A nighttime photograph of a city skyline reflected in water. A prominent feature is a bridge with blue lighting that spans across the water. The background shows various skyscrapers and buildings, some with lights on. The overall scene is dark, with the city lights providing the main illumination.

Open Cirrus Overview

Presentation at the 6th Open Cirrus Summit
Atlanta, GA, October 12, 2011

Dejan Milojevic
Senior Research Manager and Researcher
Managing Director, Open Cirrus
Hewlett-Packard Laboratories
dejan.milojevic@hp.com



Outline

- Old slide set on Open Cirrus with some analysis
- More recent research examples
- Lessons learned
- If we were to do it all over again today...?
- Summary
- What Next

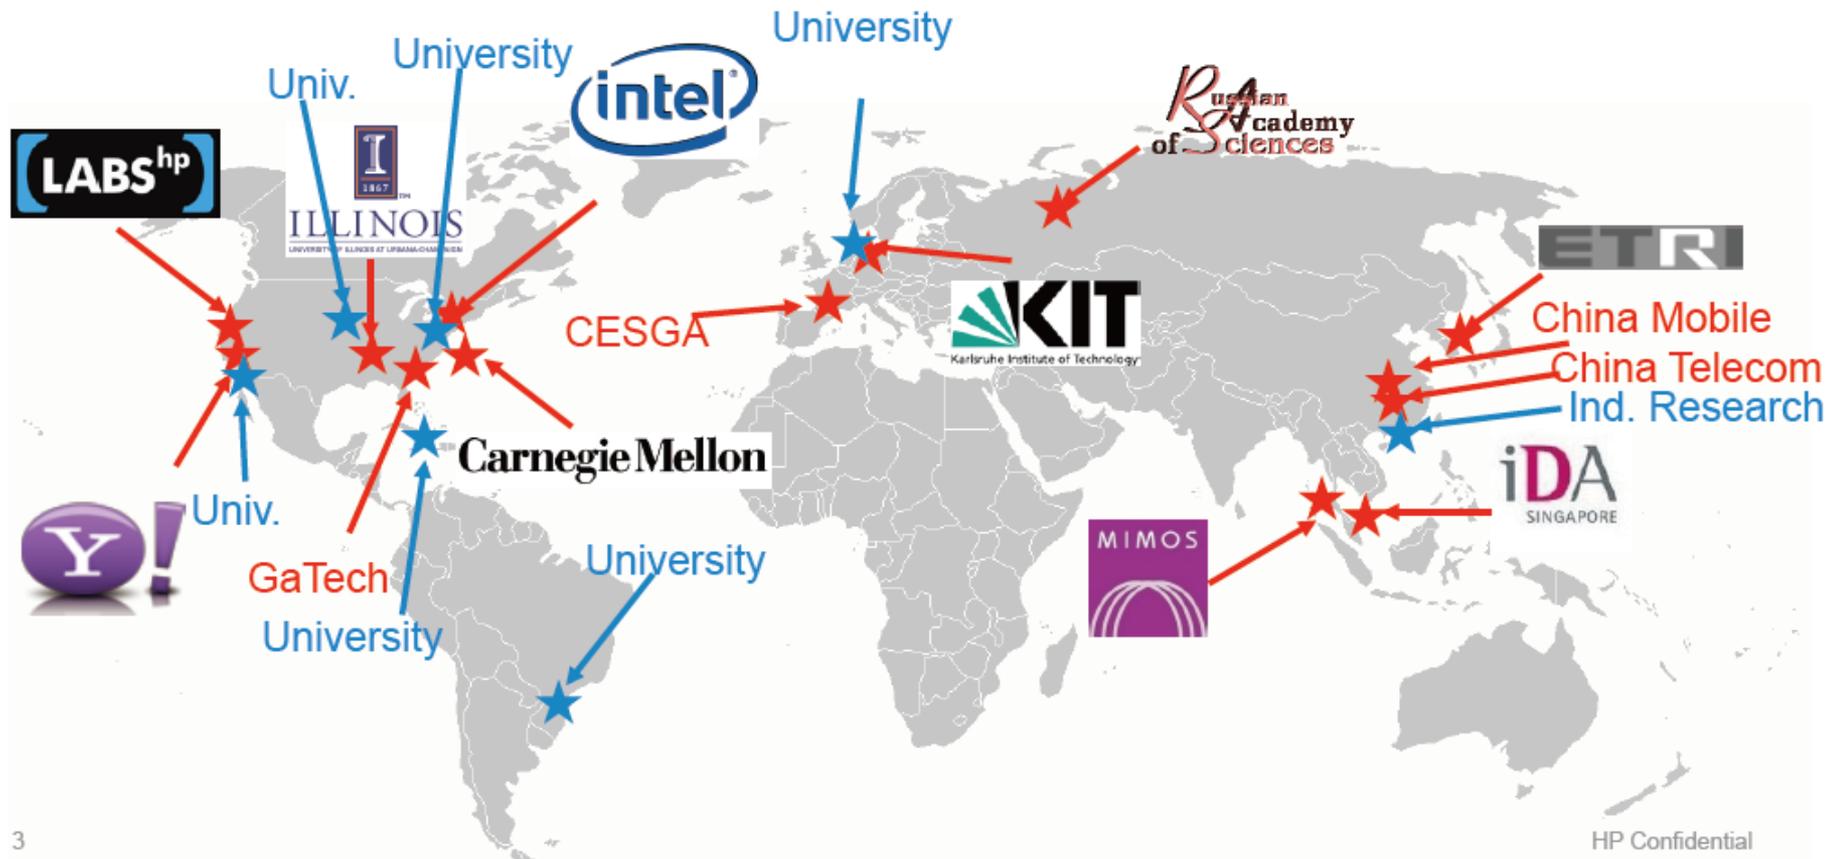
Open Cirrus™ Cloud Computing Testbed

Shared: research, applications, infrastructure (11K cores), data sets

Global services: sign on, monitoring, store. **Open source stack** (prs, tashi, hadoop)

Sponsored by **HP, Intel, and Yahoo!** (with additional support from **NSF**)

- 15 sites currently, target of around 20 in the next two years.



Open Cirrus

- Objectives

- Create an ecosystem for Cloud services modeling +
- Foster systems research around cloud computing ++
- Expose research community to enterprise level requirements +
- Provide realistic traces of cloud workloads - -
- Vendor-neutral open-source stacks and APIs for the cloud

- How are we unique

- Support for systems research and applications research ++
- Federation of heterogeneous datacenters +
- Interesting data sets -

Process

- Central Management Office, oversees Open Cirrus
- Governance model
 - Research team
 - Technical Team
 - New site additions
 - Support (legal (export, privacy), IT, etc.)
- Each site
 - Runs its own research and technical teams,
 - Contributes individual technologies
 - Operates some of the global services
- E.g. HP Site supports: Portal and PRS

Open Cirrus (TM)

Home

Search this site:

marthalyons

- My account
- ▷ CoE Collaboration Space
- Admin Forums
- Workflow summary
- ▷ Create content
- Recent posts
- ▷ Feed aggregator
- ▷ Administer
- Log out

Request Access

Requesting Access to OpenCirrus

The resources provided through the OpenCirrus Cloud Computing Testbed are a finite resource and are intended to be used for research purposes only.

Consequently, OpenCirrus computing resources are allocated to research projects that must be approved by one or more of the OpenCirrus Centers of Excellence. Project proposals are submitted by a *Principal Investigator* who is typically a university faculty member, senior staff member, or industrial researcher/technologist. Once a project is approved the Principal Investigator is able to identify additional team members who should be granted access as part of the project. This organization is similar to the arrangement for [PlanetLab](#) and nearly identical to the one used for [Emulab](#).

Project Proposal Process

The process for proposing a project is relatively straightforward.

1. First, the Principal Investigator (PI) should select one of the OpenCirrus Centers of Excellence to serve as the *Home Site* for a project.
2. The PI should email a brief description of the project to the Project Coordinator at the Home Site. This description should include at least (1) the research goals of the projects, (2) a high-level description of the OpenCirrus resources that would be involved, and (3) the expected project start/end dates. See a sample [here](#). The research coordinators for each site are listed below:

- **HP Labs Site** - Martha Lyons, martha.lyons@hp.com
- **Intel Pittsburgh Research Site** - Michael Kozach, email@intel.com
- **Yahoo! Research** - Thomas Kwan, email@yahoo.com
- **UIUC** -
- **KIT** -
- **Singapore IDA** -

New forum topics

- [Hi Everyone](#)
- [Participating in Open Cirrus!](#)
- [Why is Open Cirrus unique and interesting to the systems community?](#)
- [Why are we building Open Cirrus?](#)
- [Welcome!!!! Some facts about Open Cirrus!](#)

[more](#)

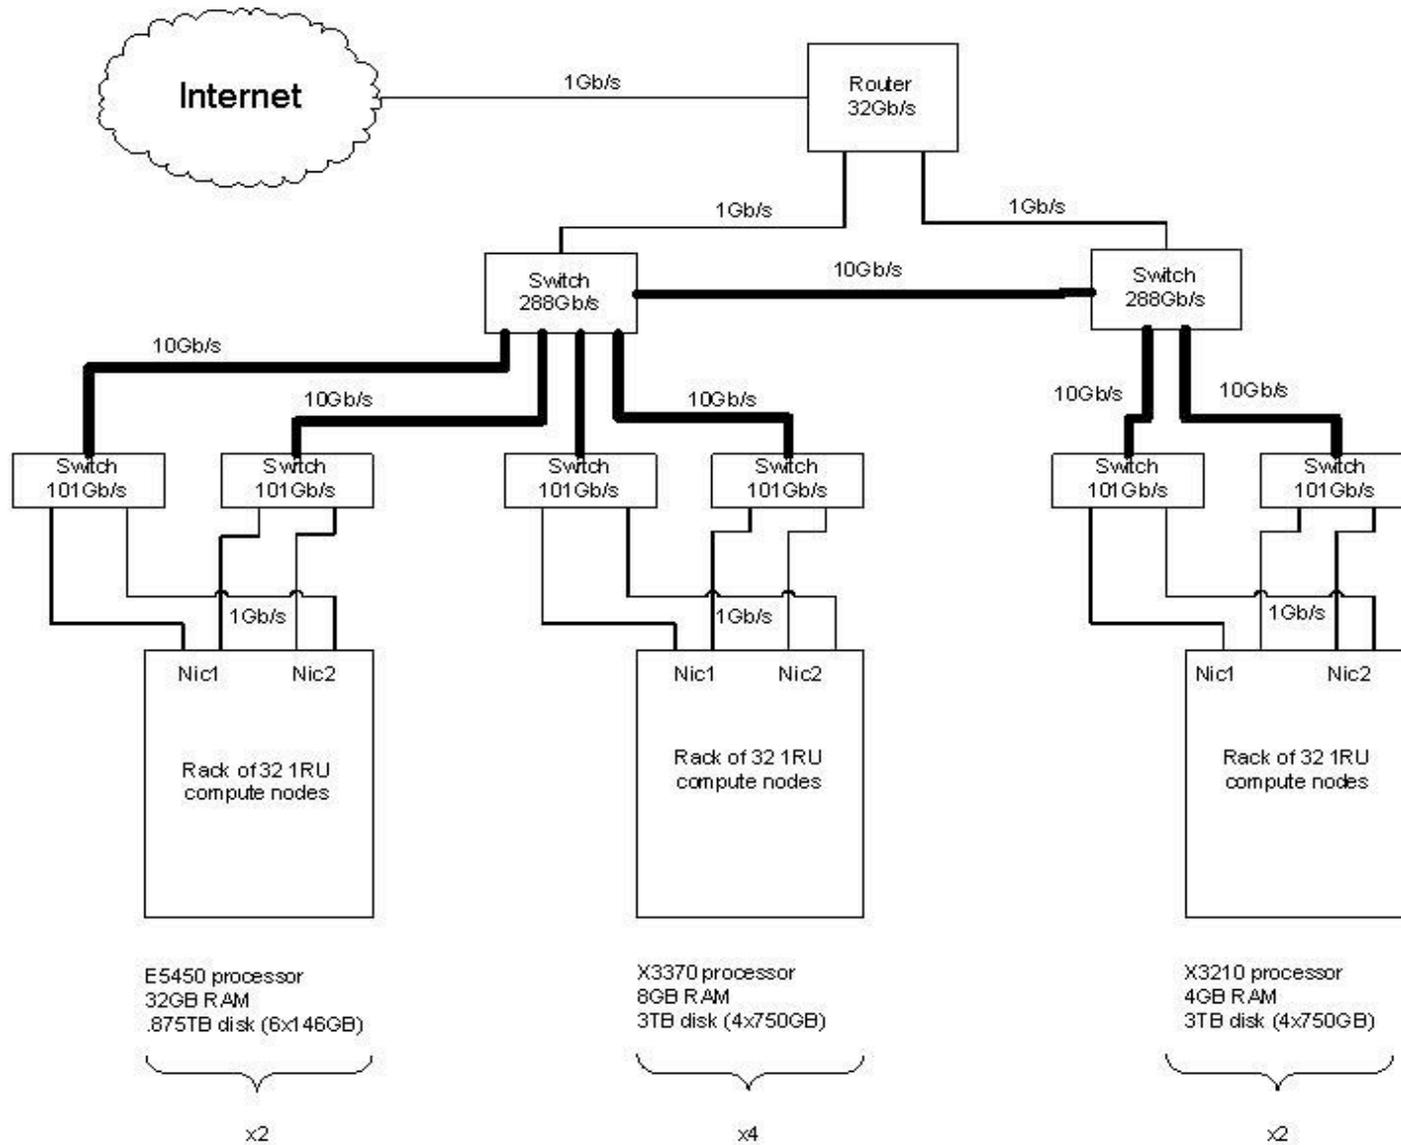
Who's online

There are currently *1 user* and *0 guests* online.

Online users

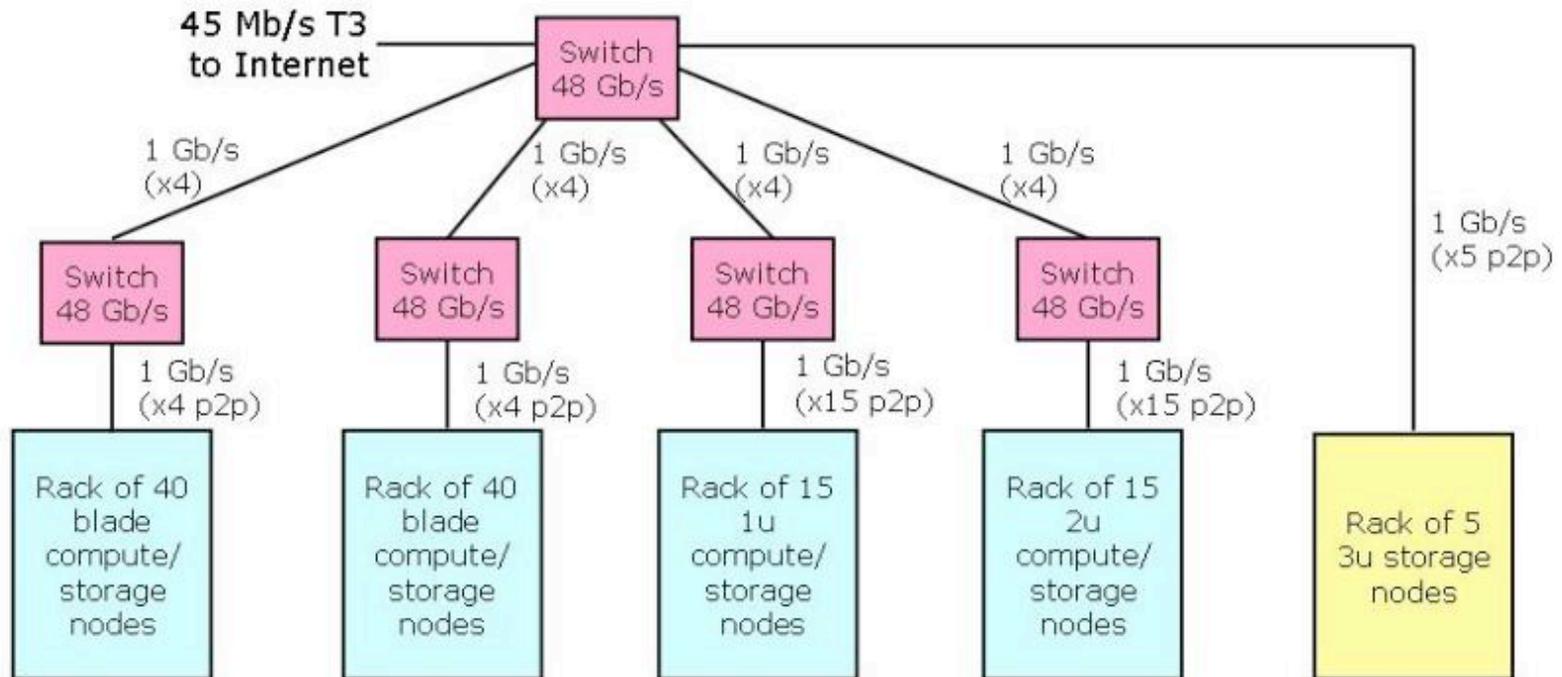
- [marthalyons](#)

HP Labs Cirrus Cluster topology



				Totals
Nodes/cores:	64/256	128/512	64/256	256/1024
RAM (GB):	2048	1024	256	3328
Storage (TB):	56	384	192	632
Spindles:	384	512	256	1152

Intel Research BigData Cluster



10 nodes: 8 Core2 cores (4x2), 8GB RAM, 0.3TB disk (2x150GB)

20 nodes: 1 Xeon core, 6GB RAM, 366GB disk (36+300GB)

10 nodes: 4 Xeon cores (2x2), 4GB RAM, 150 GB disk (2x75GB)

Node: 8 Core2 cores (4x2), 8GB RAM, 0.3TB disk (2x150GB)

Node: 8 Core2 cores (4x2), 8GB RAM, 2TB disk (2x1TB)

x2

Node: 8 Core2 cores (2x4), 8GB RAM, 6TB disk (6x1TB)

x3

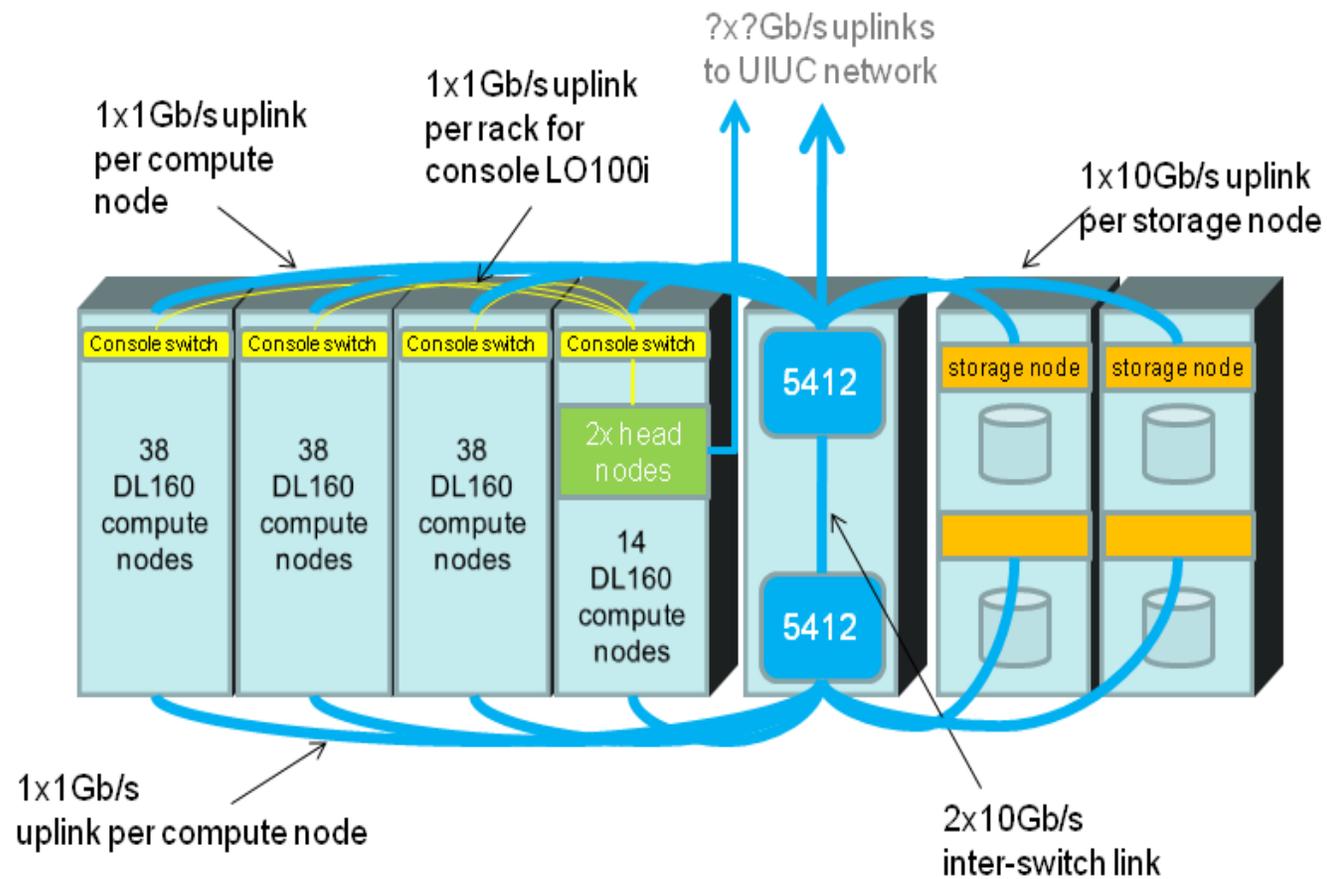
Node: 12TB disk (12x1TB)

[totals]

Nodes/cores:	40/140	40/320	30/240	45/360		[155/1060]
RAM (GB):	240	320	240	360		[1160]
Storage (TB)	11	12	60	270	60	[413]
Spindles:	80	80	60	270	60	[550]

UIUC cluster network topology

- Console switches: connect to 1 port/head node
- Links to external UIUC network connect either to both head nodes, or to both core switches
- Each storage node has 1x10Gb/s to each core switch
- Each compute node has 1x1Gb/s link to each core switch
- Core switches have 2x10Gb/s inter-switch links between them

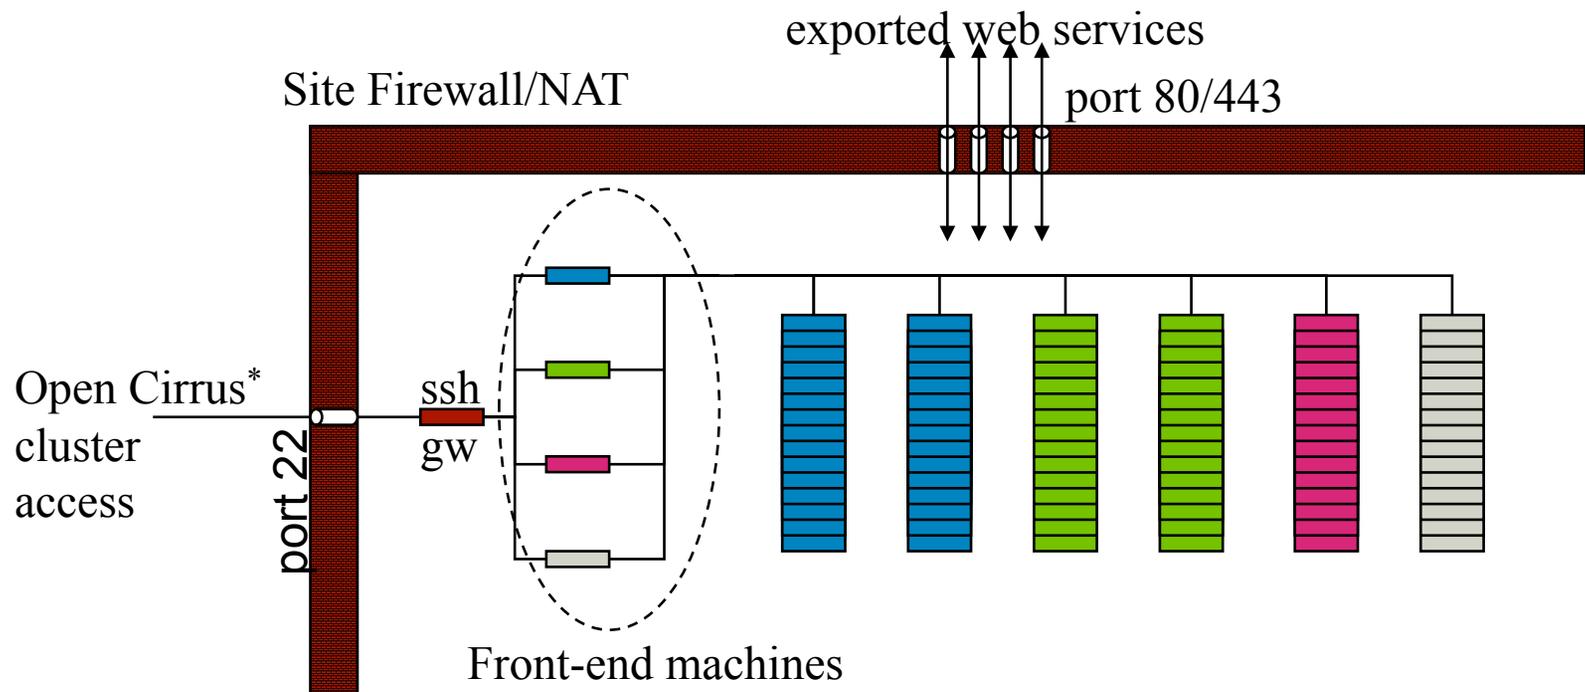


Open Cirrus Sites

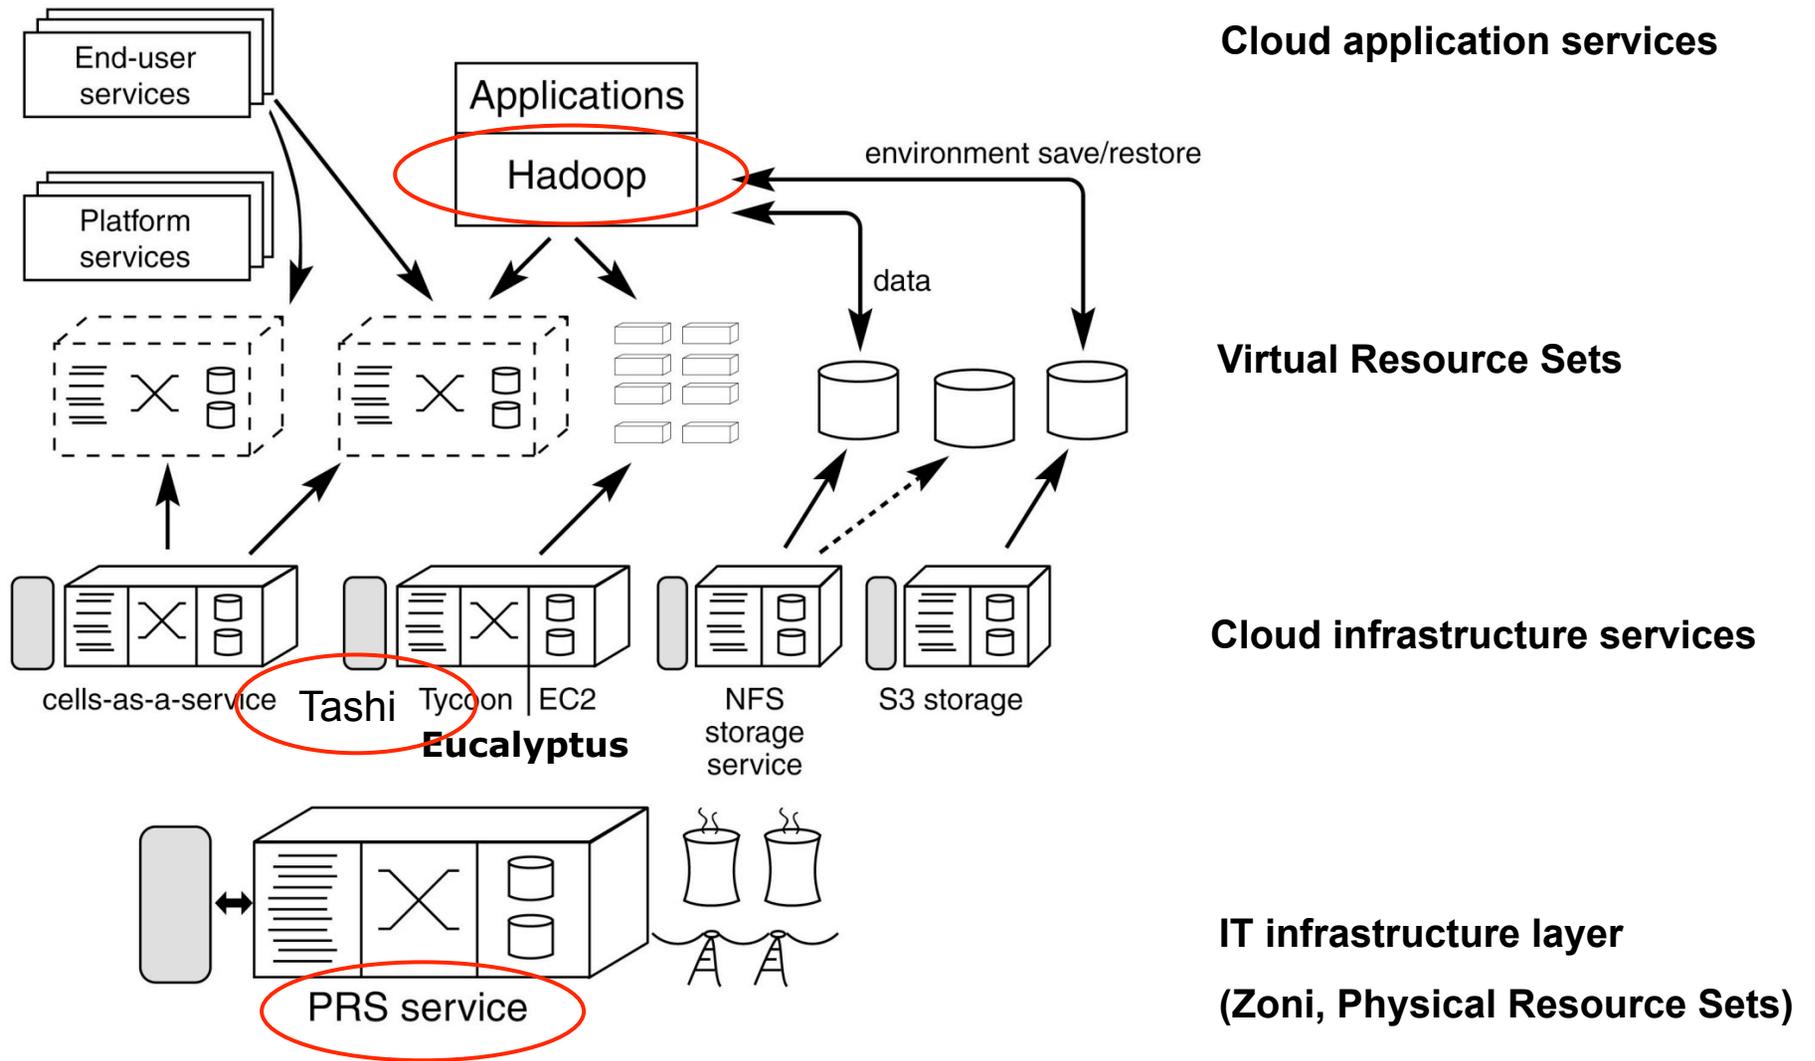
Site	Characteristics							
	#Cores	#Servers	Public partition	Memory Size	Storage Size	Spindles	Network	Focus
HP	1,024	256	178	3.3TB	632TB	1152	10G internal 1Gb/s x-rack	Hadoop, Cells, PRS, scheduling
IDA	2,400	300	100	4.8TB	43TB+ 16TB SAN	600	1Gb/s	Apps based on Hadoop, Pig
Intel	1060	155	145	1.16TB	353TB local 60TB attach	550	1Gb/s	Tashi, PRS, MPI, Hadoop
KIT	2048	256	128	10TB	1PB	192	1Gb/s	Apps with high throughput
UIUC	1024	128	64	2TB	~500TB	288	1Gb/s	Datasets, cloud infrastructure
Yahoo	3200	480	400	2.4TB	1.2PB	1600	1Gb/s	Hadoop on demand

Access Model

- At a minimum, sites must expose a ssh gateway
- Sites may also provide additional external connections
 - Some provision for web services is highly recommended
- Sites may also be divided into resource pools by service
 - Some services may require a front-end machine (e.g. hadoop)



Open Cirrus Software Stack



How do users get access to Open Cirrus sites?

- Project PIs apply to each site separately
- Contact email addresses on the Open Cirrus portal
 - <http://opencirrus.org>
- Each Open Cirrus site decides which users and projects get access to its site
- *A global sign on* for all sites
 - Users are able to login to each OpenCirrus site for which they are authorized using the same login and password.

What kinds of research projects are Open Cirrus sites looking for?

- Open Cirrus™ is seeking research in the following areas (different centers will weight these differently)
 - Datacenter federation
 - Datacenter management
 - Web services
 - Data-intensive applications and systems
 - Hadoop map-reduce applications
- The following kinds of projects are of less interest
 - Traditional HPC application development
 - Production applications that just need lots of cycles
 - Closed source system development

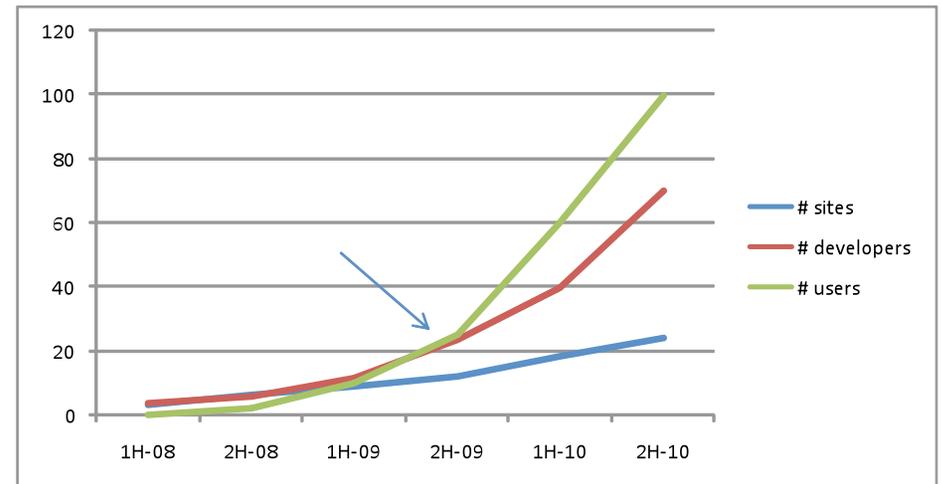
Metrics of Success

- Community

- + – Technology used
- + – # Sites, Projects, (Vibrant) Users
- + – Research Productivity (Shared Cost of Research), # papers published
- - – Cross-collaboration (Portal traffic)
- ++ – # New open source components
- – Global presence

- Technical

- + – Utilization of Open Cirrus, TCO
- + – Ease of use (e.g. provision 50% of OC nodes in < 30sec)
- – Federation transparency/adoption
- – Reliability



Open Cirrus v. Other Testbeds

	Testbeds							
	Open Cirrus	IBM/Google	TeraGrid	PlanetLab	EmuLab	Open Cloud Consortium	Amazon EC2	LANL/NSF cluster
Type of research	Systems & applications	Data-intensive applications research	Scientific applications	Systems and services	Systems	interoperab. across clouds using open APIs	Commer. use	Systems
Approach	Federation of heterog. data centers	A cluster supported by Google and IBM	Multi-site hetero clusters super comp.	A few 100 nodes hosted by research instit.	A single-site cluster with flexible control	Multi-site heterogeneous clusters	Raw access to virtual machines	Re-use of LANL's retiring clusters
Participants	HP, Intel, IDA, KIT, UIUC, Yahoo!	IBM, Google, Stanford, U.Washington, MIT	Many univ. & orgs	Many univ & organizations	University of Utah	4 centers –	Amazon	CMU, LANL, NSF
Distribution	15 sites	1 site	11 partners in US	> 700 nodes world-wide	>300 nodes univ@Utah	480 cores, distributed in four locations		1000s of older, still useful nodes at 1 site

Open Cirrus Research Summary

HP

- Mercado
- Policy Aware Data Mgmt
- Wikipedia Mining & tagging
- SPARQL Query over Hadoop (UTD)
- N-tier App Benchmark (GaTech)

- Economic Cloud Stack
- Parallel Data Series

- OpenNet
- Exascale Data Center

Intel

- Everyday Sensing and Perception
- SLIPstream/Sprout
- Parallel Machine Learning
- NeuroSys
- Computational Health
- FastBeat (w/France Telecom)

- Tashi (with CMU, Yahoo)

- PRS (with HP)

Cloud application frameworks and services

Cloud infrastructure services

IT infrastructure layer

OpenNet on OpenCirrus

- OpenNet
 - Programmable, open layer-2 network
 - Features for
 - Robust, adaptive routing over redundant layer-2 networks
 - VM machine migration *without* dropping connections
 - In-situ network monitoring
 - Quality-of-Service guarantees
 - Installed on OpenCirrus cluster at HP Fall 2009
- OpenNet on OpenCirrus
 - Full bisection bandwidth
 - Virtual machine migration
 - Platform for high energy efficiency in the Data Center
 - Based on SPAIN (HP Labs), PortLand (UC San Diego)
 - Joint project between HP Labs, UC San Diego (funded by HP Open Innovation Program)

OpenCirrus on GENI

- GENI: Global Environment for Network Innovations
 - Major National Science Foundation program to provide a national-scale experimental facility for computer science researchers
 - Currently entering Spiral Two prototyping phase
- OpenCirrus on GENI
 - Give access to GENI researchers to the OpenCirrus platform (PlanetLab Control Framework for OpenCirrus)
 - Give OpenCirrus users access to GENI resources
- Key technological challenges
 - Mutual authentication between PlanetLab Control and OpenCirrus
 - Exchange of authorization and access functions
 - Resource allocation
- Status
 - Joint proposal to GENI Project Office by HP Labs (Kevin Lai, Rick McGeer) and UC San Diego (Alex Snoeren, Amin Vahdat)
 - Accepted by GENI Project Office (GPO) for Spiral Two Funding
 - Part of GPO proposal to NSF for Spiral Two (decision early Sept)

SPARQL Query over Hadoop for Very Large RDF Datasets



- Provide a semantic web framework using Hadoop which scales for large RDF data sets.
 - Use the Lehigh University Benchmark (LUBM) data (provides 14 queries) to measure SPARQL queries implemented over Map/Reduce framework provided by Hadoop.
 - Goal: to find the best possible way to query the data (SPARQL) by Map/Reduce programming.

N-tier Application Benchmark & Evaluation over Open Cirrus



- Generate, deploy, and run N-tier application benchmarks (including non-stationary workloads)
 - Collect data on standard and custom N-tier application benchmarks such as RUBiS (e-commerce) and RUBBoS (bulletin board) over a wide range of settings and configurations (both hardware and software)
 - Collect, analyze, and evaluate performance data using statistical software tools.
 - Apply the experimental evaluation results to cloud management applications such as configuration planning and adaptive reconfiguration

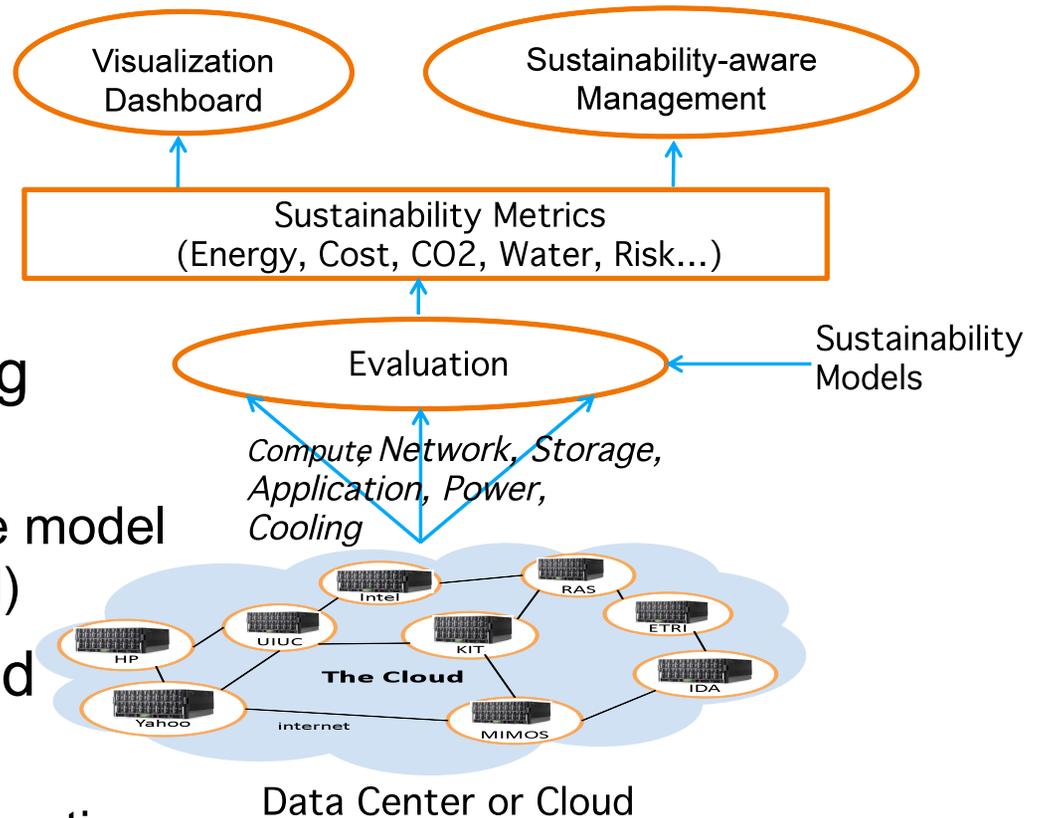
Cloud Sustainability Dashboard (CSD)

Open Cirrus Site	Economical (\$)					Ecological				Social		
	IT	cooling	ntwk	support	econo. overall	CO ₂ (tonnes-eq)	water (mill. Gal)	Resource Use (GJ-eq)	ecolog. overall	State of devt.	Risk of instability	social overall
Site 1	\$0.72	\$0.35	\$0.16	\$0.43		6.0	2.6	83		High	Low	
Site 2	\$1.27	\$0.59	\$0.21	\$1.11		6.8	3.3	96		High	Very Low	
Site 3	\$1.05	\$0.47	\$0.12	\$1.07		5.9	2.3	81		High	Low	
Site 4	\$0.75	\$0.35	\$0.12	\$0.61		6.1	2.7	85		High	Very Low	
Site 5	\$0.27	\$0.13	\$0.05	\$0.09		4.3	2.4	59		Low	High	
Site 6	\$1.82	\$0.77	\$0.11	\$1.17		10.2	4.3	142		High	Low	
Site 7	\$1.23	\$0.54	\$0.11	\$0.98		15.0	4.4	192		High	Low	
Site 8	\$0.55	\$0.26	\$0.10	\$0.16		6.9	2.6	95		Med.	Low	
Site 9	\$1.01	\$0.44	\$0.10	\$0.83		5.3	2.5	74		High	Very Low	
Bricks-and-Mortar (US)	\$0.58	\$0.70	\$0.12	\$0.83		9.0	2.1	127		High	Very Low	



CSD Summary

- A systematic approach for representing and assessing sustainability of Clouds
 - Derived from a comprehensive model (economical, ecological, social)
- Automated, real-time Cloud Sustainability Dashboard
 - Express, assess and display run-time sustainability of Cloud & Cloud services
 - Preference-based customization
- Opportunities for integration with different enterprise tools



Lessons Learned

- Entered 4th year, growing, 15 members, 6 summits, IEEE Computer paper
 - 16 HPL projects across different Labs, similar in other institutions
 - Of them, 4 directly OpenCirrus related, most remaining Cloud related
 - *Sujata Banerjee, “Our NSDI paper would not have happened without Open Cirrus”*
 - *Rick McGeer, “We couldn’t have done demo at GEC-8 without Open Cirrus”*
 - *Prakash Reddy, “BookPrep prepared 1M books on Open Cirrus on demand”*
 - *Jerry Liu “Open Cirrus enabled us to build performance models for Article Clipper”*
- ++ A lot of external excitement, people “get it,” want to join, use it, associate with it
- + Exhaustive use by individual sites, excellent setting for external collaboration
- Feels like a research project but not resourced like one
- – Getting hardware without development resources stifles progress (HP, UIUC, IDA)



Summary

- Successful as a community
- Ahead of the time as a Cloud Stack
- Contributions internally, **less so across organizations**
- Created a lot of IP and enabled a lot of research

- Invitation to Open Cirrus Summits
 - *5th Summit in Moscow (June), 6th Summit in Atlanta (October)*
IEEE Co-sponsored, papers will appear in IEEE Digital Library
 - 7th Summit in Beijing, China, June 21-22, 2012
 - 8th Summit in Palo Alto, CA, collocated with ICAC, Fall 2012

