Berkeley Research Computing

Larry Conrad, CIO

Berkeley
UNIVERSITY OF CALIFORNIA

Challenge and Opportunity

- Faculty PI's needed campus infrastructure for research computing, esp. HPC
  - Recruitment and retention challenges
  - Losing grants (no more private clusters)
- Broad need for computing support
  - Not just physical sciences, genomics, etc.
  - Not just HPC: Cloud, VMs with analytics stacks
- Innovation: Reproducible research, etc.
  - Not enough $$ to compete with privates
  - IT/CompSci innovation to advance *all* research

Berkeley
UNIVERSITY OF CALIFORNIA

Peer Benchmarking Report: Shared Research Computing

November 22, 2013

Subject Experts

Steve Masover, Patrick Schmitz, Chris Hoffman - IST-RIT; Harrison Dekker - Library Data Lab

Description

Includes provision for research *and* teaching of: "traditional" HPC (highly parallelized computing), Data Science methodologies & computational resources, high-powered workstations (including VMs) to support computation at a level between a typical desktop/laptop and an HPC cluster or VM array. Secure compute, storage, data transfer, and data archiving are also in scope.

Benchmarking Criteria

- **Coordinated program** that includes a suite of coordinated services to support computational research and teaching, including a roadmap for service evolution.
- **Support for diverse computational research techniques**, e.g., 'traditional' HPC, virtual machine arrays, and high-powered workstations (which may be virtualized); as well as data transfer and lifecycle management.
- **Training**: Availability and breadth of training.
- **Documentation**: Availability and breadth of documentation.
- **Consulting services**: Including assessment and advice on aligning research problems/needs to available computational resources; grant writing, hardware and software purchasing, and software design, tuning, and refactoring consultation.

Summary of Findings

Tier	Description	Institutions
1	• Strong across all benchmarking criteria	UC San Diego, Princeton, Northwestern
2	• Strong in most benchmarking criteria, stronger in some areas than others.	Harvard, Michigan, MIT, NYU, UCLA, Virginia
3	• Mixed assessment	Columbia, Stanford, Cornell, UW
4	• Weak assessment in most or all areas.	**Berkeley**

Draft Recommendations

Tier	Action
4 → 2	Build a comprehensive program for research computing that provides a range of services from traditional HPC to cloud VM resources to virtual workstations. Develop a community of consultants who have joint appointments in schools, colleges, centers with RIT. One time investment of approx. $1.2 million and recurring investment of up to $1.8 million.
2 → 1	Use Berkeley's strengths in innovation and partnerships with such groups as EECS/Amp Lab, D-Lab, BIDS, and science centers to grow new services in cloud-based HPC and virtual research workstations.

Description

Criteria

Findings

Recommendations

Berkeley
UNIVERSITY OF CALIFORNIA

UCB/LBNL Partnership

- Working together since 2008
- CalHPC Program
  - Founded 2008 to provide cluster support for researchers
  - IST provided colo, networking; LBL handled admin, support
  - 7 clusters managed
  - Recharge model incurred “double taxation”
- Berkeley Research Computing
  - HPC condo/institutional cluster managed by LBL team at UCB
  - Cloud compute service in development (private/hybrid/public)
  - Virtual Workstations with analytics profiles
  - Research Consulting and engagement

Berkeley
UNIVERSITY OF CALIFORNIA

Berkeley Research Computing Timeline 2013/14

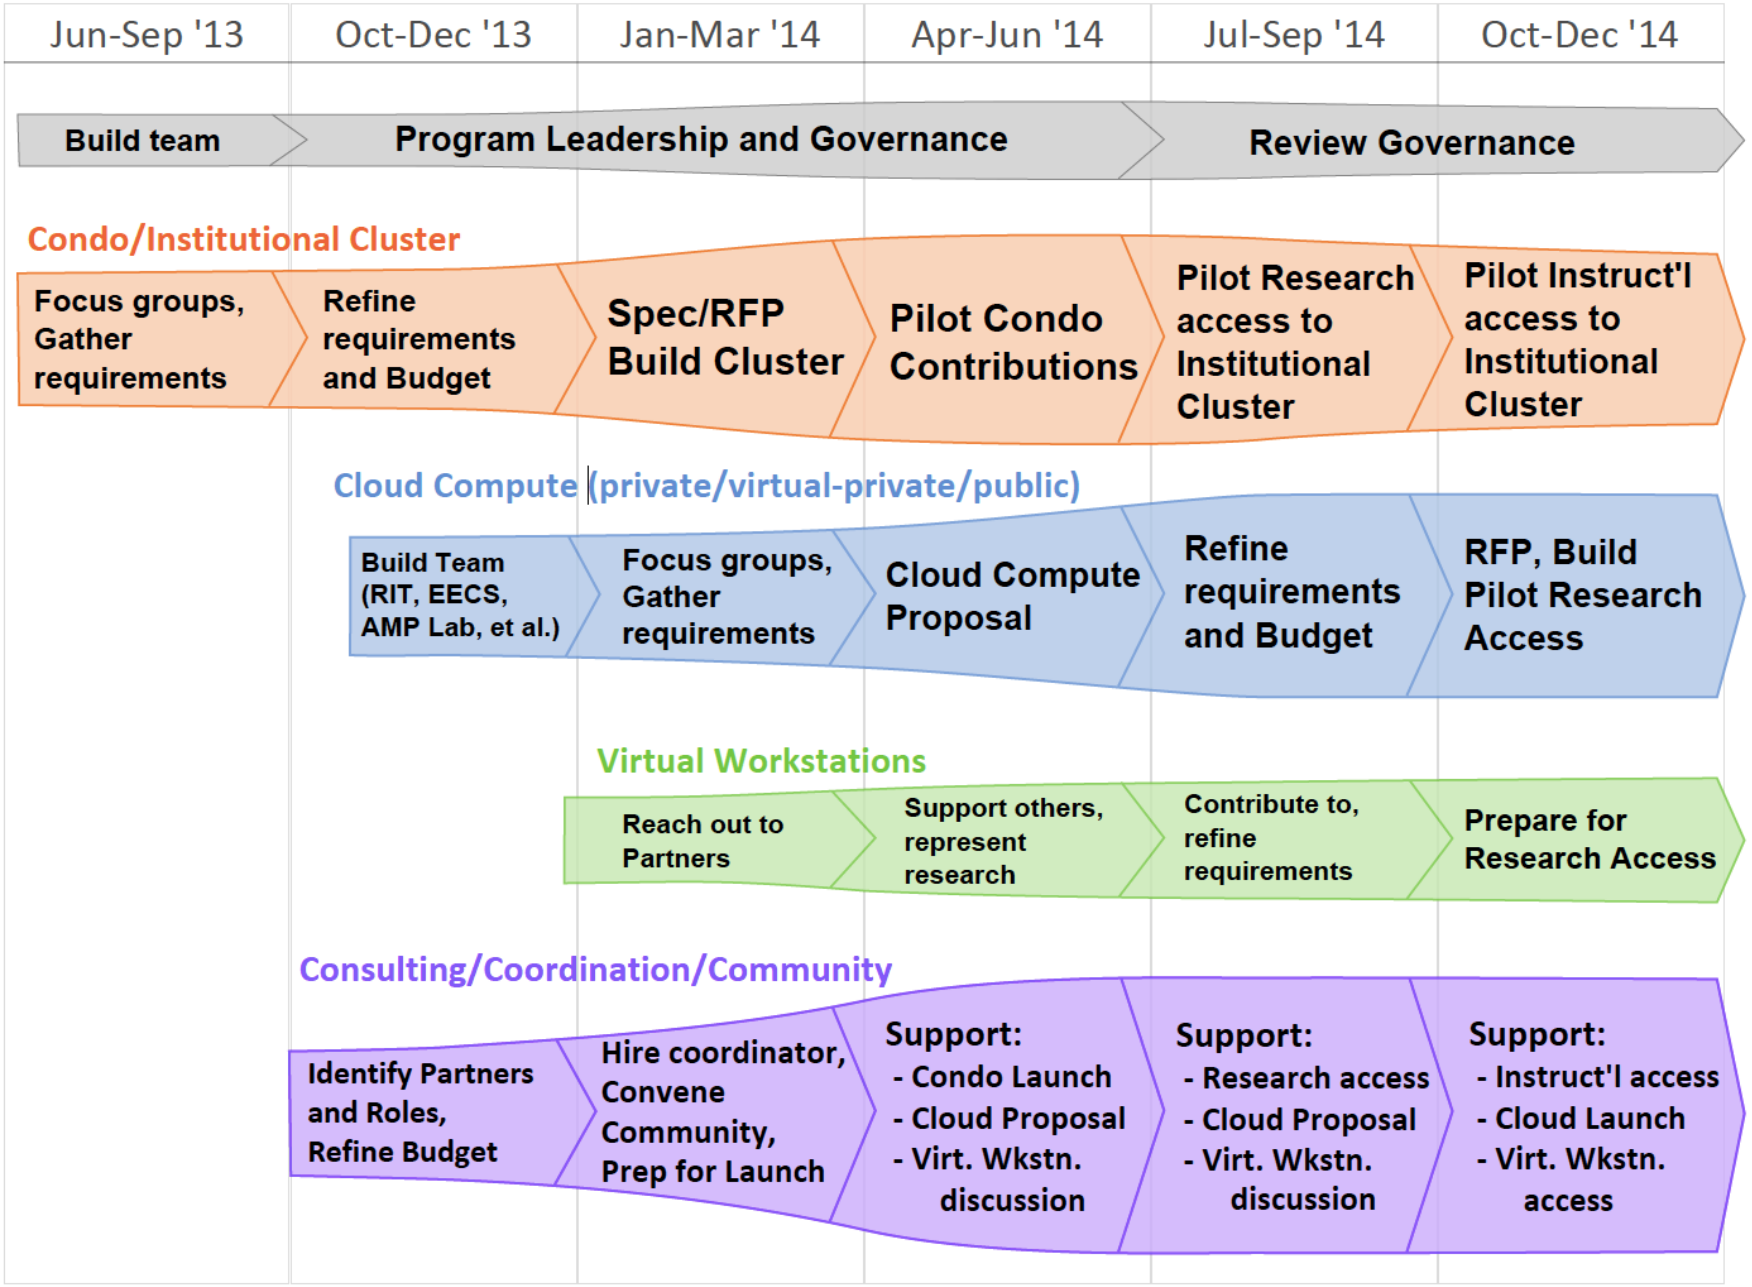

Note: width of band indicates level of activity over time.