10^{18}

BDEC

BIG DATA AND EXTREME-SCALE COMPUTING

Europe-USA-Asia
International series of Workshops on Extreme Scale Scientific Computing

Following the International Exascale Software Initiative
(IESP 2008-2012 ➔**Big Data and Extreme Computing** workshops (BDEC)
http://www.exascale.org/bdec/

Overarching goal:

1. Create an international collaborative process focused on the co-design of software infrastructure for extreme scale science, addressing the challenges of both extreme scale computing and big data, and supporting a broad spectrum of major research domains,
2. Describe funding structures and strategies of public bodies with Exascale R&D goals worldwide
3. Establishing and maintaining a global network of expertise and funding bodies in the area of Exascale computing

1 – BDEC Workshop, Charleston, SC, USA, April 29-May1, 2013

2 – BDEC Workshop, Fukuoka, Japan, February 26-28, 2014

3 – BDEC Workshop, Barcelona, Spain, January 28-30, 2015

BDEC is All About Convergence

- High-end data analytics (big data) and HPC are both essential elements of an integrated computing research-and-development agenda; neither should be sacrificed or minimized to advance the other.

- Programming models and tools are perhaps the biggest point of divergence between the scientific-computing and big-data ecosystems.

As scientific research increasingly depends on both high-speed computing and data analytics, the potential interoperability and scaling convergence of these two ecosystems is crucial to the future.

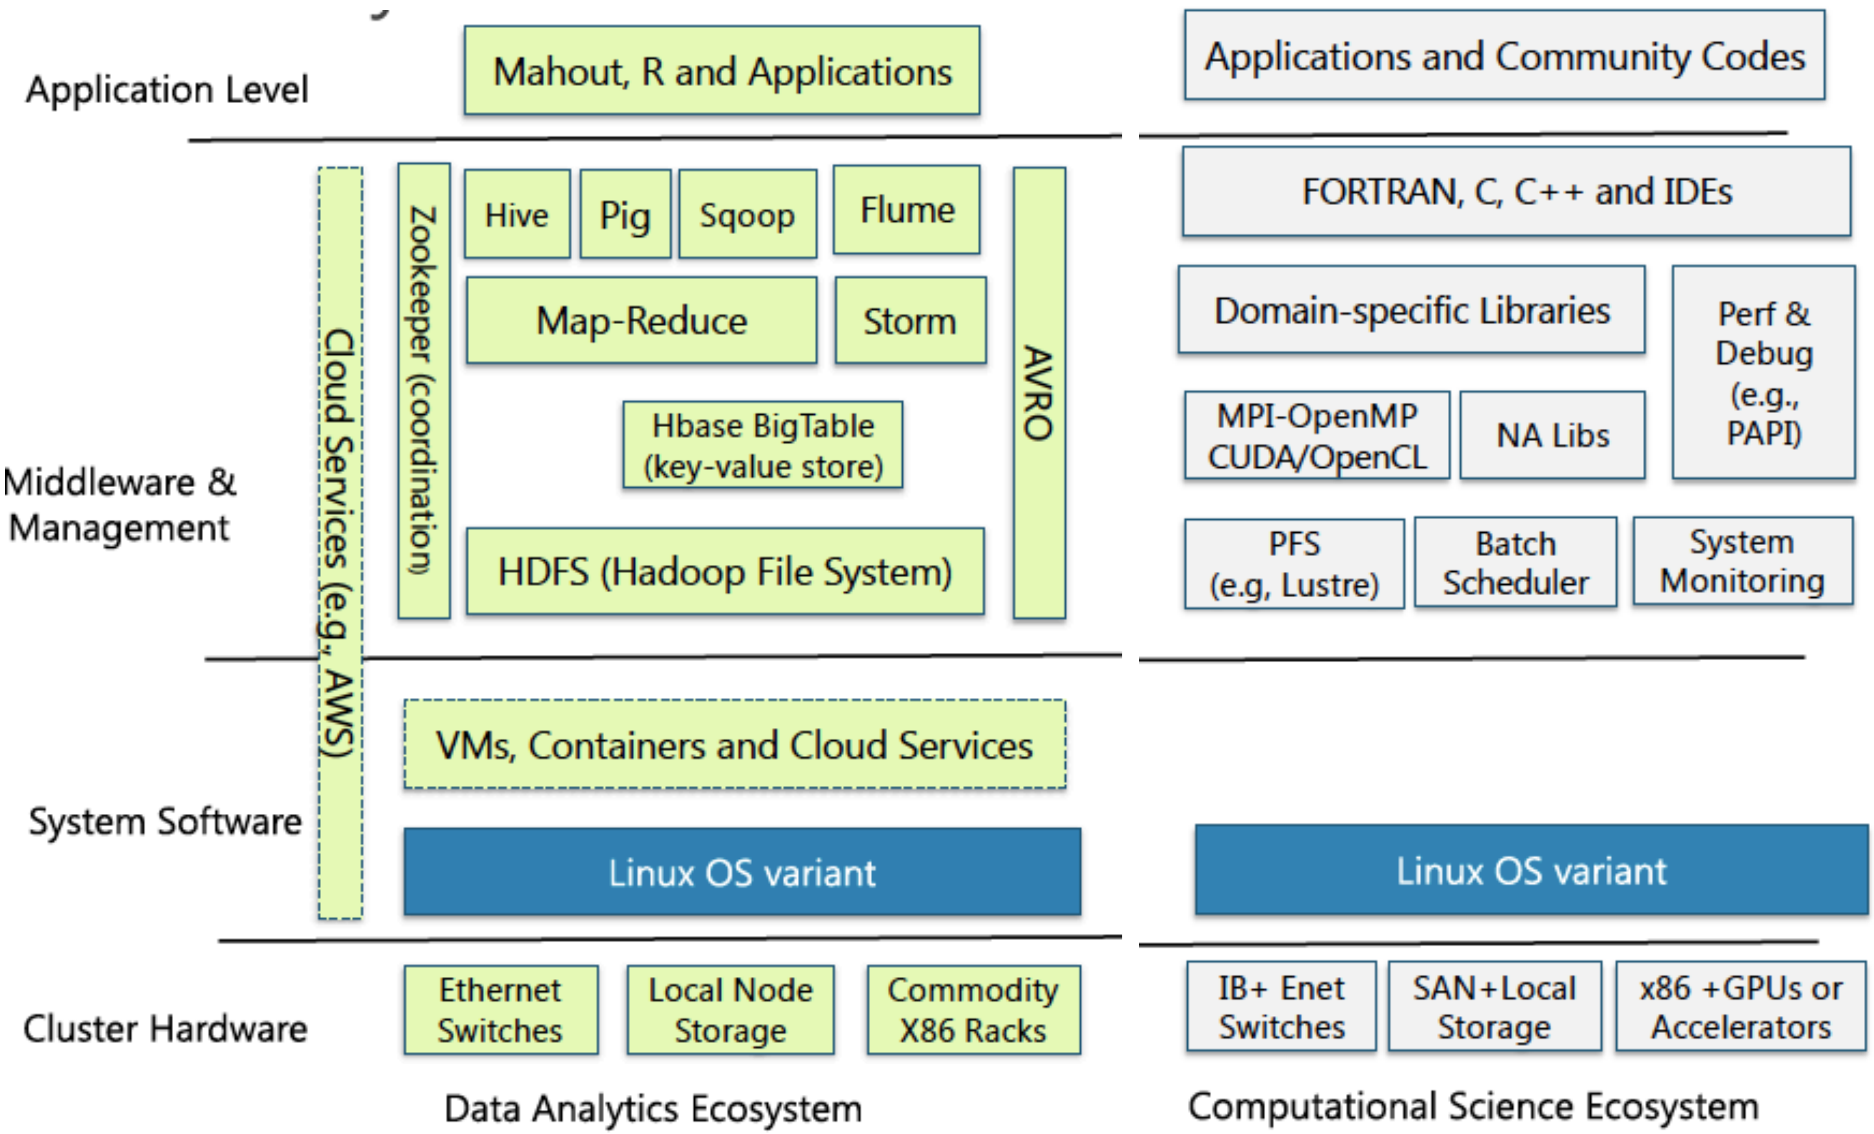

Comparing Architecture

Big Data	Extreme Computing
? Cost in memory and interconnect bandwidth	***Significant Cost*** in memory and interconnect bandwidth
Little Cost for resilient hardware in data storage	***Significant Cost*** in resilient hardware in shared file system
Little Cost for hardware to support system-wide resilience	***Significant Cost*** in resilience hardware to reduce whole-system MTTI
Significant Cost: ***increased aggregate IOP/s***	Significant Cost: ***cutting-edge CPU performance features***
Often trades performance for capacity	***Often trades capacity for performance***

Comparing Operations

Big Data	Extreme Computing
Continuous access to long-lived "services" created by science community	***Periodic access*** to compute resources via job submitted to scheduler and queue
Time-shared access to elastic resources	***Space-shared*** compute resources for exclusive access during jobs
New hardware capacity ***purchased incrementally***	New tightly integrated system ***purchased every 4 years***
Users charged for all resources (storage, cpu, networking)	***Users charged for CPU hours, storage and networking is free***

Comparing Software

Big Data	Extreme Computing
Software responds to elastic resource demands	After allocation, ***resources static until termination***
Data access often ***fine-grained***	Data access is ***large bulk*** (aggregated) requests
Services are resilient to fault	***Applications restart after fault***
Often ***customized*** programming models	Widely ***standardized*** programming models
Libraries help ***move computation to storage***	Libraries help ***move data to CPUs***
Users routinely deploy their own services	***Users almost never deploy customized services***

Comparing Data

Scientific Big Data	Extreme Computing
Inputs ***arrive continuously***, streaming workflows	Inputs ***arrive infrequently***, buffering carefully managed
Data is ***unrepeatable*** snapshot in time	Data often ***reproducible*** (repeat simulation)
Data generated by sensors (***error: from measurement***)	Data generated from simulation (***error: from simulation***)
Data rate ***limited by sensors***	Data rate ***limited by platform***
Data often ***shared and curated*** by community	Data ***often private***
Often ***unstructured***	***Semi-structured***

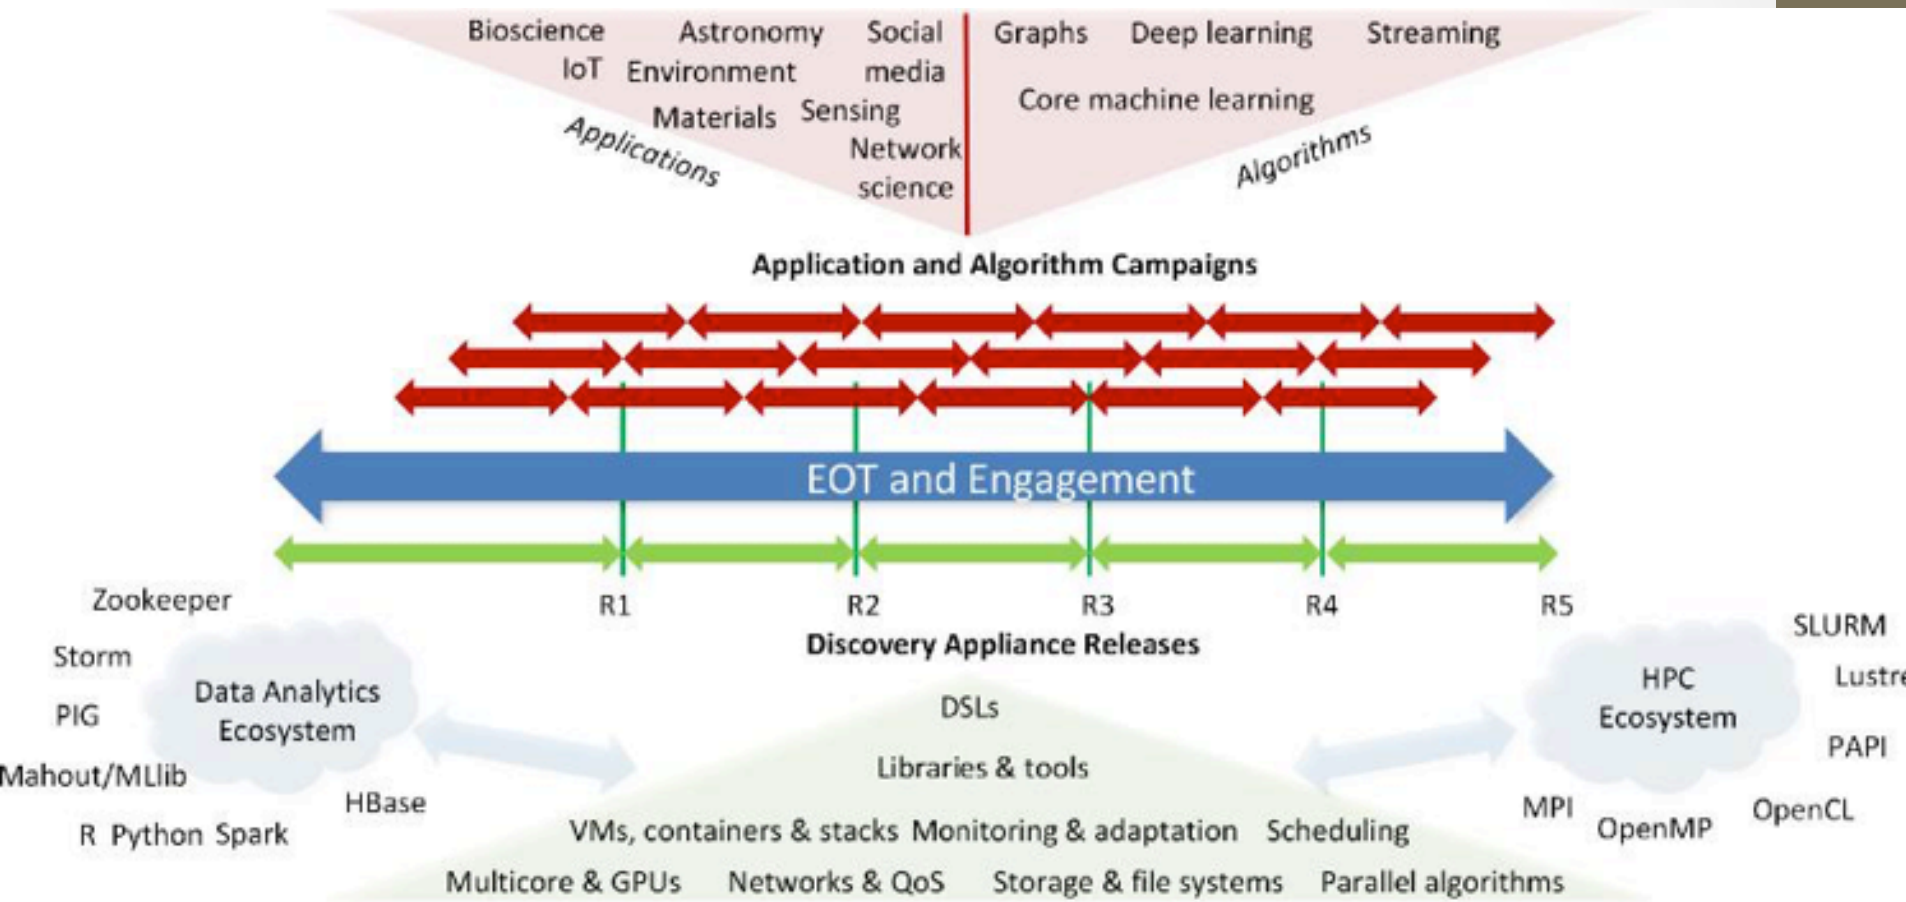

Bioscience
Astronomy
Social media
IoT
Environment
Materials
Sensing
Network science
Applications
Graphs
Deep learning
Streaming
Core machine learning
Algorithms
Application and Algorithm Campaigns
EOT and Engagement
R1
R2
R3
R4
R5
Discovery Appliance Releases
Zookeeper
Storm
PIG
Mahout/MLlib
R Python Spark
HBase
Data Analytics Ecosystem
DSLs
Libraries & tools
VMs, containers & stacks
Monitoring & adaptation
Scheduling
Multicore & GPUs
Networks & QoS
Storage & file systems
Parallel algorithms
HPC Ecosystem
SLURM
Lustre
PAPI
MPI
OpenMP
OpenCL

Goals for This Meeting: How do we Converge?

Breakout groups

Applications and Science
Chairs: David Keyes, Rosa Badia, Jean-Claude Andre

Architecture and Operation/Comprehensive Production Services
Chairs: Bill Kramer, Ewa Deelman, Francois Bodin

Algorithm and Applied Mathematics
Chairs: Hiroshi Nakashima, Philippe Ricoux, Alison Kennedy

Software Stack
Chairs: Franck Cappello, Kate Keahey, Satoshi Matsuoka

Plan for the Day

- Reports on current strategies from US, Europe, Asia
- Opportunities for community engagement

Introduction
09:00 am - 09:10 am
Pete Beckman, Argonne National Laboratory

Jack Dongarra, University of Tennessee & ORNL

BDEC Software
09:10 am - 09:30 am
Kate Keahey, Argonne National Laboratory

Satoshi Matsuoka, Tokyo Institute of Technology

BDEC Algorithms
09:30 am - 09:50 am
Mike Heroux, Sandia National Laboratories

BDEC Architecture
09:50 am - 10:10 am
Bill Kramer, NCSA

BDEC Applications
10:10 am - 10:30 am
Rosa Badia, BSC

David Keyes, KAUST

EXDCI (EESI3)
10:30 am - 11:00 am
François Bodin, IRISA

SPEXXA2: A Success Story of Multi-Agency Collaboration
11:30 am - 12:00 pm
Mark Asch, ANR

Wolfgang E. Nagel, TU Dresden

Marcus Wilms, DFG

DOE Perspective
12:00 pm - 12:15 pm
William Harrod, DoE

NSF Perspective via WebEx
12:15 pm - 12:30 pm
Irene Qualters, NSF

US Japanese
12:30 pm - 01:00 pm
Pete Beckman, Argonne National Laboratory

Yutaka Ishikawa, RIKEN AICS & University of Tokyo

Jeffrey Vetter, ORNL & Georgia Institute of Technology

Welcome & Goals/Objectives for BDEC Cooperation Opportunities to Engage the Broader International Community
02:00 pm - 02:10 pm
James Ang, Sandia National Laboratories

Opening Remarks & A Personal Perspective from an International BDEC Road Warrior
02:10 pm - 02:30 pm
Thomas Sterling, Indiana University

SKA Project Update & Collaboration Opportunities
02:30 pm - 03:00 pm
Happy Sithole, CSIR

SKA, DOME & Astron Project
03:00 pm - 03:30 pm
Ronald P. Luijten, IBM Research Zurich

BDEC Collaboration Opportunities in Russia
03:30 pm - 04:00 pm
Vladimir Voevodin, MSU

tba
04:30 pm - 05:00 pm

BDEC Collaboration Opportunities at A*Star/Singapore
05:00 pm - 05:30 pm
Marek T. Michalewicz, A*STAR

BDEC Collaboration Opportunities at KAUST/Saudi Arabia
05:30 pm - 06:00 pm
David Keyes, KAUST

Wrap-Up
06:00 pm - 06:15 pm
James Ang, Sandia National Laboratories
Thomas Sterling, Indiana University