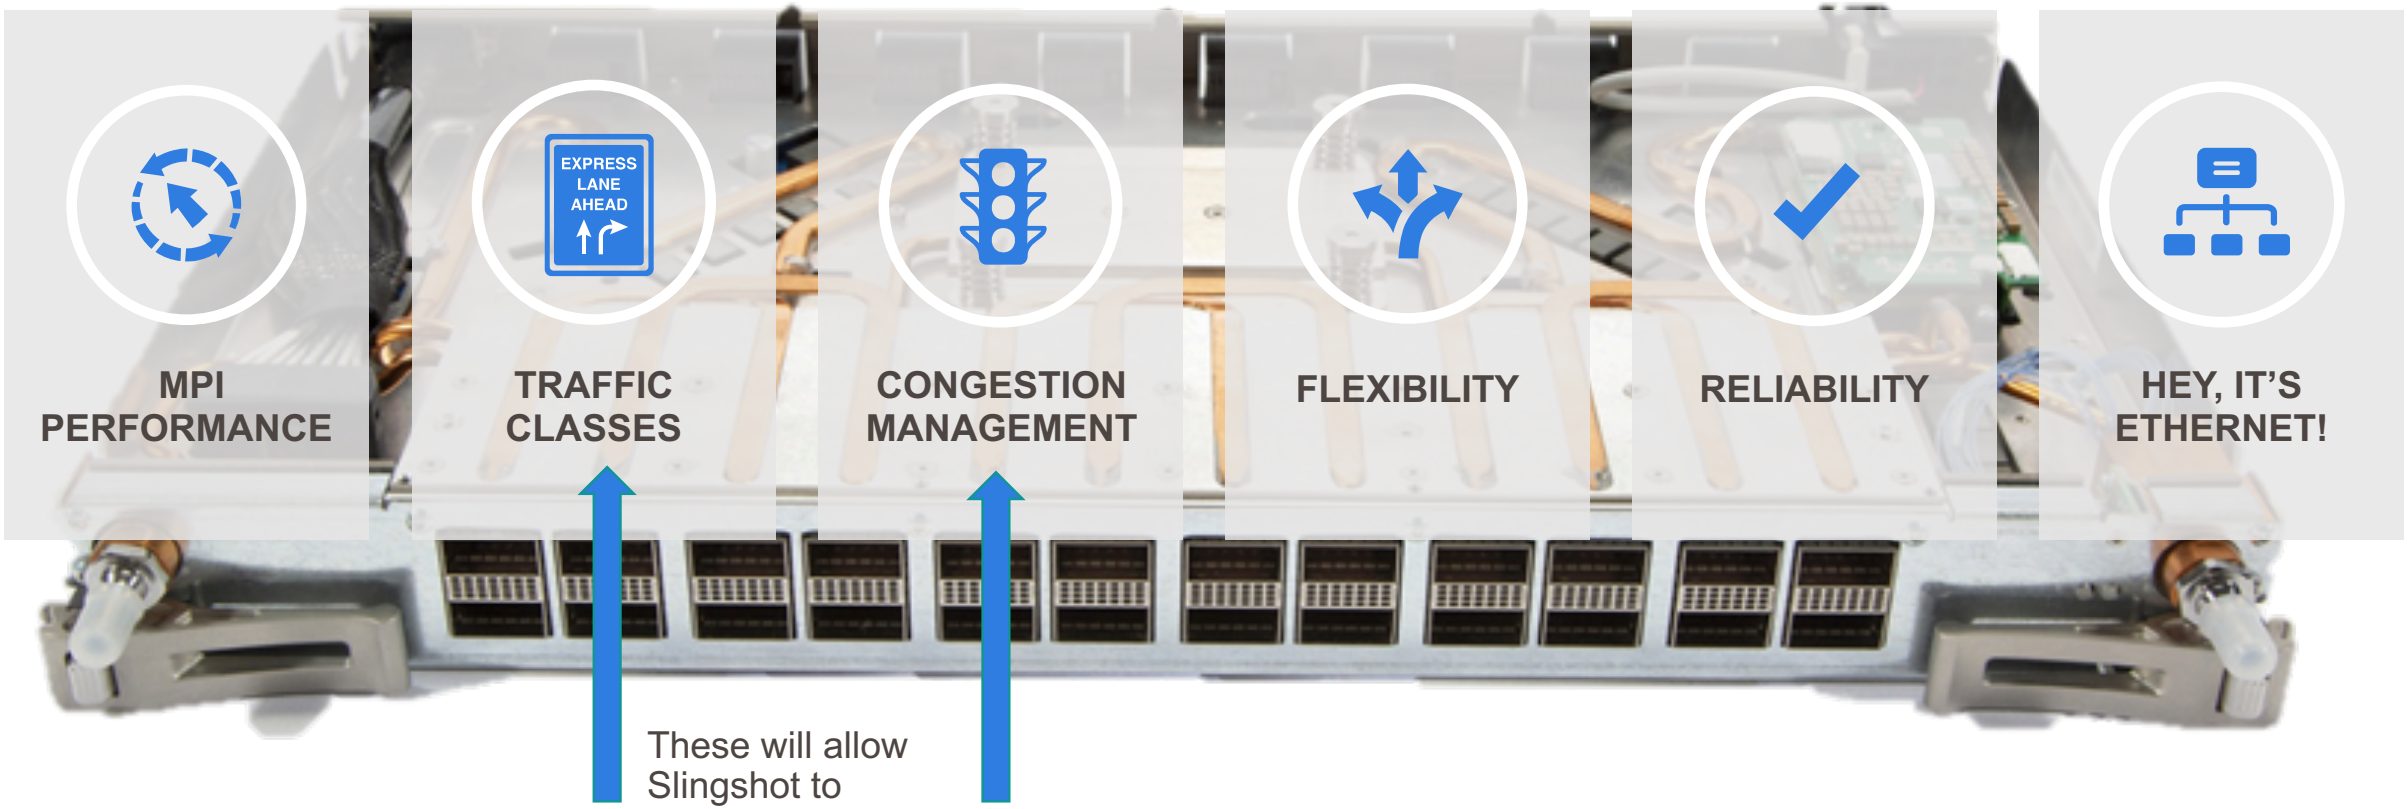

Slingshot Quality of Service Classes

Jobs Assigned to QoS classes

- Dedicated network resources for high priority jobs
- Limit bandwidth used by background jobs

Multiple QoS Classes with a Job

- Bulk data traffic in one class, low latency traffic in another
- My reductions don’t get stuck behind your large messages

I/O Traffic in its own QoS Class

- Assured bandwidth to the I/O system
- Less interference to compute network traffic

Slingshot Congestion Management

Determines the Causes of Congestion

- Pushes back on its source(s)
- Frees up buffer space for everyone else

Predictable Run Times

- Applications are much less vulnerable to other traffic patterns on the network
- Much tighter tail-latency distributions – a big benefit in apps with global synchronization

Interconnect Efficiency

- Simulations show that links on a loaded system can run at 90-95% of peak rates

All Existing Interconnects Can Suffer From Congestion

How the Met Office Solved a Weather Forecasting Runtime Scare

MARCH 9, 2018 | BY PAUL SELWOOD, MET OFFICE | LEAVE A COMMENT

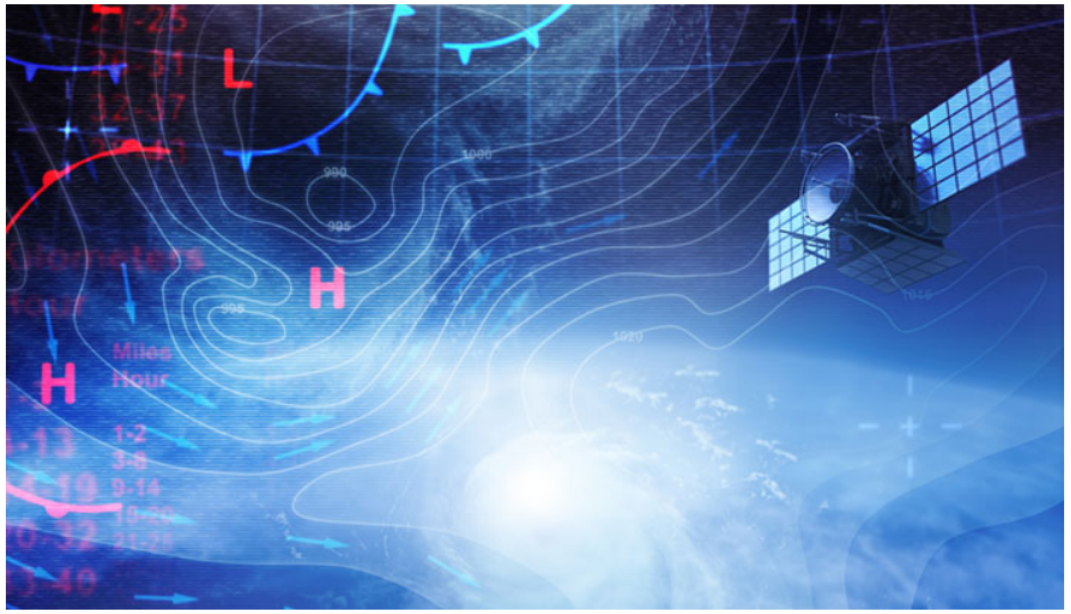

About Paul Selwood, Met Office

Paul Selwood is the Manager of HPC Optimisation at the Met Office, U.K.

- Team was increasing the UM resolution to 10km (N1280L70) from 12km (N1024L70)
- On brand new system, runtime was 55 minutes on 518 nodes but it was taking up to 74 minutes and it needs to run in 60 min or less.
- Analysis revealed that most time steps ran in ~1 second, *but some ran as long as 20 seconds*…
  - This was solved with some new LNET code sent out as an update to all XC customers
  - Eventually, the team got the 10km model running in 47 minutes

SHASTA FOR EARTH SCIENCES
MAJOR SLINGSHOT ENHANCEMENTS

Slingshot will be a **great interconnect** for earth sciences customers

This is Cray's **8th** supercomputing interconnect

MPI PERFORMANCE

TRAFFIC CLASSES

CONGESTION MANAGEMENT

FLEXIBILITY

RELIABILITY

ETHERNET COMPABILITY

Better data ingress/egress

SLINGSHOT – ETHERNET COMPATIBILITY

ETHERNET COMPABILITY

Connectivity

- Network attached storage can be connected directly to the interconnect
- Slingshot can be extended to the data center Ethernet network

Analytics workloads

- Many packages / workflows use IP socket-based communication

Ethernet protocols

HPC Ethernet protocol

- High performance and reliability

Reliability at Scale

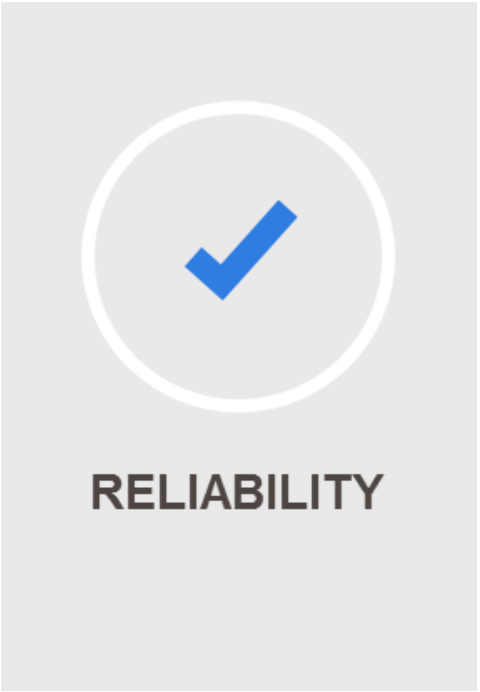

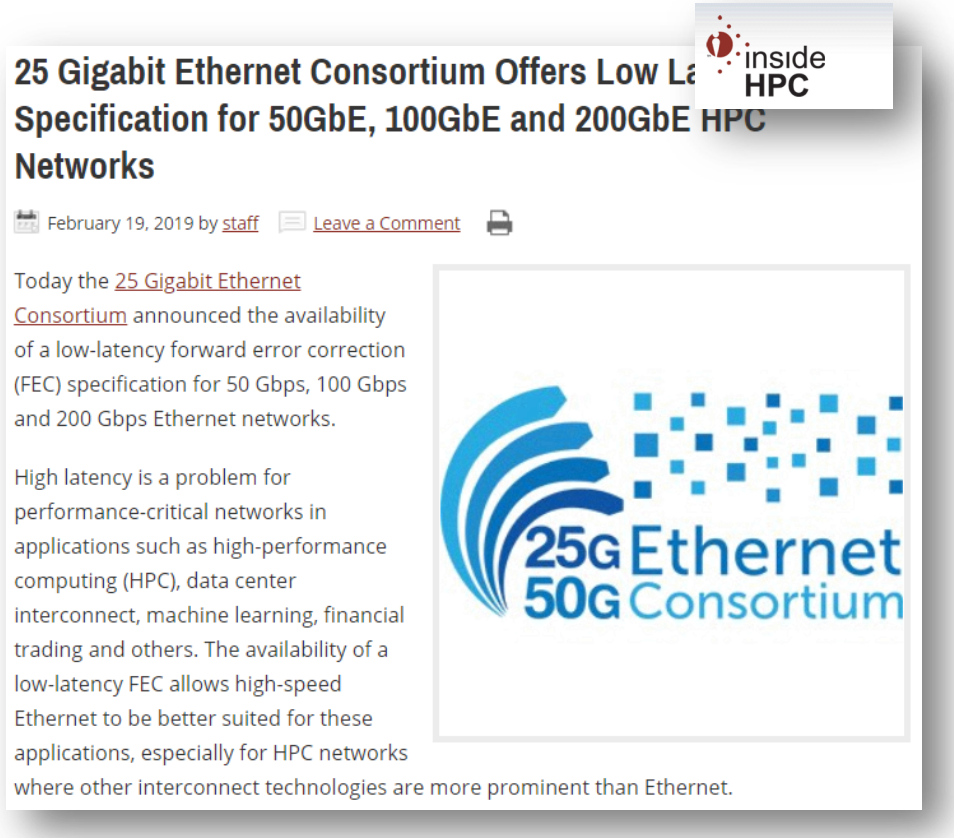

inside HPC

25 Gigabit Ethernet Consortium Offers Low La Specification for 50GbE, 100GbE and 200GbE HPC Networks

February 19, 2019 by staff Leave a Comment

Today the 25 Gigabit Ethernet Consortium announced the availability of a low-latency forward error correction (FEC) specification for 50 Gbps, 100 Gbps and 200 Gbps Ethernet networks.

High latency is a problem for performance-critical networks in applications such as high-performance computing (HPC), data center interconnect, machine learning, financial trading and others. The availability of a low-latency FEC allows high-speed Ethernet to be better suited for these applications, especially for HPC networks where other interconnect technologies are more prominent than Ethernet.

25G Ethernet 50G Consortium

Slingshot uses a new low-latency Forward Error Correction standard in addition to Link-Level Retry

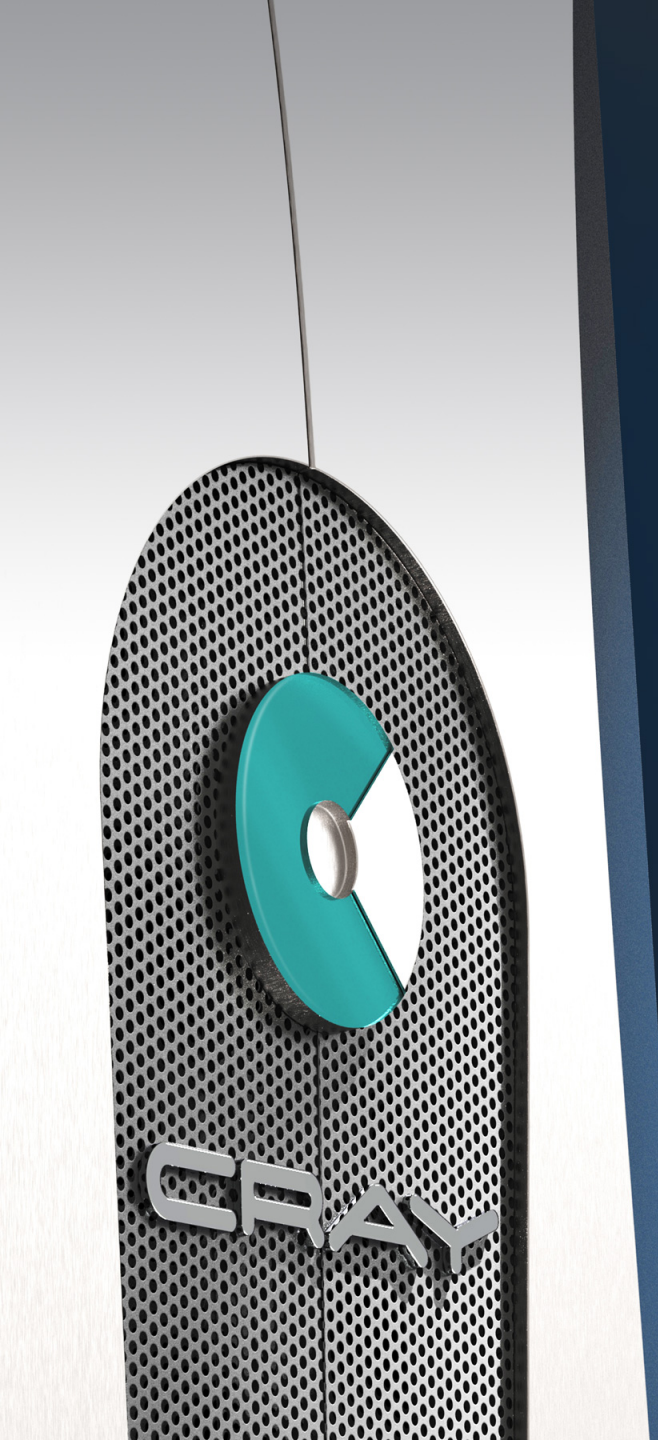

SHASTA SOFTWARE

CONVERGED WORKFLOWS

MULTI-TENANT CONTAINER BASED

EXTENSIBLE OPEN APIs

Integration of AI, Analytics and Simulation Software Challenges

Different software stacks

- Python data science
- TensorFlow, etc.
- Fortran/C/C++, MPI

Difficult to manage dependencies with environment modules

Container systems are being widely adopted

Different usage styles, different workload orchestration methods

- HPC WLM (batch queueing systems)
- Cloud-style (VMs, Kubernetes)

Containers for Flexibility

The ability to create your own world supports a diverse workload on a shared supercomputer

Build with Certified Software Stacks

Bundle libraries and dependencies

Build a consistent Environment from desktop to supercomputer

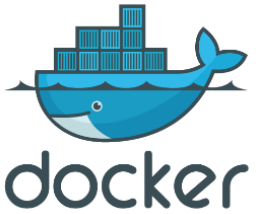

Multiple container systems with different orchestration/runtimes and different image formats

Singularity, Docker, Shifter, Sarus


PERFORMS LIKE A SUPERCOMPUTER
RUNS LIKE A CLOUD

CRAY SOFTWARE ARCHITECTURE EVOLUTION

Purpose built and optimized for performance and reliability		Extending Cray software capability with cloud flexibility
Foundational HPC batch workflow scheduling for modeling and simulation	⇨	Added converged HPC and AI workflows with addition of Kubernetes container orchestration
Support for partitioning system for HPC jobs	⇨	Added multi-tenancy between HPC and AI partitions, and sub-partitioning within AI jobs
Tightly integrated supercomputing system required coordination and downtime for upgrades	⇨	Modernized code to add containerized and self-healing services with separate compute and management planes to minimize downtime and allow for seamless upgrades
Tightly controlled system interfaces and limited access to system data driven by usage model	⇨	Added open APIs for integration, data access, and datacenter interoperability
Varied and discrete logs to track system and user activities and performance	⇨	Fully integrated system, application, and user level telemetry to quickly correlate and remediate issues

CRAY SHASTA SOFTWARE PLATFORM ARCHITECTURE

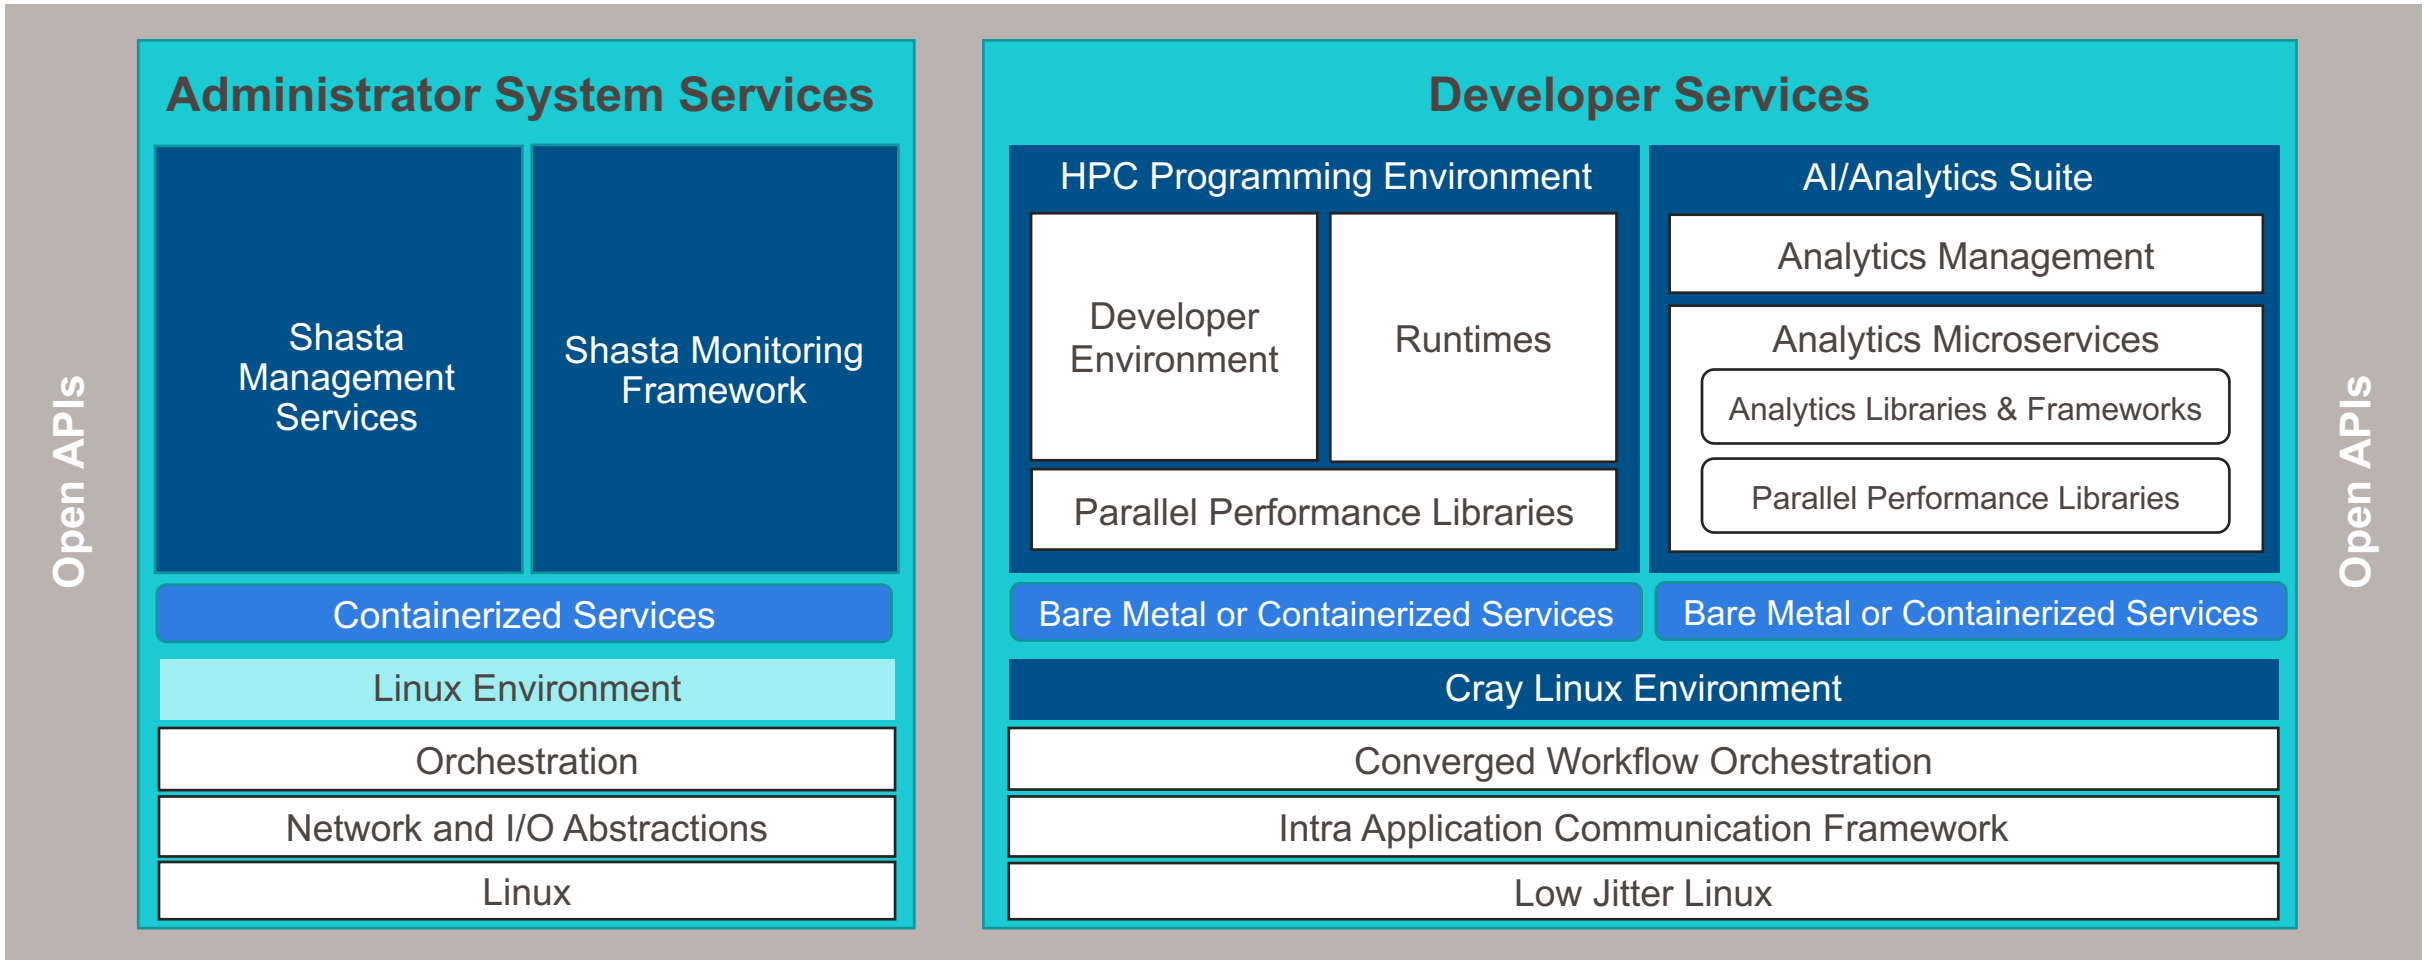

Expanding the power of supercomputing with the flexibility of cloud and full datacenter interoperability

Convergence of Supercomputing and Big Data

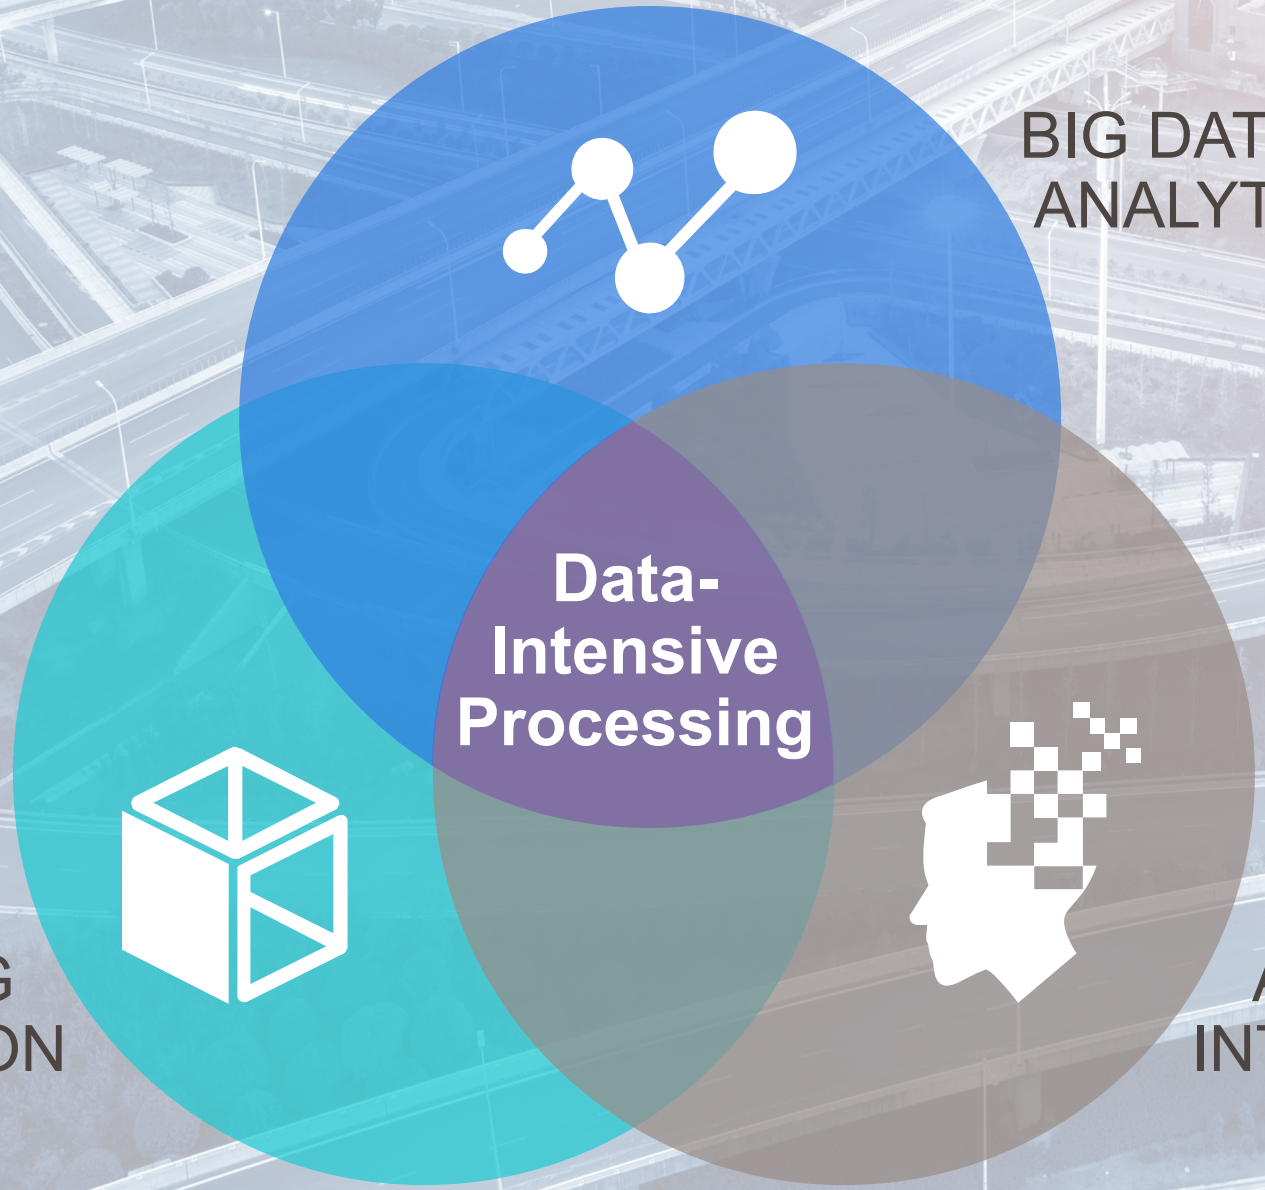

NEW SYSTEM ARCHITECTURE

NEW INTERCONNECT

NEW SOFTWARE PLATFORM

A NEW ERA OF COMPUTING:
BEYOND SUPER

3 EXASCALE WINS TOTALING OVER $1.3 BILLION

CRAY®