

Exascale and Big Data Convergence

SCIENCE, TECHNOLOGY AND CULTURE



Executive Order -- Creating a National Strategic Computing Initiative

EXECUTIVE ORDER

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

Section 1. Policy. In order to maximize the benefits of HPC for



National Strategic Computing Initiative (NSCI)

White House Executive Order



1. Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
2. *Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.*
3. Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post- Moore's Law era").
4. Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
5. Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.



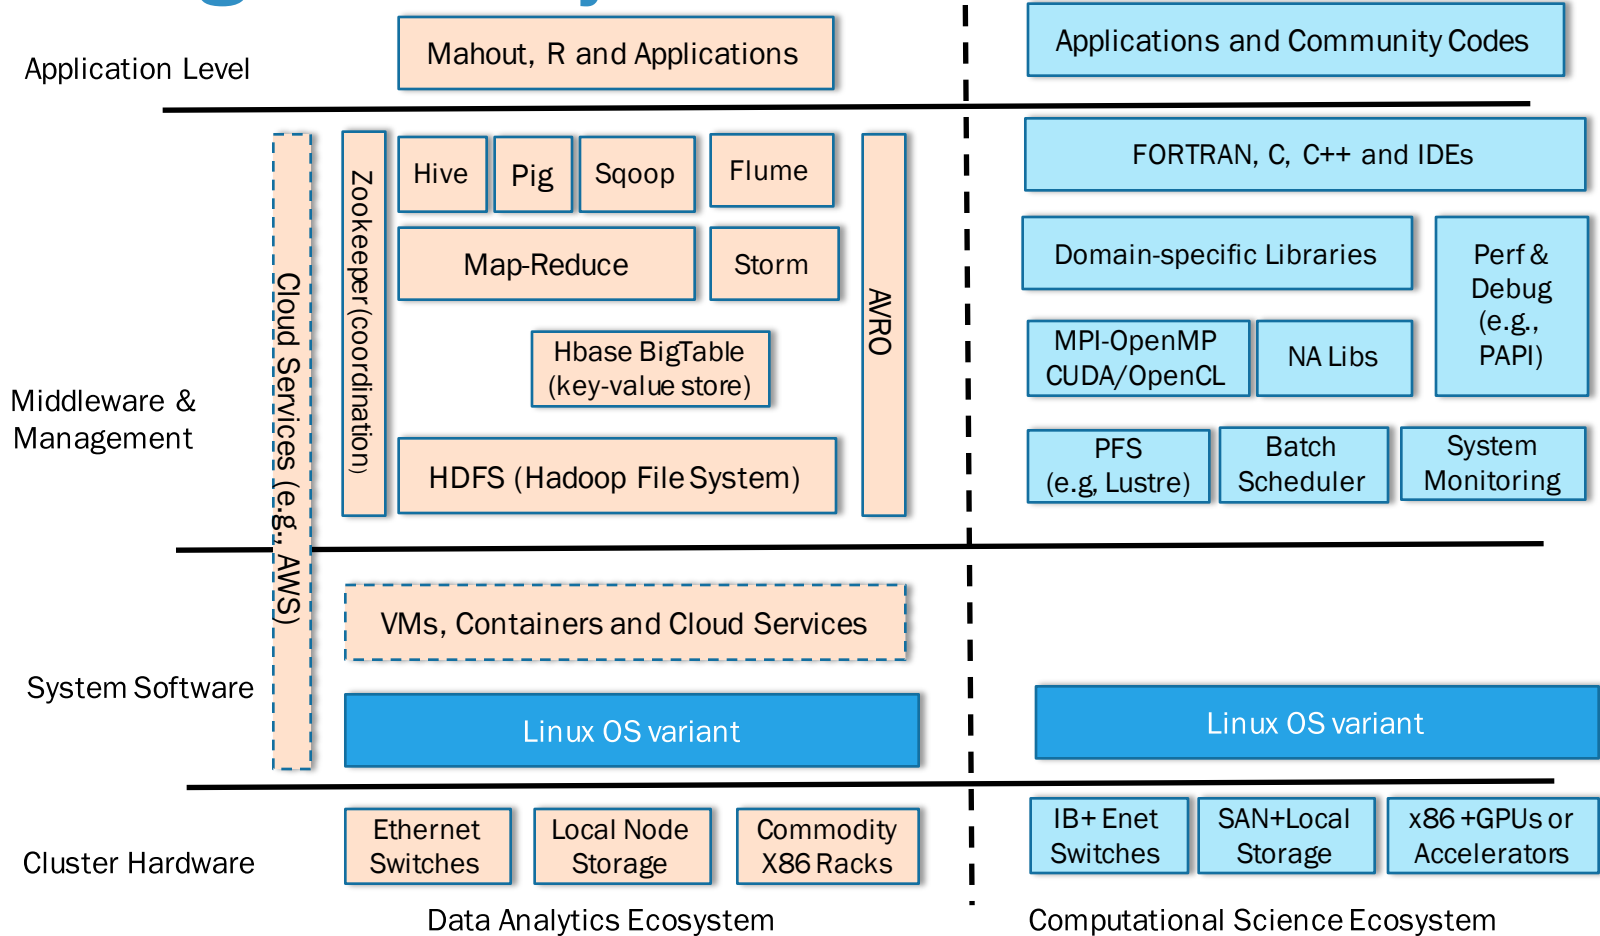
NSCI workshop (October 2015)

Opening panels

- Convergence of Data Analytics and Computationally-Intensive Computing

“As NSCI drives forward these two goals of exaflop computing ability and exabyte storage capacity, it will also find ways to combine large-scale numerical computing with big data analytics. This will enable new forms of computation, including simulations of weather that are coupled with actual observations from weather satellites. It will also enable new analytic methods that require more extensive numerical processing, such as emerging techniques that use artificial intelligence to automatically learn new capabilities from large numbers of examples.”

Divergent ecosystems



Have you ever ...



Requested compute and storage for years of *continuous data analysis*?

Sought to run a heterogeneous workflow across your system?

Found your code performance limited by sensor network & I/O bandwidth?

Struggled with transient local storage and data staging?

Asked your technical computing system for deep learning recommendations?

Wished you could load containers rather than just applications?

Envied the tools and size of the cloud and machine learning communities?

Wondered why R came after S and C doesn't matter?

Previous BDEC recommendations

Submit complete software stack with computation / analysis pipeline

Computing jobs that dynamically change (grow/shrink) resource needs

Soft real-time analysis of sensor or instrument data

Manage semi-persistent services for data analysis with virtual orgs

Extreme-scale databases as basic infrastructure, and with allocation

Triggered workflows linked to data arrival

Support for HPC curation/provenance of millions of binary data objects