Three Issues for Global e-Infrastructure From the Japanese HPCI Viewpoint

Satoshi Matsuoka, Dr. Sci.

Professor, Global Scientific Information and Computing Center (GSIC)

Tokyo Institute of Technology, Japan

ICRI 2012 presentation, Mar 22, 2012, Copenhagen

Three Issues in Global e-Infrastructures

- *How do the e-Infrastructures coordinate globally, for real?*
  - Membership, Governance, Operations, Support, Training/Education, Legal Issues, updates, etc. etc.
- *How will the global e-Infrastructures advance towards Exascale?*
  - Not just FLOPS but also Bytes, bps, …
- *How do we pay for all this?*
  - Convincing story for the world citizens in this tough economic climate

HPCI-Japan: A National Supercomputing Infrastructure

SINET4: Science Information NETwork 4

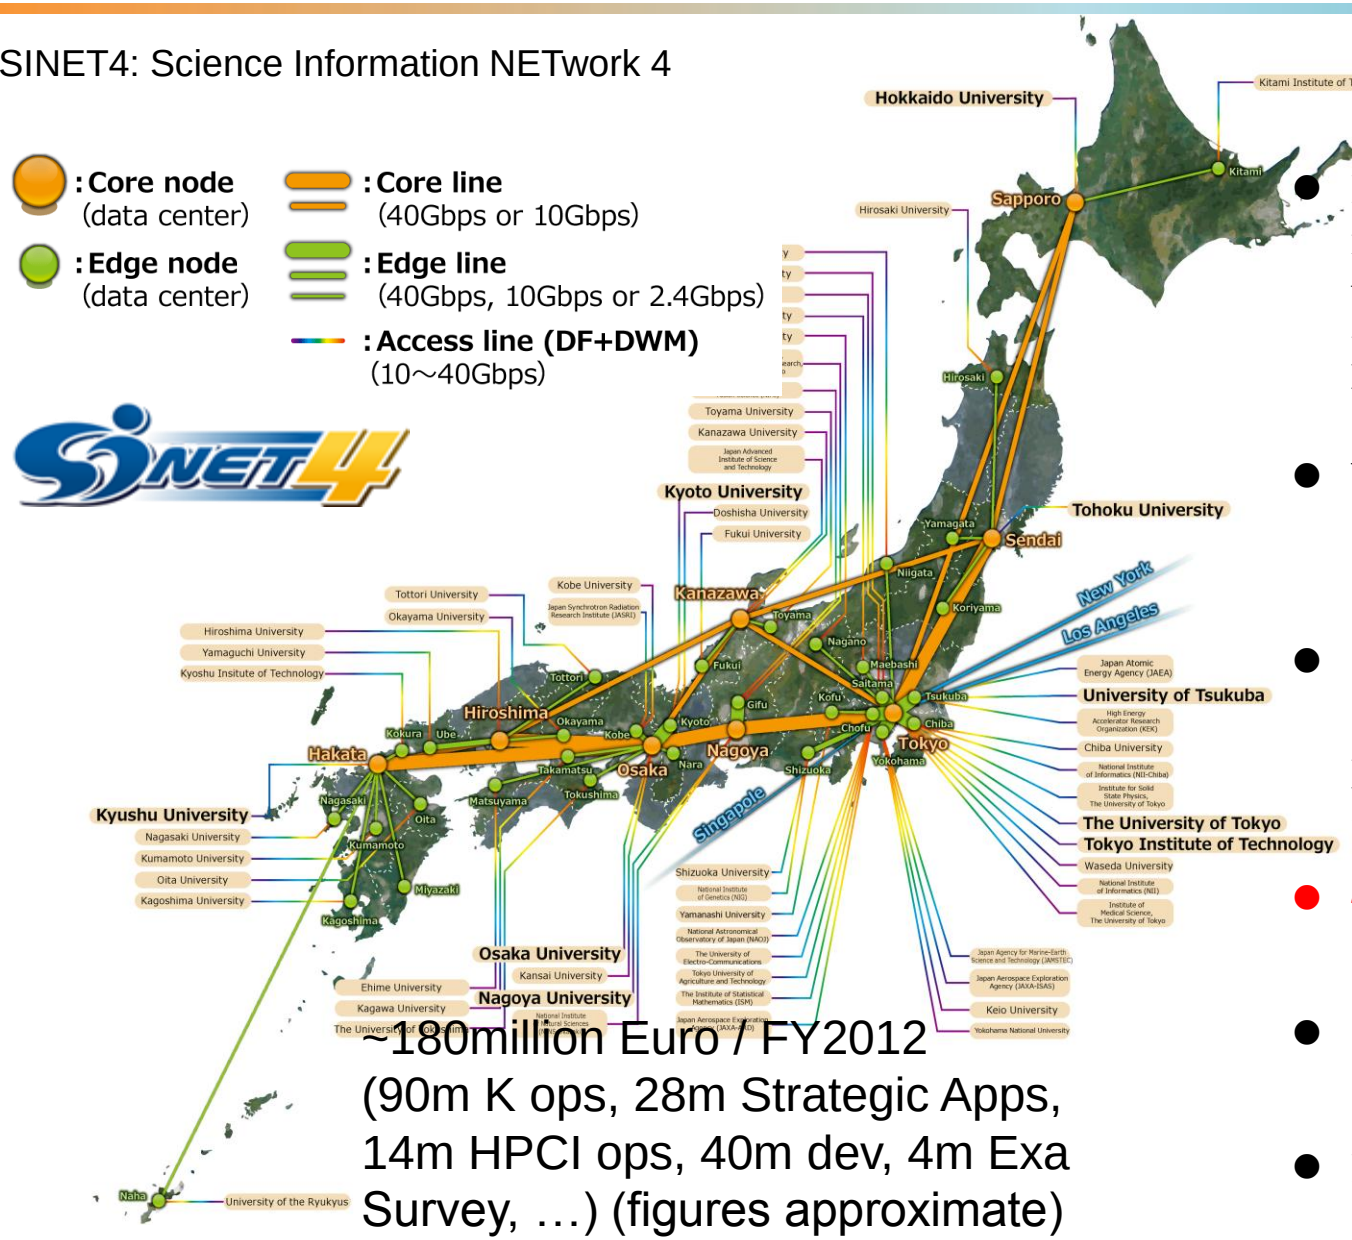

~180million Euro / FY2012
(90m K ops, 28m Strategic Apps, 14m HPCI ops, 40m dev, 4m Exa Survey, …) (figures approximate)

- HPCI is an e-Infrastructure of national supercomputing centers, with the K-Computer as the focal point, in order to cope with varying user needs.
- Will be operational from Spring 2012. Partial use of NAREGI and other Grid/Cloud MW,
- Official govt. consortium for extensive definition of operational procedures, governance, support, documents, user education, ...
- **~100PB Dedicated Shared Storage & Processing**
- Counterpart to PRACE, XSEDE, …
- Persist & Work towards Exascale

HPCI Strategic Application Areas

Life science/Drug manufacture

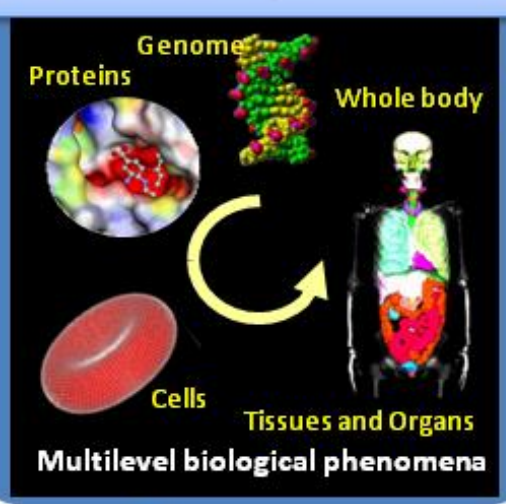

New Material and Energy Creation

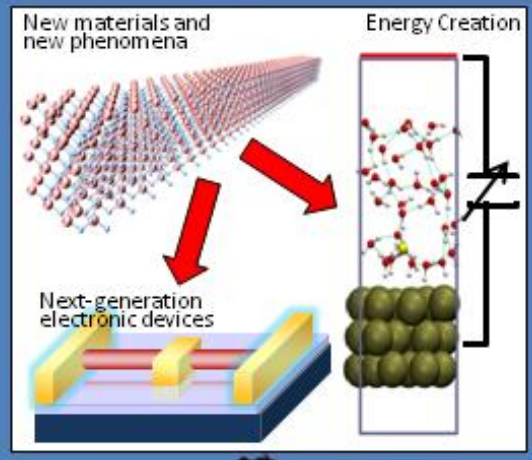

Global change prediction for disaster prevention/reduction

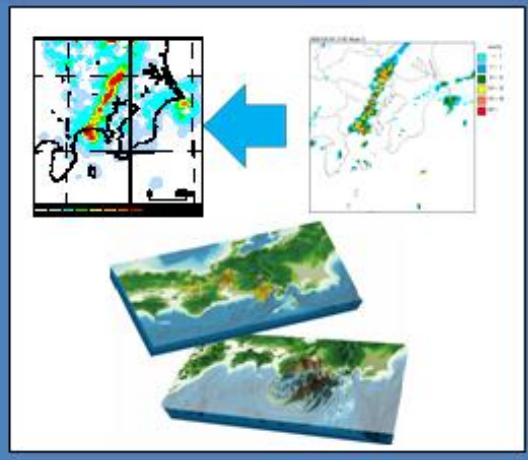

MONOZUKURI (Industrial Innovation)

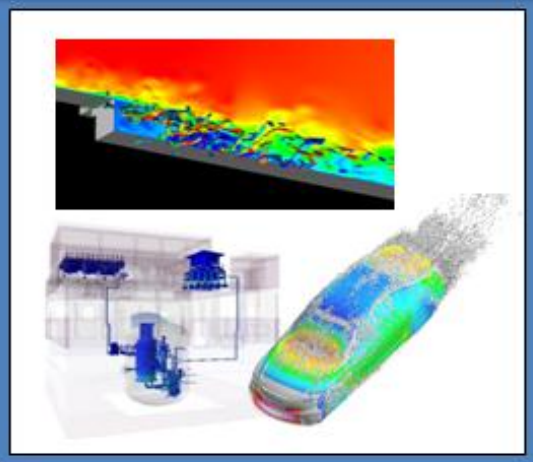

京

K computer

RIKEN
Advanced Institute for
Computational Science

The origin of matter and the universe

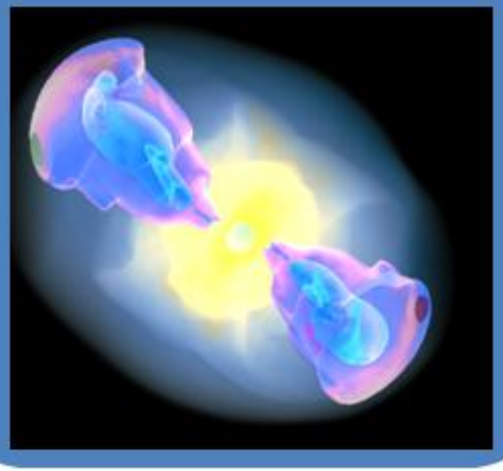

HPCI Centers: Riken AICS / K-Computerand Supercomputer Centers in Japanese University Centers

Approx 16.9PFlops total
+ National Labs SCs
pending (2~3PFlops)

AICS, RIKEN :
K computer (11.3 PFlops, 1.4PB)
Available in June 2012

Hokkaido Univ. :
SR16000(172Tflops, 22TB)
X86 Cloud (40Tflops, 12TB)

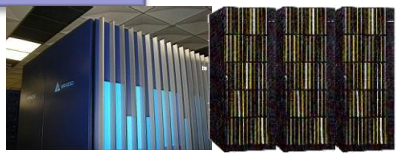

Kyoto Univ.
Appro Cluster (250 Tflops, 60TB)
Cray XE6 (300TFlops, 60TB)

Tohoku Univ. :
NEC SX-9(29.4Tflops, 18TB)
NEC Express5800 (1.74Tflops, 3TB)

Osaka Univ. :
SX-9 (16Tflops, 10TB)
SX-8R (5.3Tflops, 3.3TB)
PCCluster (23.3Tflops, 2.9TB)

Univ. of Tsukuba :
HA-PACS (800TFlops, 40TB)
T2K Open Supercomputer
(95.4Tflops, 20TB)

Kyushu Univ. :
PC Cluster (55Tflops, 18.8TB)
SR16000 L2 (25.3Tflops, 5.5TB)
PC Cluster (18.4Tflops, 3TB)

Univ. of Tokyo :
FX-10 (1.1 Pfflops, 150TB)
T2K Open Supercomputer
(140TFlops, 31.25TB)
SR16000 (50TFlops, 11TB)

Nagoya Univ. :
FX1(30.72Tflops, 24TB)
HX600(25.6Tflops, 10TB)
M9000(3.84Tflops, 3TB)

Tokyo Institute of Technology
Tsubame 2.0
(2.4 PFlops, 100TB)

Japanese HPCI AAAI, Globally Shared Storage, User Management/Support and Academic Cloud System

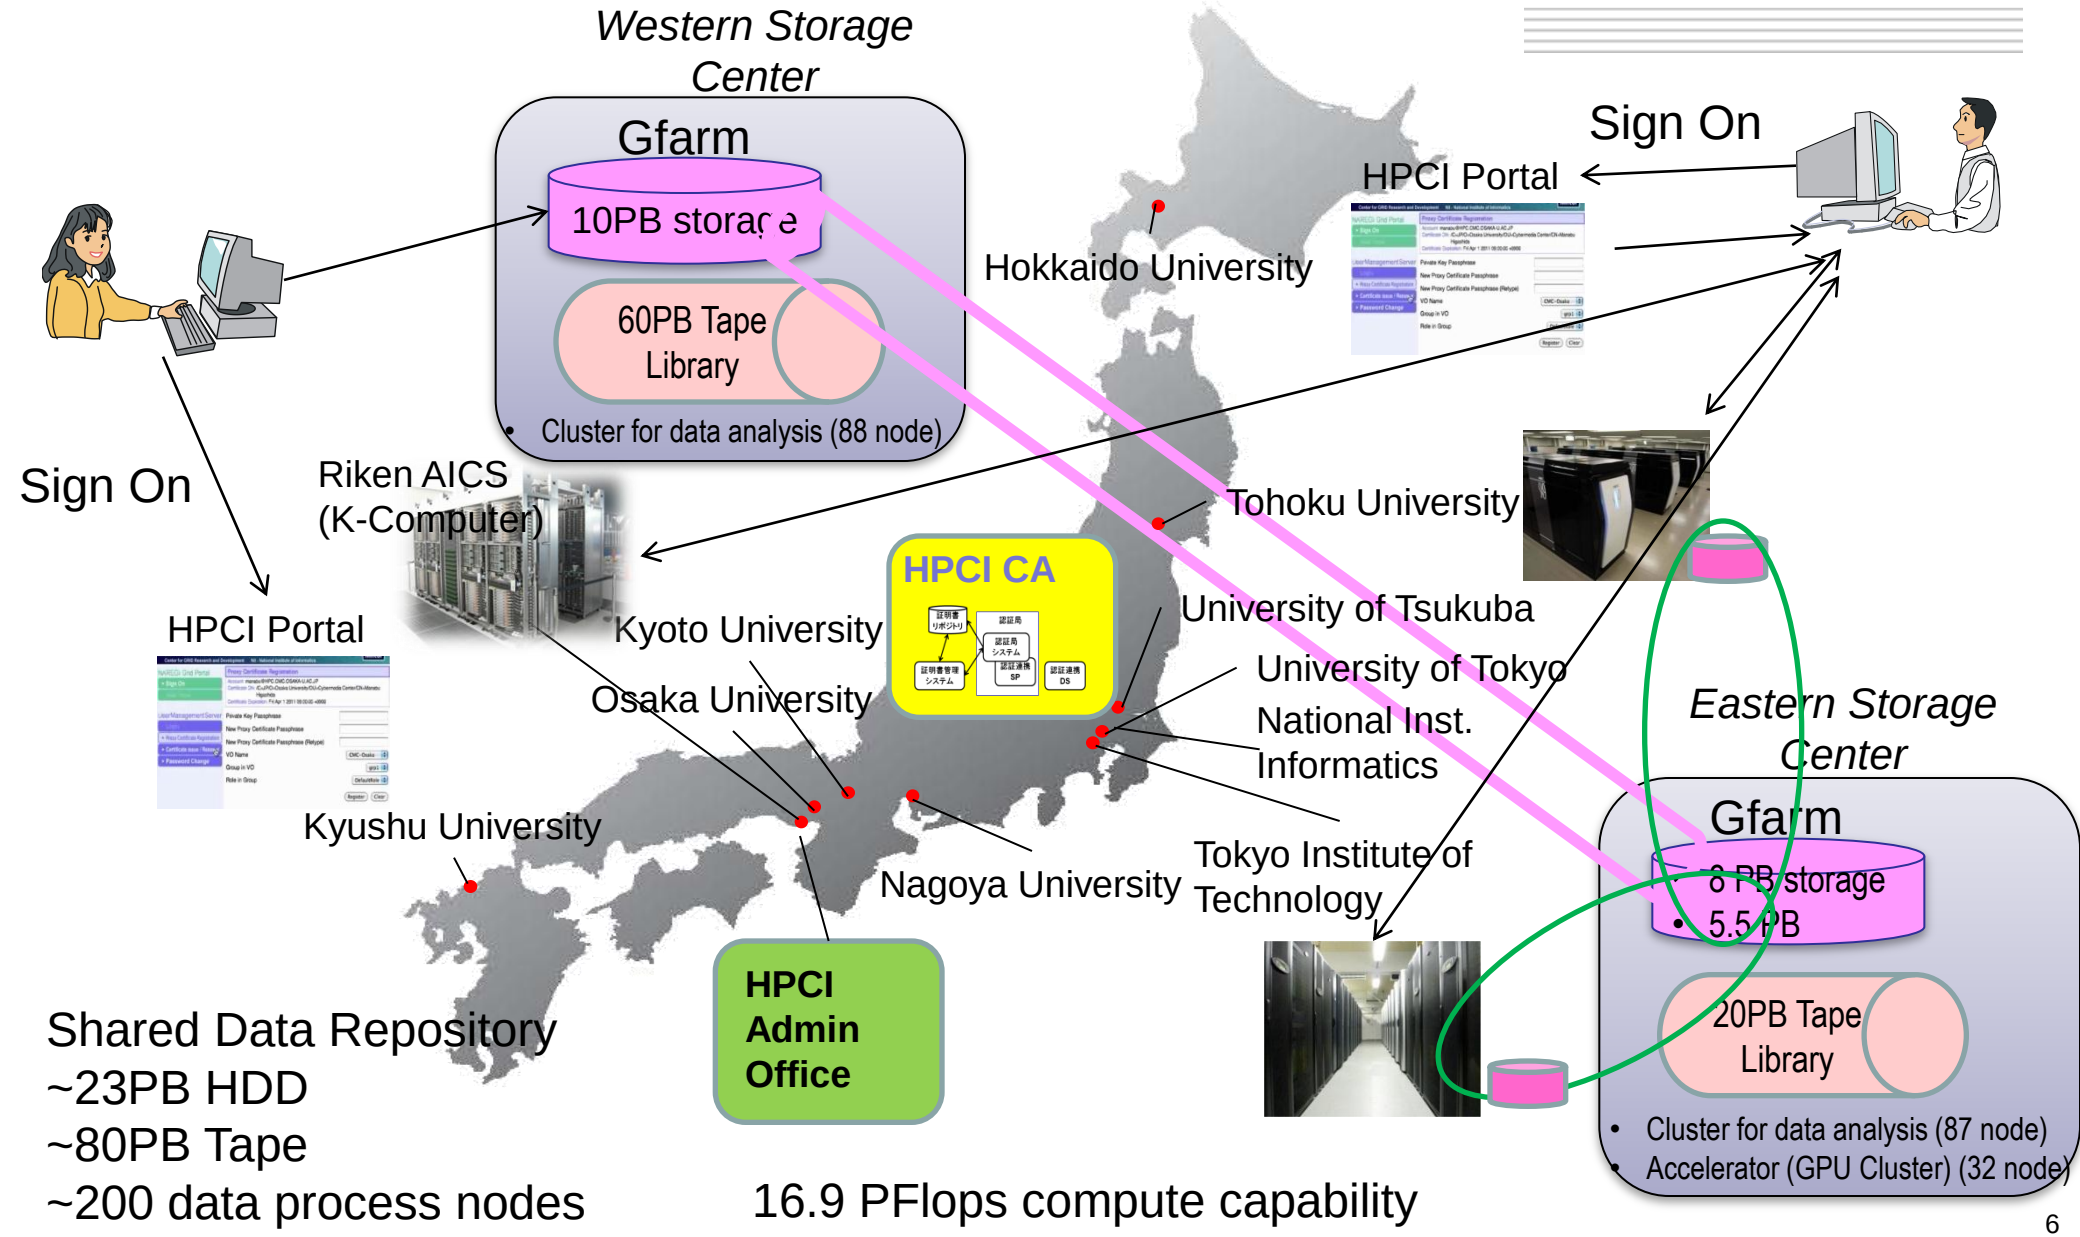

Evolving the e-Infrastructure Towards Exascale in 2019-2020 and Beyond

- U.S. – still frontrunner?
  - DoE: Exascale Co-Design Center, Software Center
  - DoE: various exascale research funds
  - DARPA: UHPC – 4 teams
- EU – Survey and Research
  - FP7: EESI- European Exascale Software Initiative
  - FP7: Exascale Computing
    - Mont-Blanc (BCS-ARM), DEEP (Julich-MIC), CRESTA (EPCC-Cray)
- China (Details not known), Russia, …
  - China announces Godson, 100PF by 2015, …
- Japan - survey and basic research only
  - HPCI-SDHPC and Feasibility Study, JST-PostPeta, etc.
- IESP (International Exascale Software Project)

NVIDIA Echelon Architecture

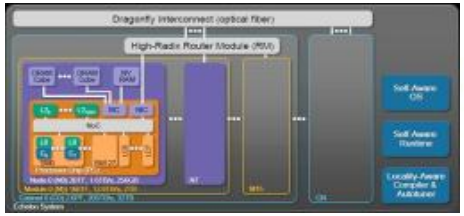

X-caliber

Challenges of Exascale (FLOPS, Byte, …) (10^{18})!

Various Physical Limitations Surface All-at-Once

- # CPU Cores: 10B
  Low Power
  c.f. Total # of Smartphones sold globally = 400Mil
- # Nodes 100K~*x*M
  c.f. The K Computer ~100K
  Google ~ 1 Mil

- Mem: *x*00PB~ExaB
  c.f. Total mem all PCs (300Mil) shipped globally in 2011 ~ ExaB
  BTW $2^{64\sim}=1.8\text{x}10^{19}=18$ExaB

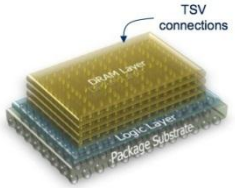

- Storage : xExaB c.f. Gmail Storage
  2 Exabyte (200Mil x 7GB+)

- All of this at 20MW (similar to K), reliability (MTTI=days), ease of programming (billion cores?), cost… in 2018?!

IBM SC vs. IT Leverage Strategy

- NCSA Blue Waters (would have been)


- LRZ SuperMUC (Warm Water Cooling)


- LANL Roadrunner


- LLNL/ANL/Julich BlueGene


Enormous Cost of e-Infrastructure for Science & Research

Who will pay the cost?

Indirect research benefits too slow

Exascale and other future infrastructural elements must exhibit immediate benefits as IT incubators

Mutual leverage with various aspects of IT industry as a whole

- Mainframes

- Mid-Low x86 Servers & PCs


- Gaming&Multimedia (PS3/Cell, XBOX360)


- Embedded (PowerPC)