

# HPCとAnalyticsに向けた クレイ製品のご紹介

梅田宏明, Cray Japan Inc.

PCクラスタシンポジウム, 12/15/2016

COMPUTE

STORE

ANALYZE

12/15/2016

Copyright 2016 Cray Inc.

1

## Cray Systems for Deep Learning



### CS-Storm: Dense GPU Cluster Options

- 8 x M40 or P100 NVIDIA
- 512 GB – 1 TB of RAM, up to 6 SSDs
- M40 has 1.2 – 1.8x workload improvement
- P100 Memory Bandwidth 3x of M40
- Optimizations in CUDA not available with K40 or K80
- Building Block for Dense Deep Learning Compute Solution



### XC Scalable Deep Learning Supercomputer

- Worlds most scalable supercomputer for deep neural network training
- XC50 Features the Telsa P100 GPU accelerator for PCIe



COMPUTE


STORE

AN

12/15/2016

Copyright 2016 Cray Inc.

2



# Cray XC Series



**CRAY  
XC40**

**CRAY  
XC50**

COMPUTE

STORE

ANALYZE

12/15/2016

Copyright 2016 Cray Inc.

3

## The XC Series For Deep Learning




Cray Programming Environment



**NVIDIA® Tesla® P100 GPU  
Accelerator Compute Blade**


Extreme single  
precision performance,  
and large GPU-memory



High Performance Parallel Storage

High Performance Packet  
Switch Network

Scalable network able  
to handle neural  
network node-to-node  
communication



**NVIDIA® Tesla® K40 GPU  
Accelerator Compute Blade**

Combine large on-node  
memory and mini-batch  
processing to mid-  
performance GPU

For large-scale Deep Learning workflows where Data Parallelism is the preferred mode, the Cray XC Series offers exceptional deployment flexibility

COMPUTE

STORE

ANALYZE

12/15/2016

Copyright 2016 Cray Inc.

4

# Cray CS-Storm: Accelerated Computing



## Accelerated

Over 800 peak GPU teraflops in one 48U rack leveraging the power of **NVIDIA® Tesla®** GPU accelerators

## Integrated

**Integrated** hardware, software, packaging and cooling for reduced **TCO** and **production** reliability

## Reliable

Supercomputer **reliability, redundancy** and **serviceability** suitable for large systems, architected, built, delivered and serviced by Cray

COMPUTE | STORE | ANALYZE

Copyright 2016 Cray Inc.



12/15/2016

5

# CS-Storm – Dense Accelerated Compute



Six local SSDs

4+4 NVIDIA® Tesla® K40, K80, M40, and P100 GPUs

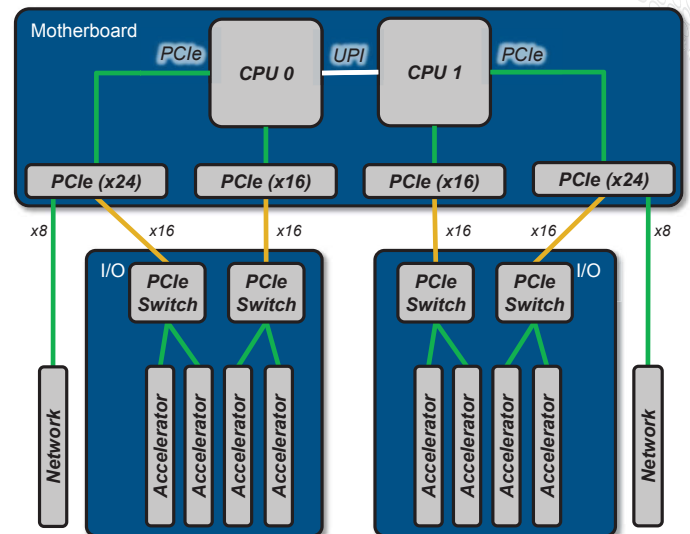
Dual Intel Xeon® host processors

16 DDR4 DIMM slots

Dense 2U chassis – 24" rack / air-cooled

### 48U rack

- 22 nodes
- 176 GPUs
- 850+ TF\*



COMPUTE | STORE | ANALYZE

Copyright 2016 Cray Inc.

\*176 Tesla P100 + 44 Intel Xeon processors

12/15/2016

6

# Cray Deep Learning Accelerators



## Validated Toolkits for Current Customers

- **Makes it easier & faster to move existing models to Cray platforms**
- **CS-Storm**
  - Downloadable Docker images & configuration guide for Deep Learning packages
    - All of the above
- **XC Series**
  - Makefiles and build documentation for Deep Learning packages
    - MS Cognitive Toolkit (CNTK) and TensorFlow

---

COMPUTE

STORE

ANALYZE

# Urika-GX Agile Analytics Platform

12/15/2016

Copyright 2016 Cray Inc.

8

**SERIOUS AGILITY** ...that's enterprise-accessible



## Example Software Layers

Dynamically repurpose resources and various analytics tools like these	kafka*	hadoop	spark	R	Processing Layer
	CRAY GRAPH ENGINE	hadoop HDFS	cassandra*	lustre™	Data Services Layer
	Apache MESOS™	YARN	openstack	whale*	Infrastructure Layer

\* Validated now, additional support to follow  
12/15/2016 Not included in Urika-GX software stack

COMPUTE

STORE

ANALYZE

Copyright 2016 Cray Inc.

# PERVASIVE SPEED


Supercomputing Experience



Up and Running in Days!  
Pre-integrated and validated software + hardware

- 42U single rack
- Cray Aries network
- 16/32/48 2-socket Intel® Xeon® E-5 2600 v4 family processor nodes
- Up to 1,728 cores per system
- Up to 22 TB DRAM
- 35-176 TB PCIe SSDs on-nodes
- 192 TB HDD local storage

- Enables large-scale memory-intensive workloads
- Node-local SSDs for deep local memory hierarchy
- Cray Aries fabric with high I/O throughput and low latency
- Options to attach external POSIX-compliant global storage:
  - Cray® Sonexion® storage (Lustre®)
  - Other existing file systems (NFS, GPFS)



Deep local memory hierarchy accelerates performance

ANALYZE


SERIOUS AGILITY

PERVASIVE SPEED

# HIGH FREQUENCY INSIGHTS



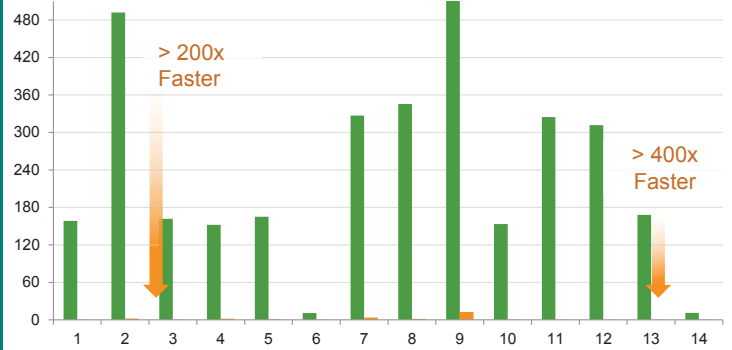
- System is tuned for demanding analytics
- Dynamic workflows for comparison and adaptation
- Cray Graph Engine for fast, complex, pattern matching
  - In-memory semantic graph database based on W3C standards such as RDF & SPARQL
  - Tuned for serious speed on noisy data with complex relationships



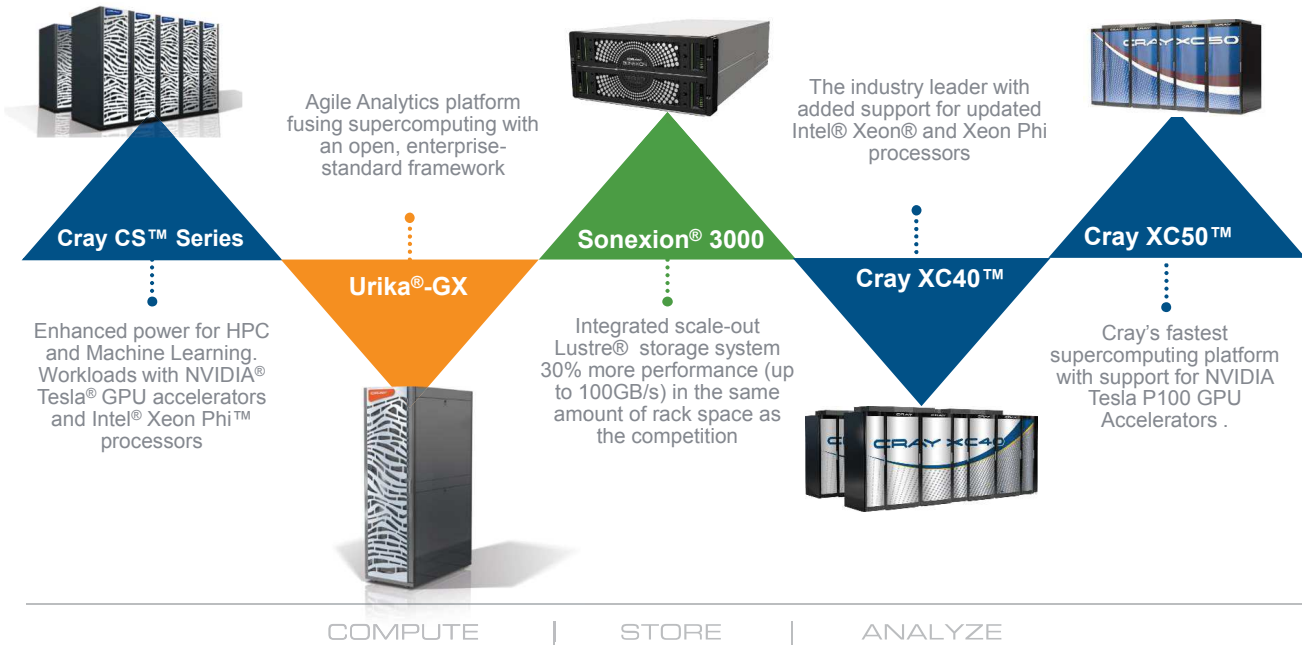
3 X Faster than prior Urika generation on complex queries

## Cray Graph Engine Outperforms Average Test Speedup of 280X

LUBM25K with >3 Billion Triples  
Spark GraphX vs. Cray Graph Engine  
Results in Seconds on Urika-GX system



# The Cray Portfolio



12/15/2016

Copyright 2016 Cray Inc.

# 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 publicly 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, and URIKA. 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.

COMPUTE | STORE | ANALYZE

12/15/2016

Copyright 2016 Cray Inc.