# Breakout session
# Architecture and Operation

29 January 2015

# Questions

- What are the main differences and commonalities between the HPC and BDA requirements/technologies/working-assumptions in this area?
- Are there common needs/problems/interfaces could serve as the basis (or as stepping stones) along a path to (some reasonable level of) infrastructure and application convergence?
- Are there interdomain testbeds that combine BDA and HPC workflows in ways that could help uncover pathways toward convergence?
- What is/are the technology or new research that may be a game changer?
- What action would be your number one priority to be taken rapidly to ensure success of the converge of Extreme computing and Big Data infrastructures?
- What action would be your number one priority to be taken rapidly to ensure the emergence of efficient Extreme computing and Big Data applications?
- How would you measure the success of the BDEC initiative?

# What are the main differences and commonalities between the HPC and BDA requirements/technologies/working-assumptions in this area?

- ## Data apps
  - Database, realtime, virtual machine, cloud, not waiting for resources, interactive data analytics, languages (e.g. Java), aggregate I/O, importance of time to solution rather than raw performance, shared memory machines, in-memory, integer performance (?), dynamic data, unstructured data

- # Commonalities
  - Visualization, memory system convergence(?), data reduction capabilities

- ## HPC apps
  - Multi-many cores, arithmetic, Hadoop useful?, static data, structured data

# Are there common needs/problems/interfaces could serve as the basis (or as stepping stones) along a path to (some reasonable level of) infrastructure and application convergence?

- Energy cost
- Data movements dominating energy
- New kinds of memories, persistent storage
- HPC output is a big data issue
- Workflow is more complex in BD apps (heterogeneous machines)
- Heterogeneous "data view" (same data different hardware)
- Common APIs energy aware
- Pathways for data movements
  - New file systems
  - IO system useful for HPC and data analytics
  - Common flexible resiliency and consistency model
  - Heterogeneity (multi-physics), flexible hardware (including virtualization / scheduling), ability to specify compute and memory apps needs
  - data reduction capabilities

# Are there interdomain testbeds that combine BDA and HPC workflows in ways that could help uncover pathways toward convergence?

- Climate modeling type workflows
  - Requires different types of resources
- HPC creates BD, BD reduces data
- Intelligent cities generates BD problems
- Testbeds
  - System with mix types of nodes, scheduling capabilities to access the different types in a coordinated manner
  - Interactive use of resources from desktops, batch jobs
  - Potentially different network topologies
  - Data reduction capabilities
  - Enabling studying (dynamic) tradeoff between re-compute vs store
  - Monitoring tools (performance, energy, IO, …)
  - Flexible repair modes (nodes and communication)
  - Hierarchical storage, flexible sub-systems composition, strongly and loosely coupled sub-systems
  - Scheduling / QoS / resources (interactive) management
  - Software stack issues aware

# What is/are the technology or new research that may be a game changer?

- Research beneficial for HPC and BD
  - High capacity, high bandwidth cheap memories
  - Processors making use of 3D memory (string matching)
  - High speed to storage (HDD is a bottleneck)
  - Interconnects speed order magnitude (inter-processors, inter-clusters)
  - Novel data representation (floating point)
  - High level abstractions for computation and data
  - Percipient storage
- Research to promote convergence
  - High performance file systems
  - Speed in/off chips
  - Avoiding data movements using active storage
  - Software defined provisioning and management of resources
  - Co-existing VM in traditional HPC systems / software and application stack control in HPC
  - Ease of application validation in different environments
  - Methods for co-location of computation and data
  - Benchmark application models, traces
  - Automated, easier way to express optimal/efficient use of deep memory hierarchies, on-the fly data processing
  - Efficient graph libraries

# What action would be your number one priority to be taken **rapidly** to ensure **success** of the converge of Extreme computing and Big Data **infrastructures**?

- 1 billion euros!
- New APIs
- HP file-object/storage systems
- Dynamic integration of memory and storage resources API
- Benchmark/mini-apps (HPC & BD workload)
- More collaboration between HPC and BD researchers
- Automated data movements
- Realizing convergence is not needed
- Knowledge dissemination between the communities
- HP virtualized I/O
- Virtualized high speed interconnect for all parts of the system
- Energy efficient resource management
- Elaborating a value proposition (ROI)  for convergence
- Work with selective set of applications to make them work
- - 1 billion euros
- Find a good example where such a convergence is useful/necessary

- **HP virtualized memory & storage system with open API for transparent data movement and on-the-fly processing**

What action would be your number one priority to be taken **rapidly** to ensure the emergence of **efficient** Extreme computing and Big Data **applications**?

- Setting up a file system testbed
- Programming model / API with data abstraction fitting HPC – BD needs
  - Both type of applications would use the API/language
  - Starting standardization
- Repository of reference components and workflow systems useful for HPC and BD
- Establish relationship with major "data creators"
  - Large instruments (LSST, SKA, LIGO, …)
- Defining new benchmarks (and methods)
- Revisiting policies for resource allocation on HPC systems

# How would you measure the success of the BDEC initiative?

- New calls for proposals targeting convergent systems, software and applications

- New synergies between countries and organizations

- A strawman proposal that outlines a set of activities that would lead to "convergence"

- Determine a clear vision on how commercial and publicly (national and regional ) funded resources may brought together to integrate HPC and BD techniques to solve socially relevant research problems

- Address trans-national policies to encourage collaborations and flexible efficient resources allocations

# Executive Summary

- **1 General**
  - 1.a We need to identify different levels of convergence and understand their ROIs. We need to conduct a cost benefit analysis to determine where and how convergences would benefit the user communities and how to best prioritize the activities in a way that reflect the needs of the users communities and the priorities of the funding organizations.
  - 1.b Address trans-national policies to encourage collaborations and flexible efficient resources allocations, Exploit new synergies between countries and organizations.
  - 1.c Other

# Executive Summary

- ## 2 Architecture (3)

  - 2.a More robust and dynamic methods to move the data where they are needed

  - 2.b. Ensure I/O and storage technology research and productization targeting needs of converged systems receive sufficient focus and funding

  - 2.c Other

# Executive Summary

- 3 Production / services (3)
  - 3.a We need to setup a repository of reference components and workflow systems useful for HPC and BD
  - 3.b We have a clear need for convergence of resources allocation and management mechanisms and services
  - 3.c Encourage resource providers to adopt a user-centric model that includes support for convergent BD/HPC applications
  - 3.d Testbeds and Benchmarks
  - 3.e Other

# Who's doing what

- Current state: Bill
- Motivation: Bill
- General: MCS
- Architecture: AG
- Production and service: Ewa

- Bronis will take an integration pass when other

# To keep in mind

- Users support capturing HPC and BD operational issues

- Orchestration may be another architectural attribute

- Need to understand the initiatives in the BD community that may be closed to BDEC

- There are a lot of production service/operational aspects of HPC and BD that are already synergistic

# Thanks to the cats

- Ewa Deelman
- Francois Bodin
- Piyush Mehrotra
- Marie-Christine Sawley
- Giovanni Erbacci
- Yutaka Ishikawa
- Toshio Endo
- Bill Kramer

- Osman Unsal
- Jamie Kinney
- Bronis de Supinski
- Masaaki Kondo
- Marek Michalewicz
- Malcolm Muggeridge
- Jean-Francois Lavignon