RED HAT SUMMIT

BOSTON, MA
JUNE 23-26, 2015

THE MASSACHUSETTS OPEN CLOUD

Orran Krieger, Principal Investigator, Boston University
Jan Mark Holzer, CTO Office, Red Hat
Brent Holden, Chief Architect, Red Hat
June 24th, 2015

BOSTON UNIVERSITY | VE RI TAS | MGHPCC MASSACHUSETTS GREEN HIGH PERFORMANCE COMPUTING CENTER | MIT | Northeastern | UMASS

MASSACHUSETTS TECHNOLOGY COLLABORATIVE

Today’s clouds are owned operated, controlled by a single provider

- Limiting research, innovation by third parties
  - ➡technology comp
  - ➡performance ser above
- No visibility/auditing operations:
  - ➡Major security ch critical datasets
- Vendor lock in by features, interfaces, and pricing model

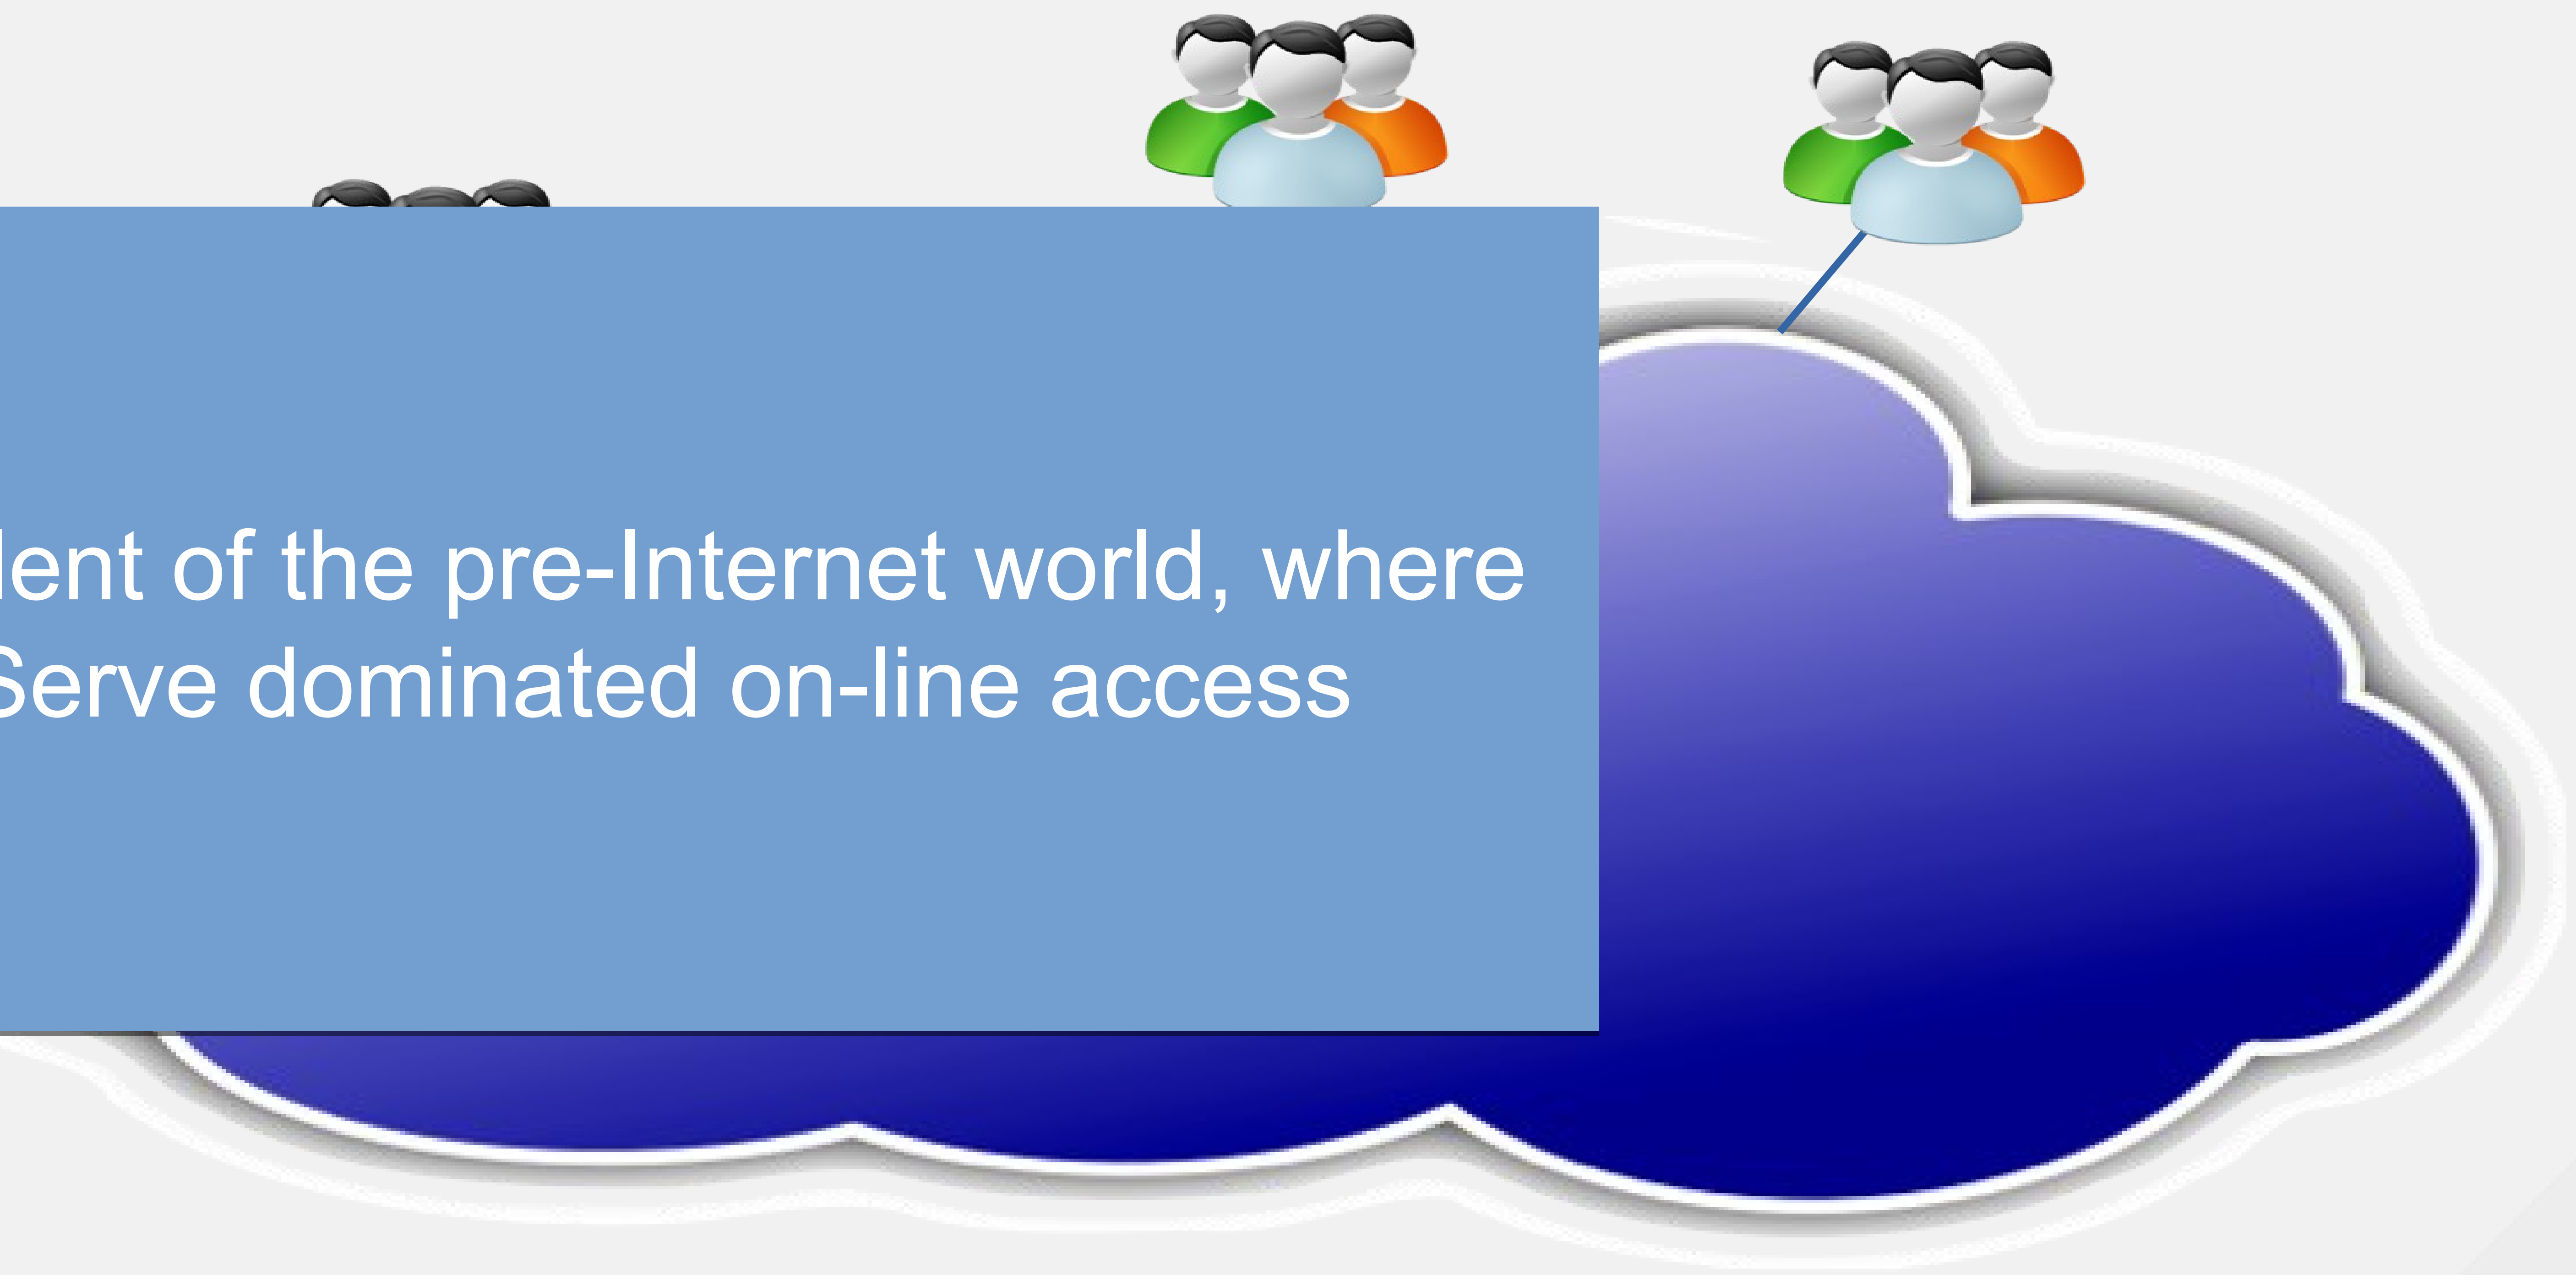

We are in the equivalent of the pre-Internet world, where AOL and CompuServe dominated on-line access

THE VISION

A new model is required: an “Open Cloud eXchange (OCX)”

- Multiple “partners” participate in implementing and operating cloud:
  - eXchange Service Providers (XSP)
- Each XSP determines how to charge for her services
- Customers can select and move between services
- Domain specific intermediaries:
  - provide customers with simple model
  - enable optimization
- Multi-sided marketplace

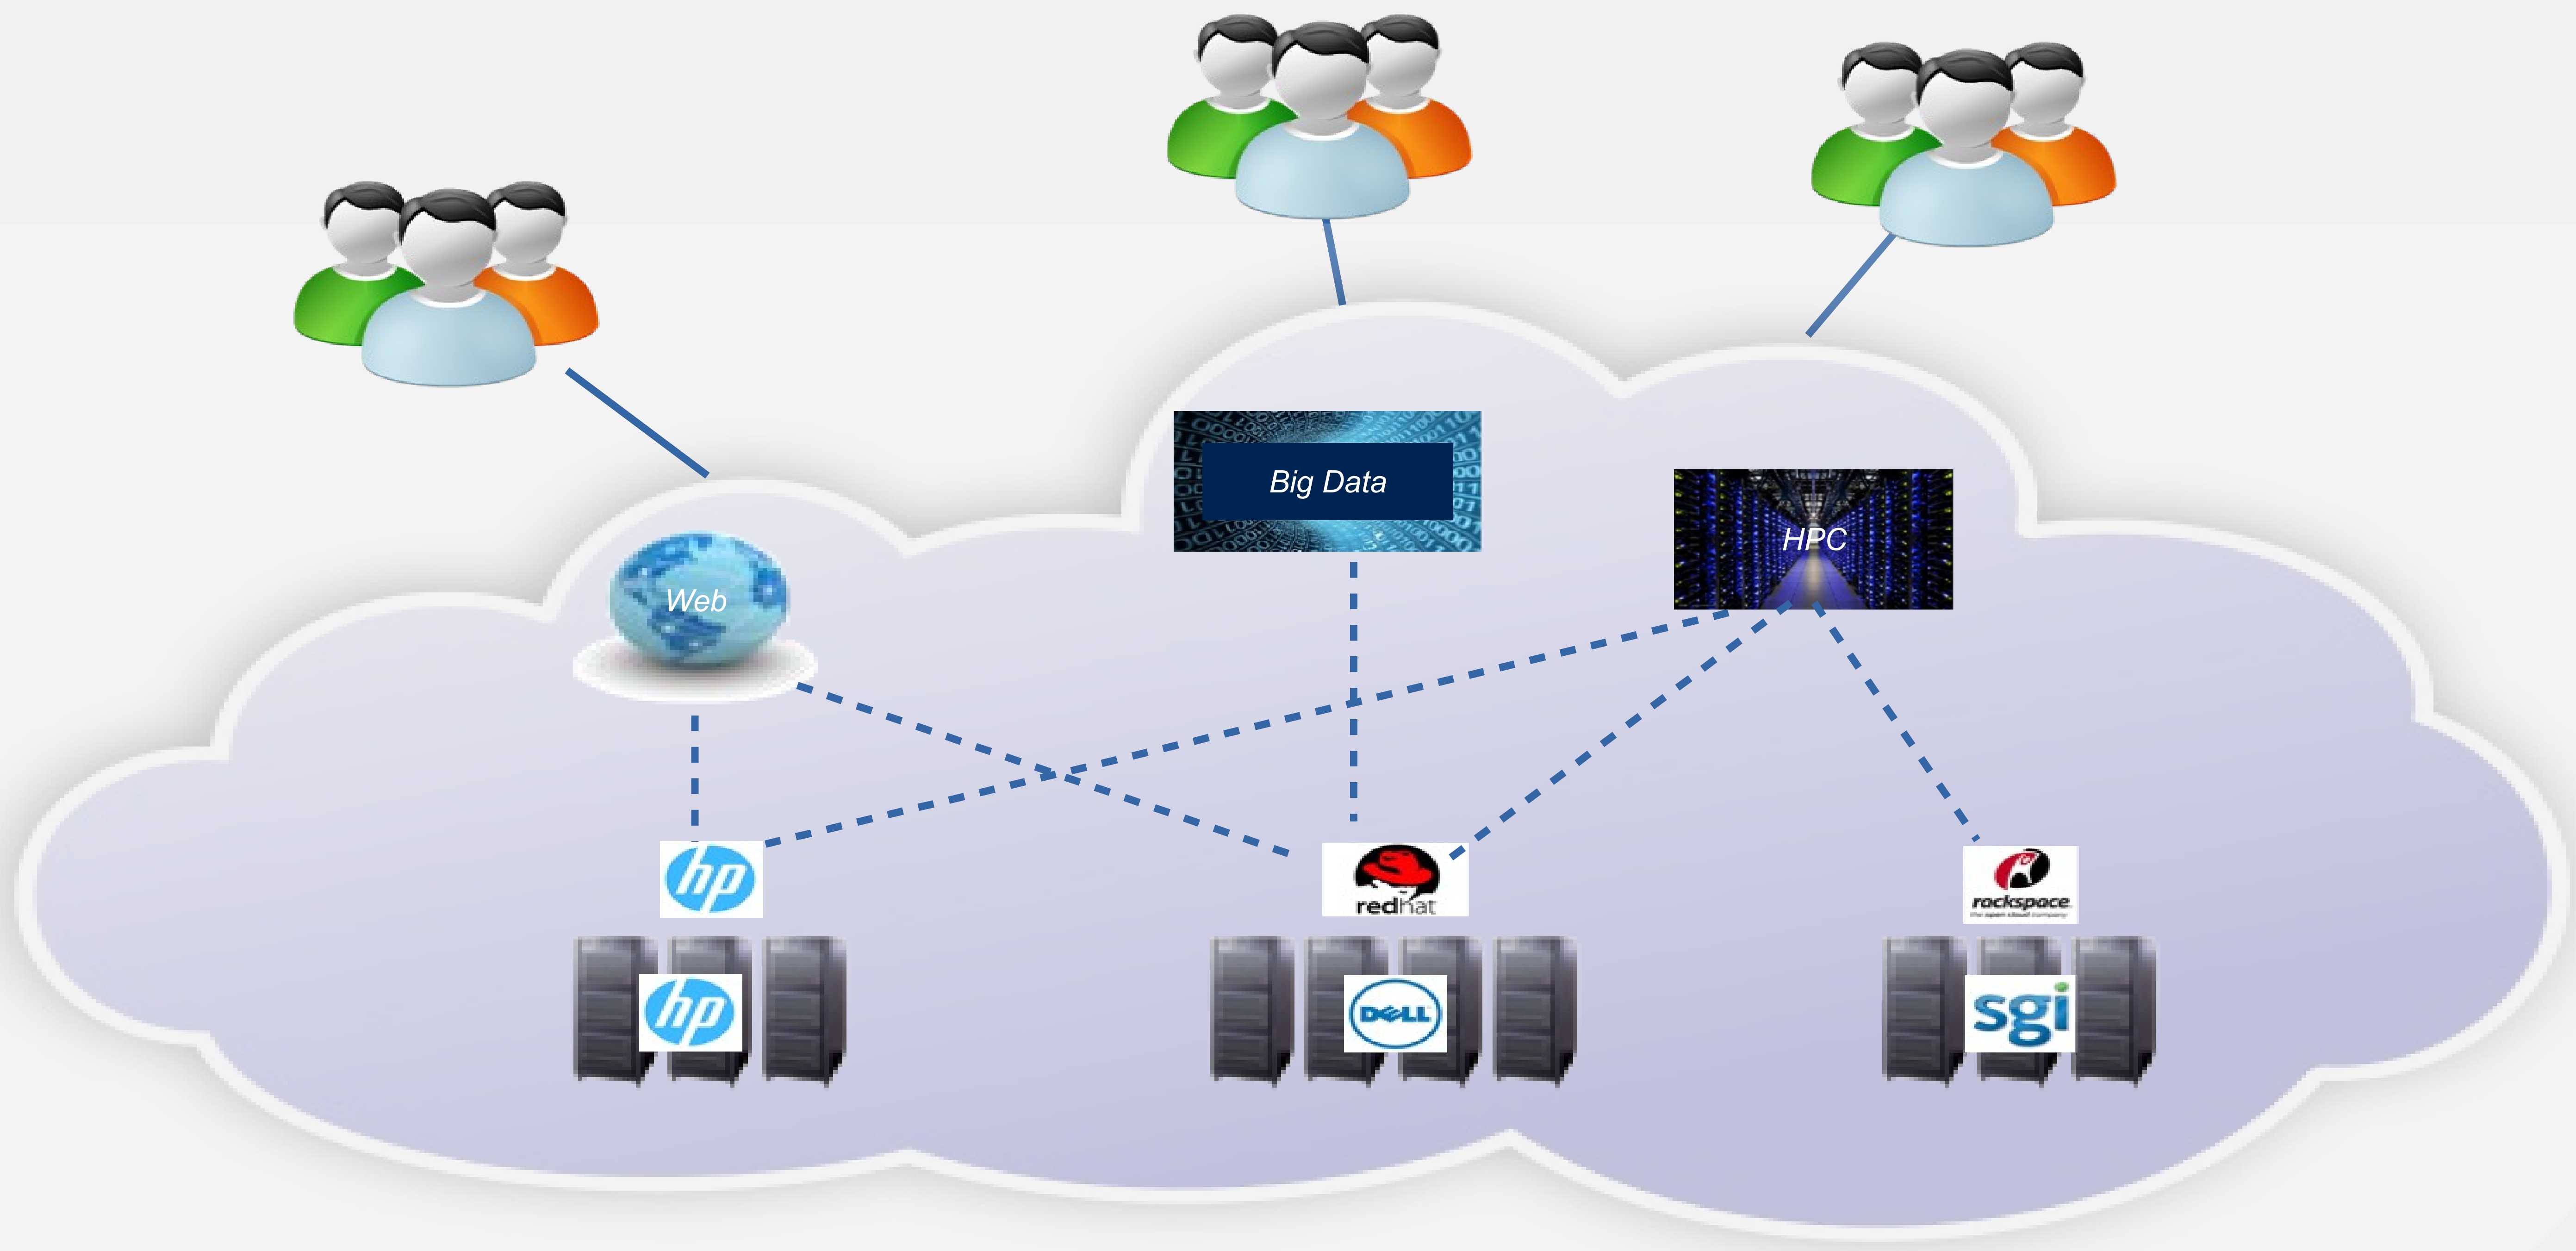

USE CASES

Examples

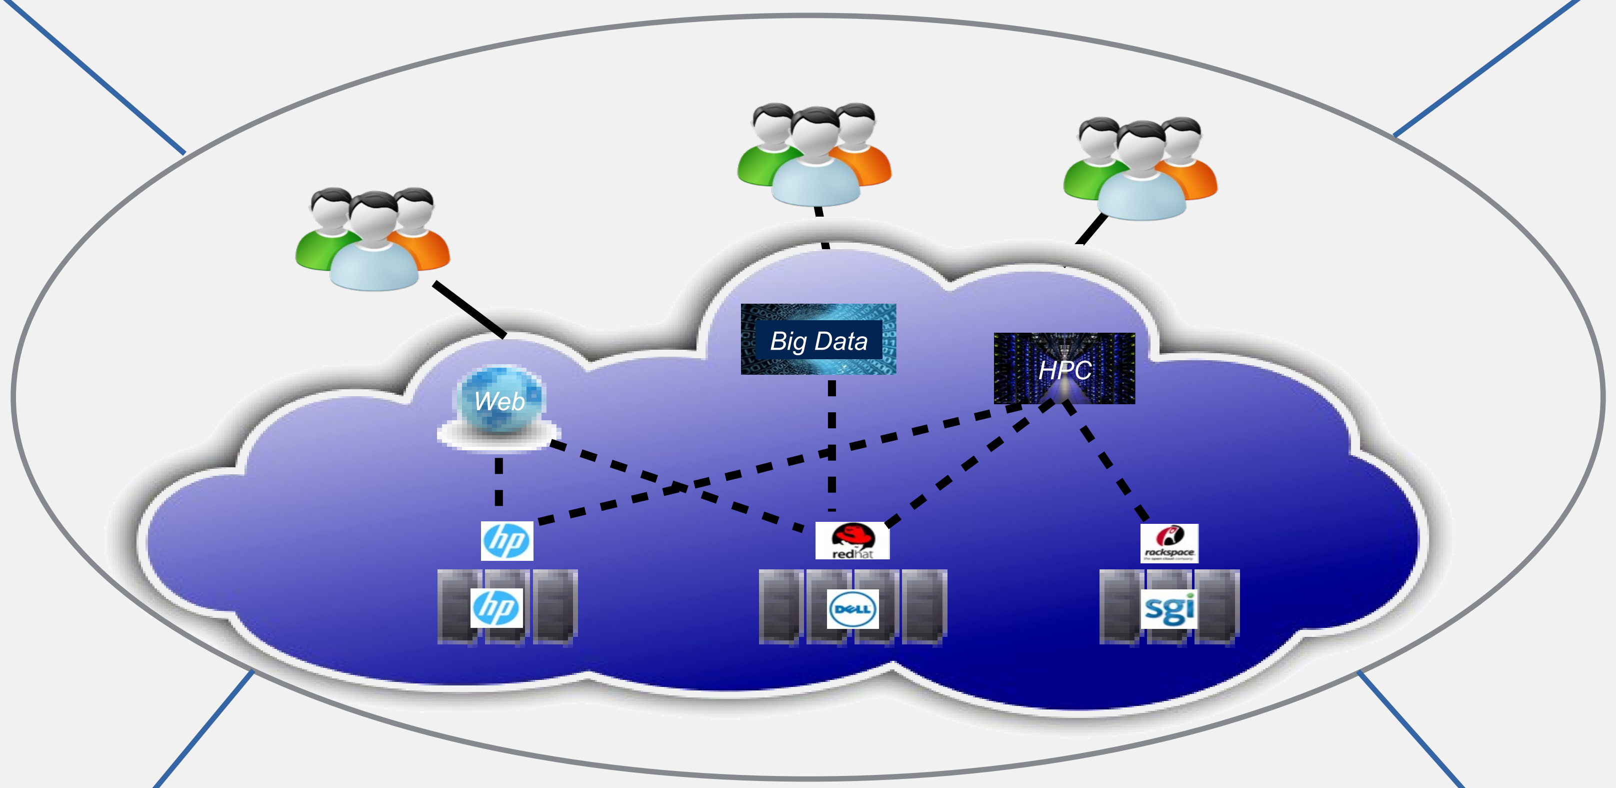

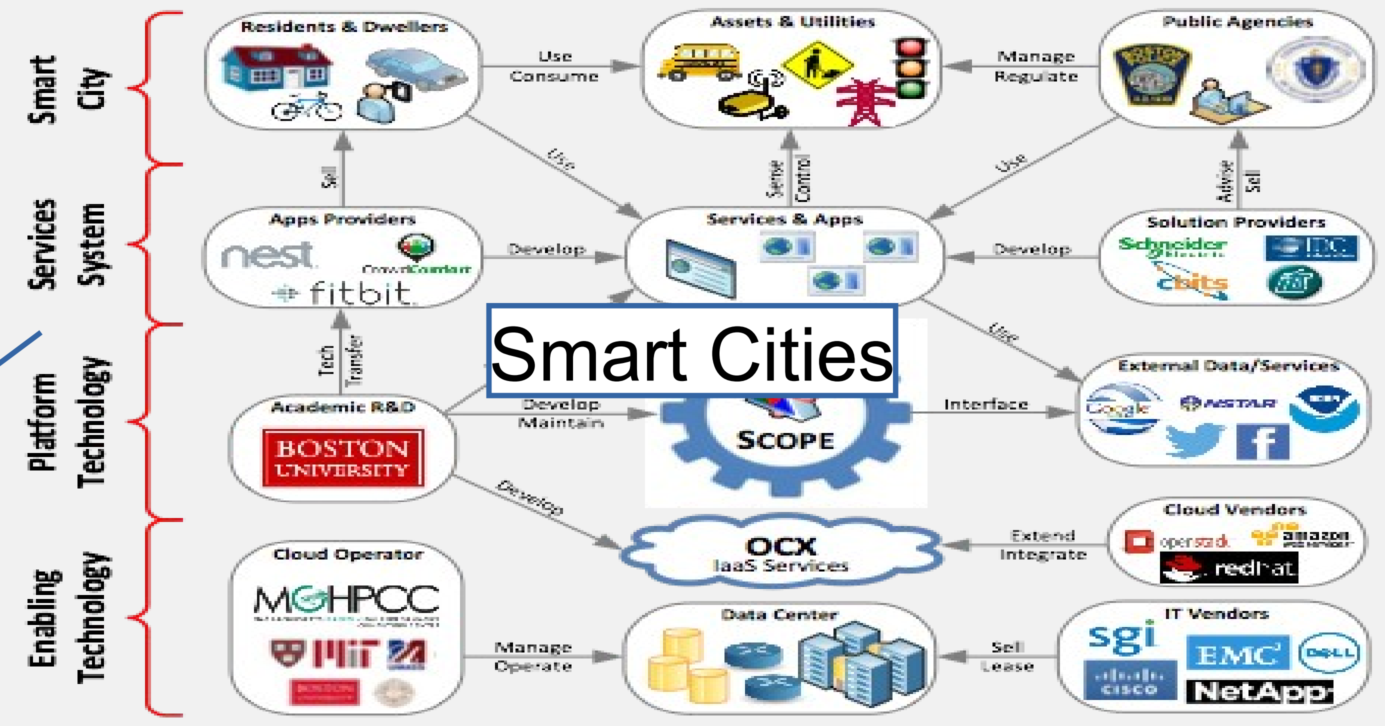

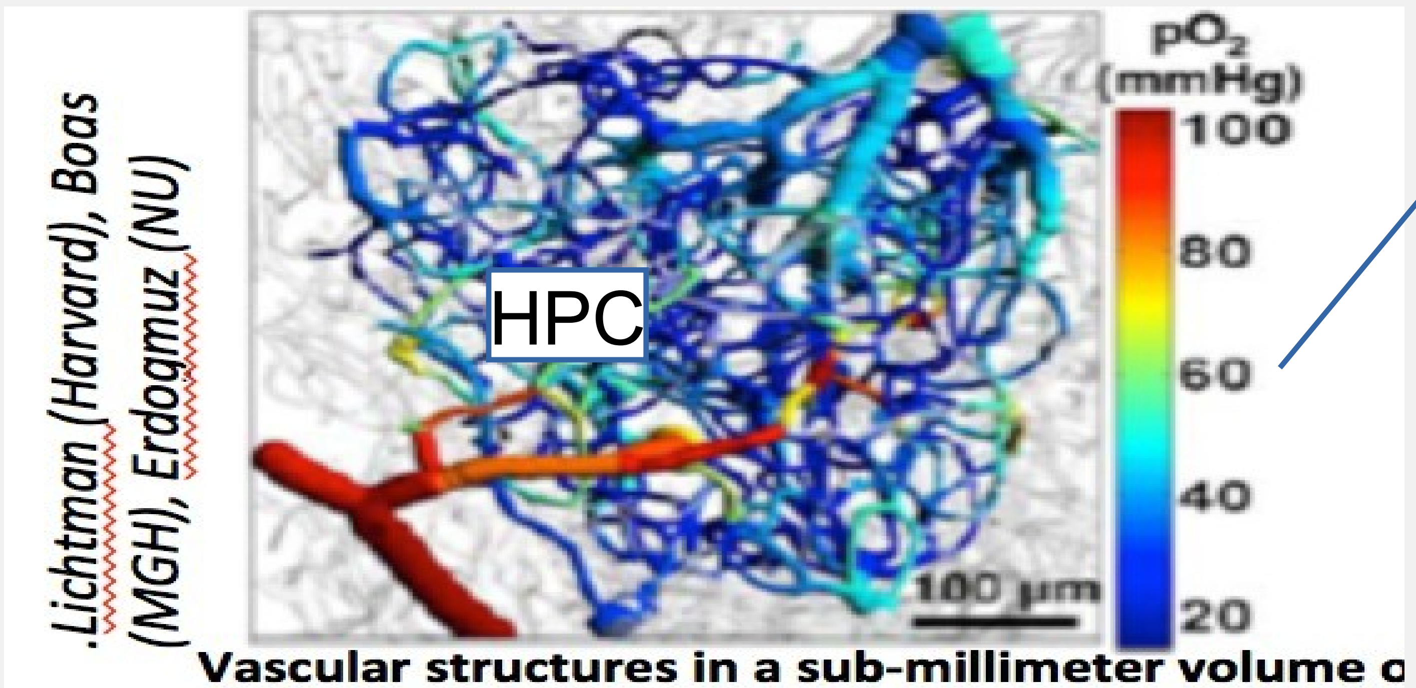

Vascular structures in a sub-millimeter volume o neural (mouse) tissue

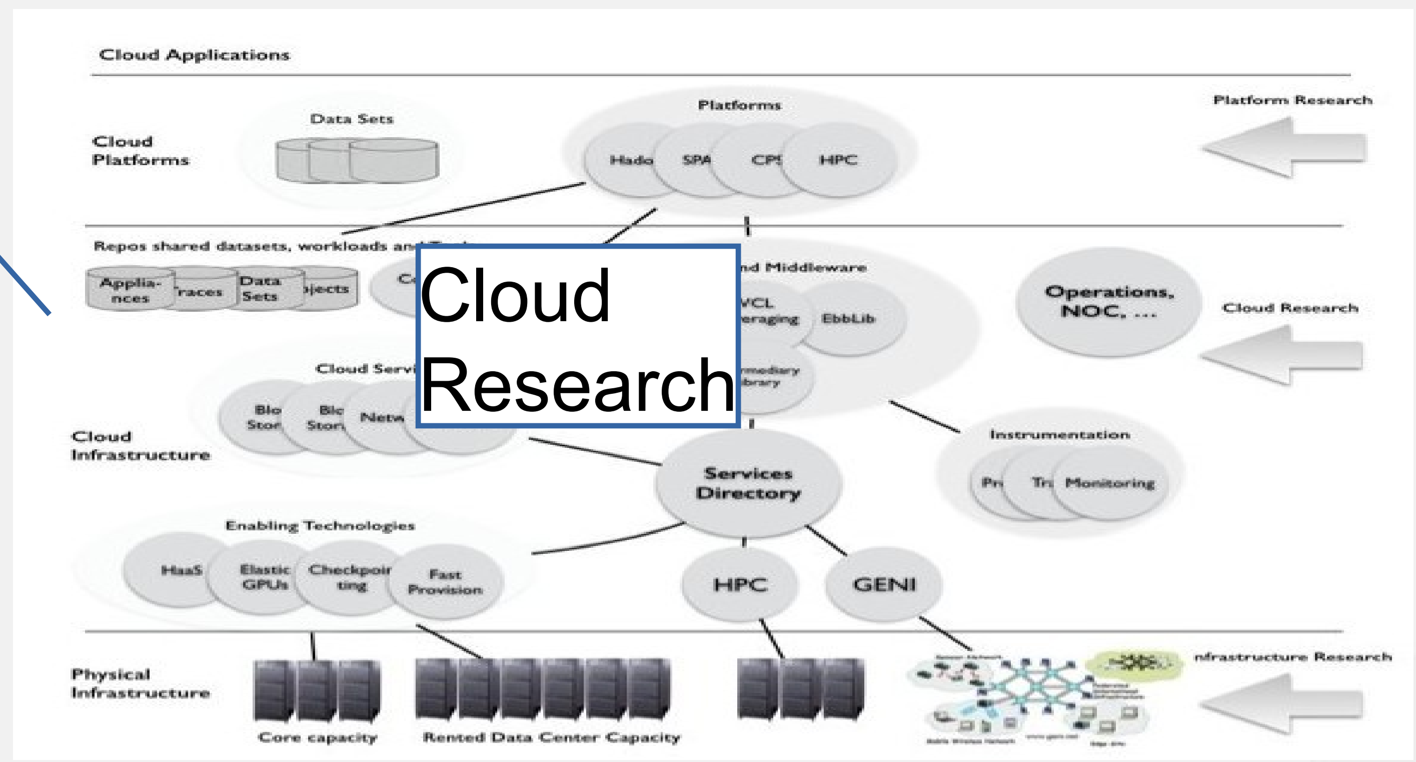

Examples


Spark
hadoop
VERTICA
An HP Company
redhat
Industry
sgi
GlusterFS

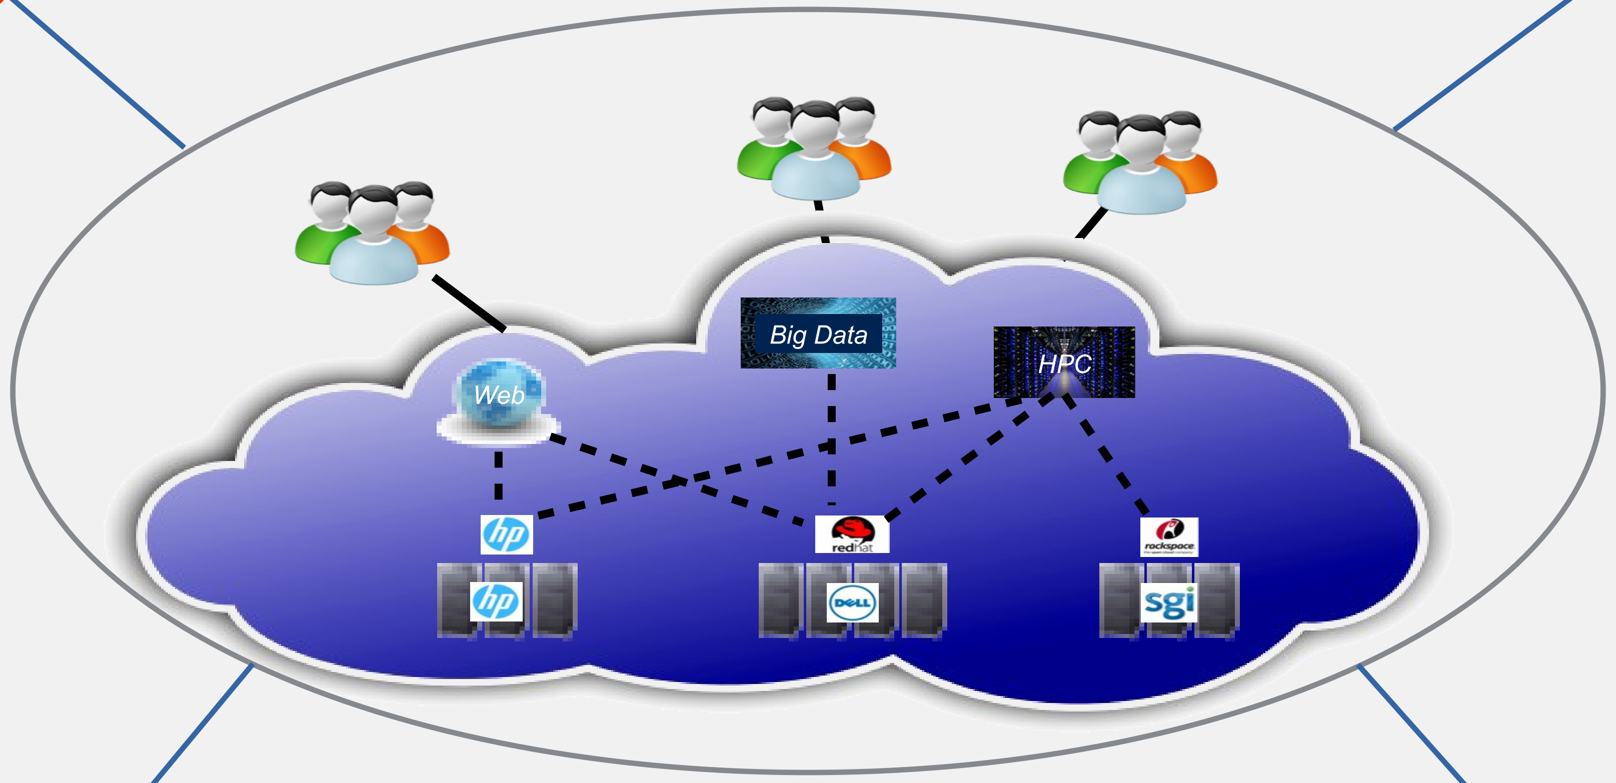
Big Data
Web
HPC
sgi

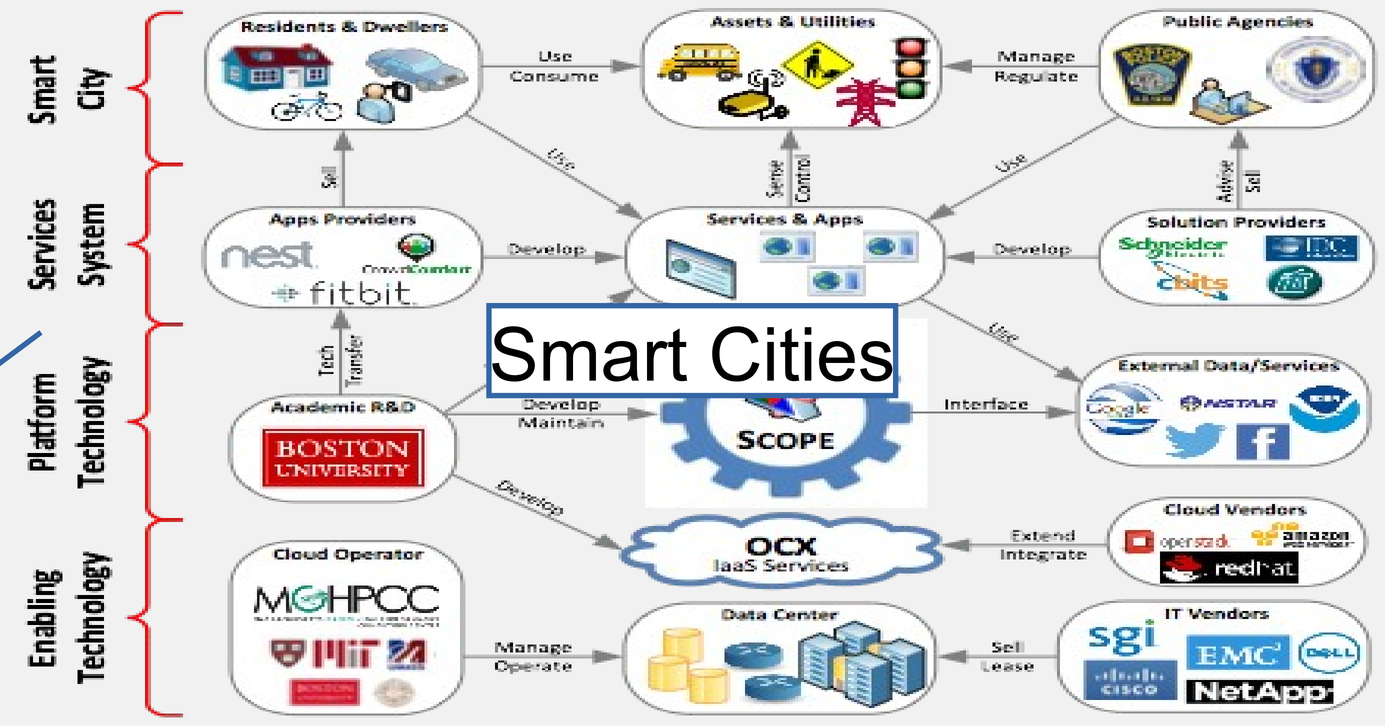
Smart City
Services System
Platform Technology
Enabling Technology
Residents & Dwellers
Use Consume
Assets & Utilities
Manage Regulate
Public Agencies
Apps Providers
Develop
Services & Apps
Develop
Solution Providers
nest
fitbit
Smart Cities
Academic R&D
BOSTON UNIVERSITY
Develop Maintain
SCOPE
Interface
External Data/Services
Cloud Operator
MGHPCC
OCX
IaaS Services
Extend Integrate
Cloud Vendors
redhat
Data Center
Manage Operate
Sell Lease
IT Vendors
sgi
EMC
NetApp

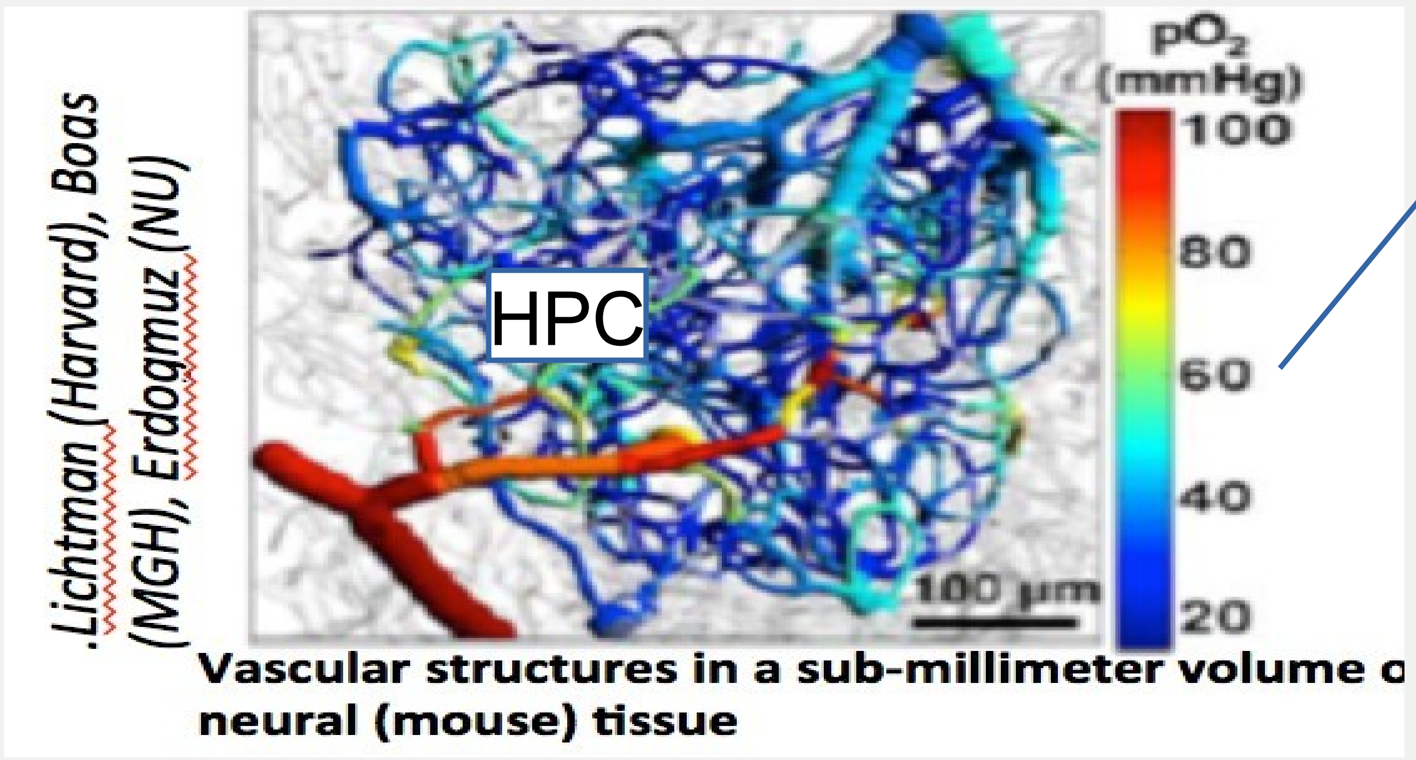
pO2 (mmHg)
100
80
60
40
20
100 μm
HPC
Lichtman (Harvard), Boas (MGH), Erdogmuz (NU)
Vascular structures in a sub-millimeter volume o
neural (mouse) tissue

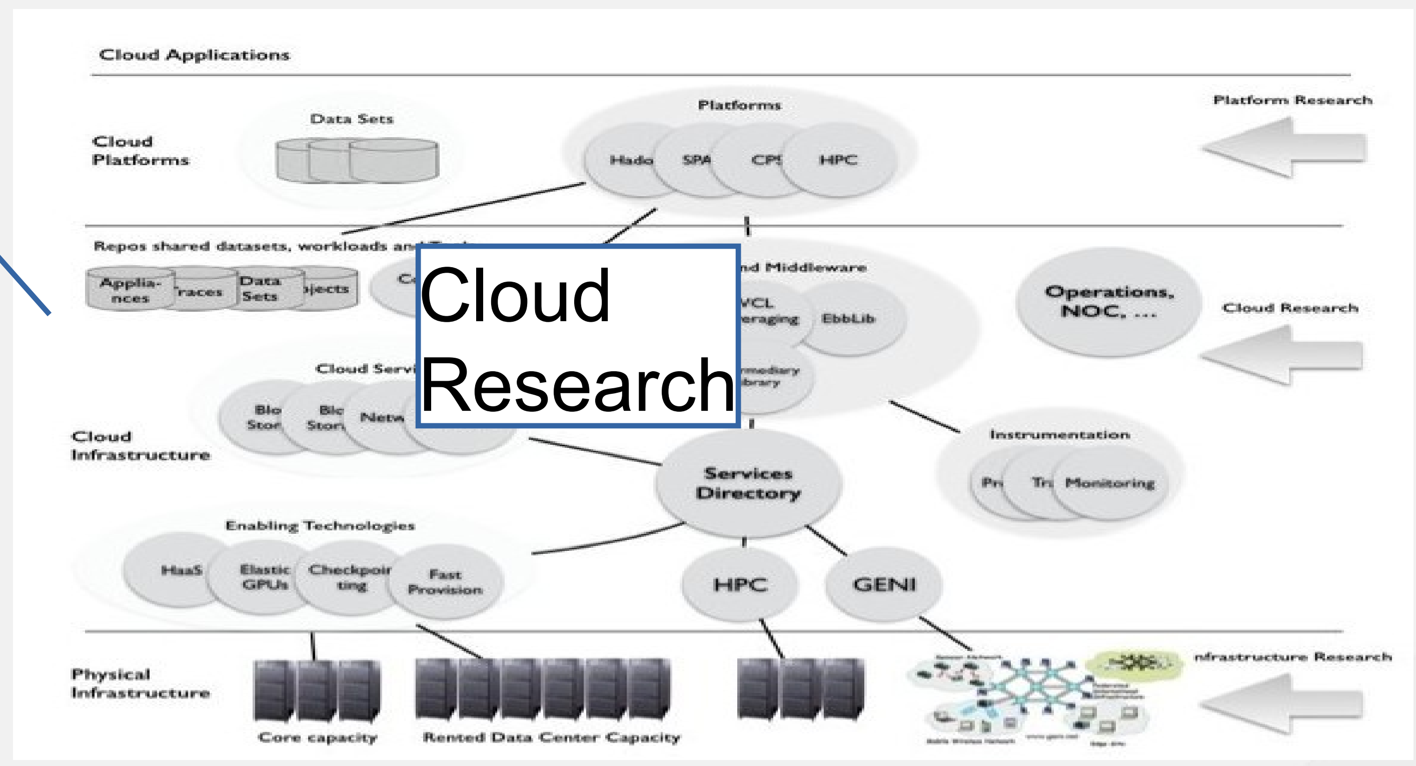
Cloud Applications
Data Sets
Platforms
Platform Research
Cloud Platforms
HPC
Cloud Research
Operations, NOC, ...
Cloud Infrastructure
Services Directory
Instrumentation
Monitoring
Enabling Technologies
HaaS
Fast Provision
HPC
GENI
Physical Infrastructure
Core capacity
Rented Data Center Capacity

Example: Industry

Use case by the CTO of Communispace:

We gather vast amounts of … quantitative data every day on 100,000 consumers from 150 countries around the world – a classic "Big Data" opportunity. Rather than … a traditional (costly, lengthy, ineffective) vendor selection process … make some of our data available in the MOC and have vendors leverage that data to both perform their own R&D on real-word data sets, but also compete for the opportunity to become part of our analysis solution.

Example: Industry

- Communispace puts representative dataset in MOC
- HackReduce hosts hackathon
- Communispace purchases a cloud in a box from VCE (Cisco/EMC) to run “hourly review” application
- Demanding application uses Vertica with Plexxi low-latency network: stays on cloud
- Communispace, a Verizon customer, uses Juniper’s SDN solution to extend its network into the MOC

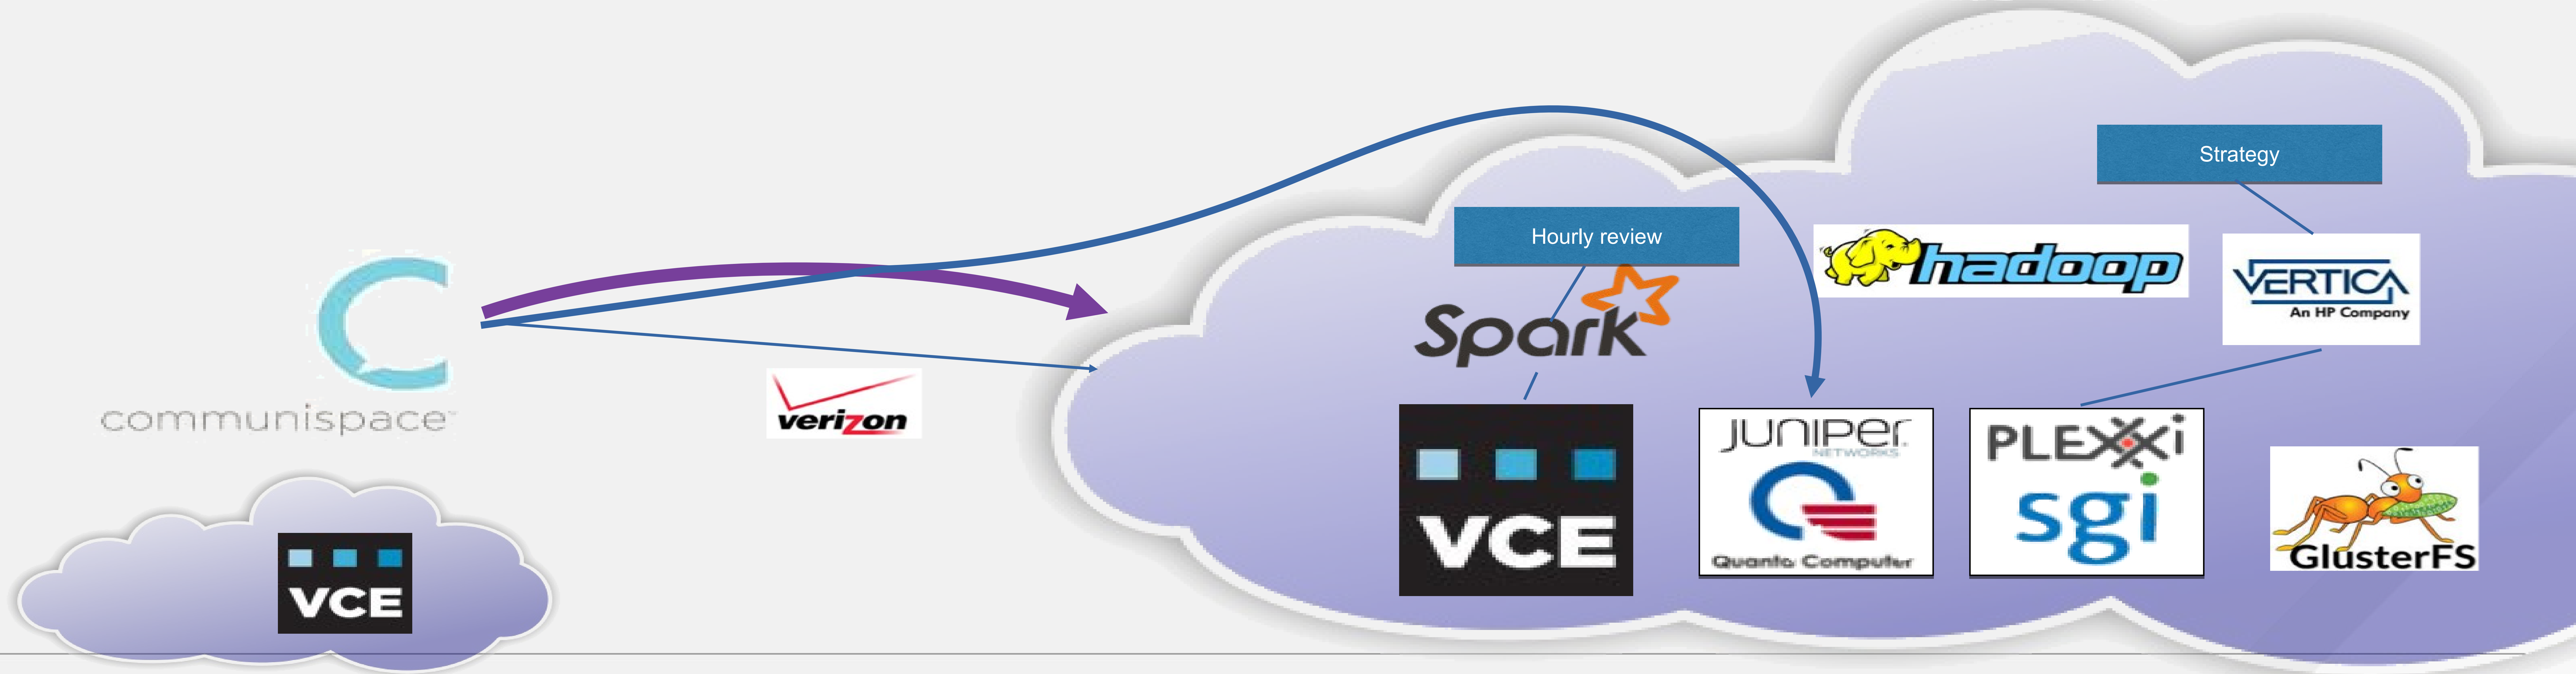

Examples

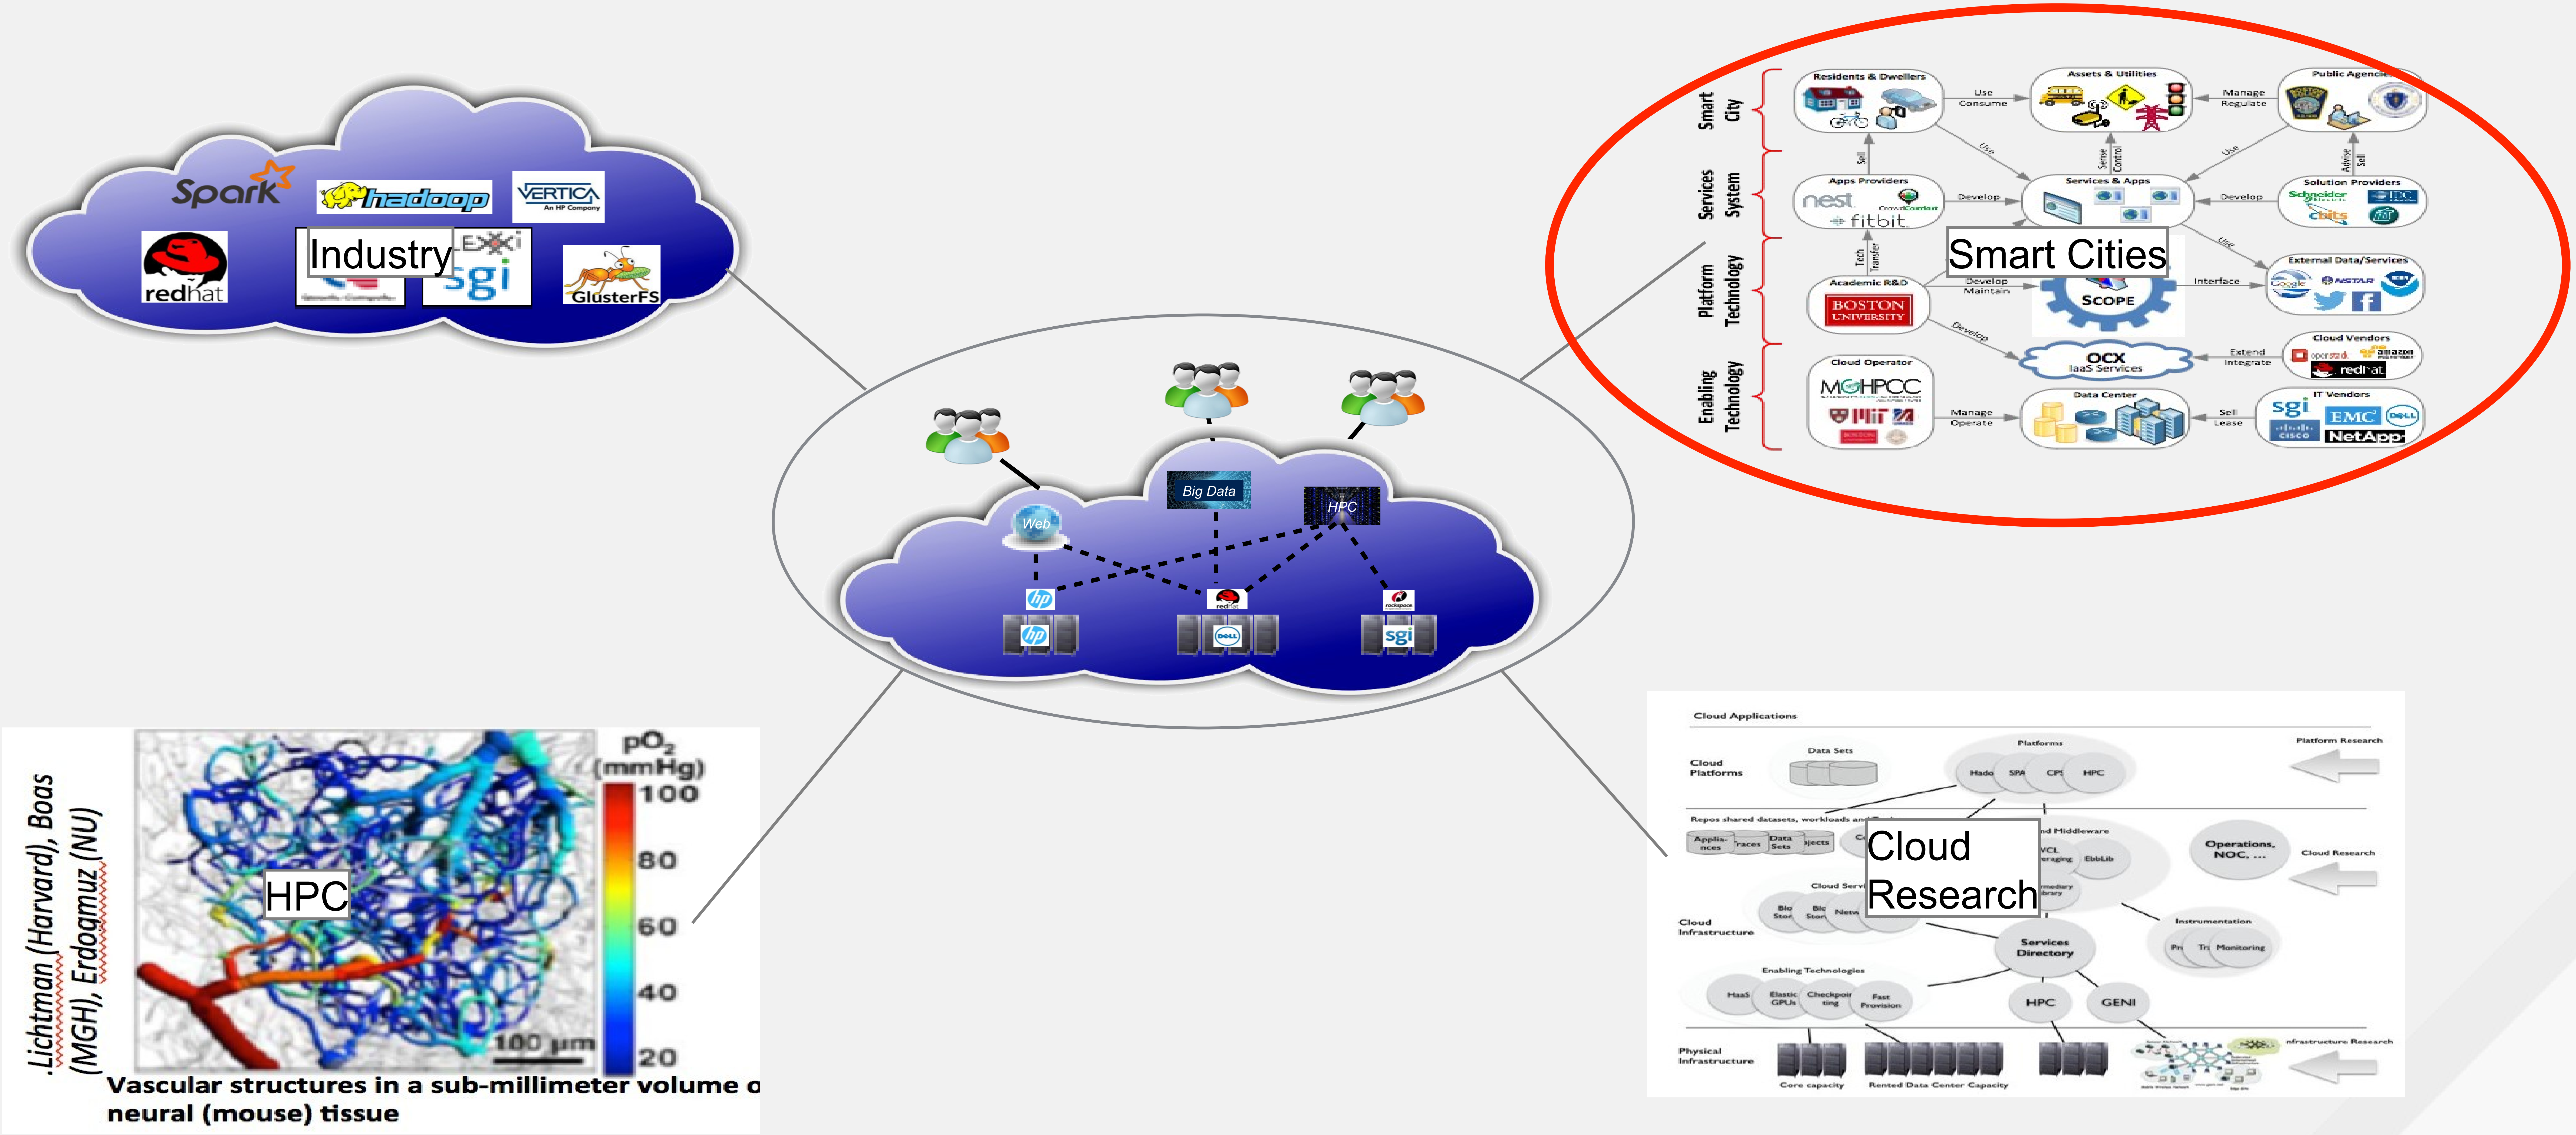
Industry
Spark
hadoop
VERTICA
redhat
sgi
GlusterFS
Web
Big Data
HPC
Smart Cities
Smart City
Services System
Platform Technology
Enabling Technology
Residents & Dwellers
Assets & Utilities
Public Agencies
Apps Providers
Services & Apps
Solution Providers
Academic R&D
SCOPE
External Data/Services
Cloud Operator
OCX
IaaS Services
Cloud Vendors
Data Center
IT Vendors
HPC
Lichtman (Harvard), Boas (MGH), Erdogmuz (NU)
pO_2 (mmHg)
100
80
60
40
20
100 μm
Vascular structures in a sub-millimeter volume of neural (mouse) tissue
Cloud Research
Cloud Applications
Cloud Platforms
Platforms
Platform Research
Cloud Research
Operations, NOC, ...
Cloud Infrastructure
Services Directory
Instrumentation
Enabling Technologies
HPC
GENI
Physical Infrastructure
Core capacity
Rented Data Center Capacity

Example: Smart cities

http://www.bu.edu/hic/research/scope/

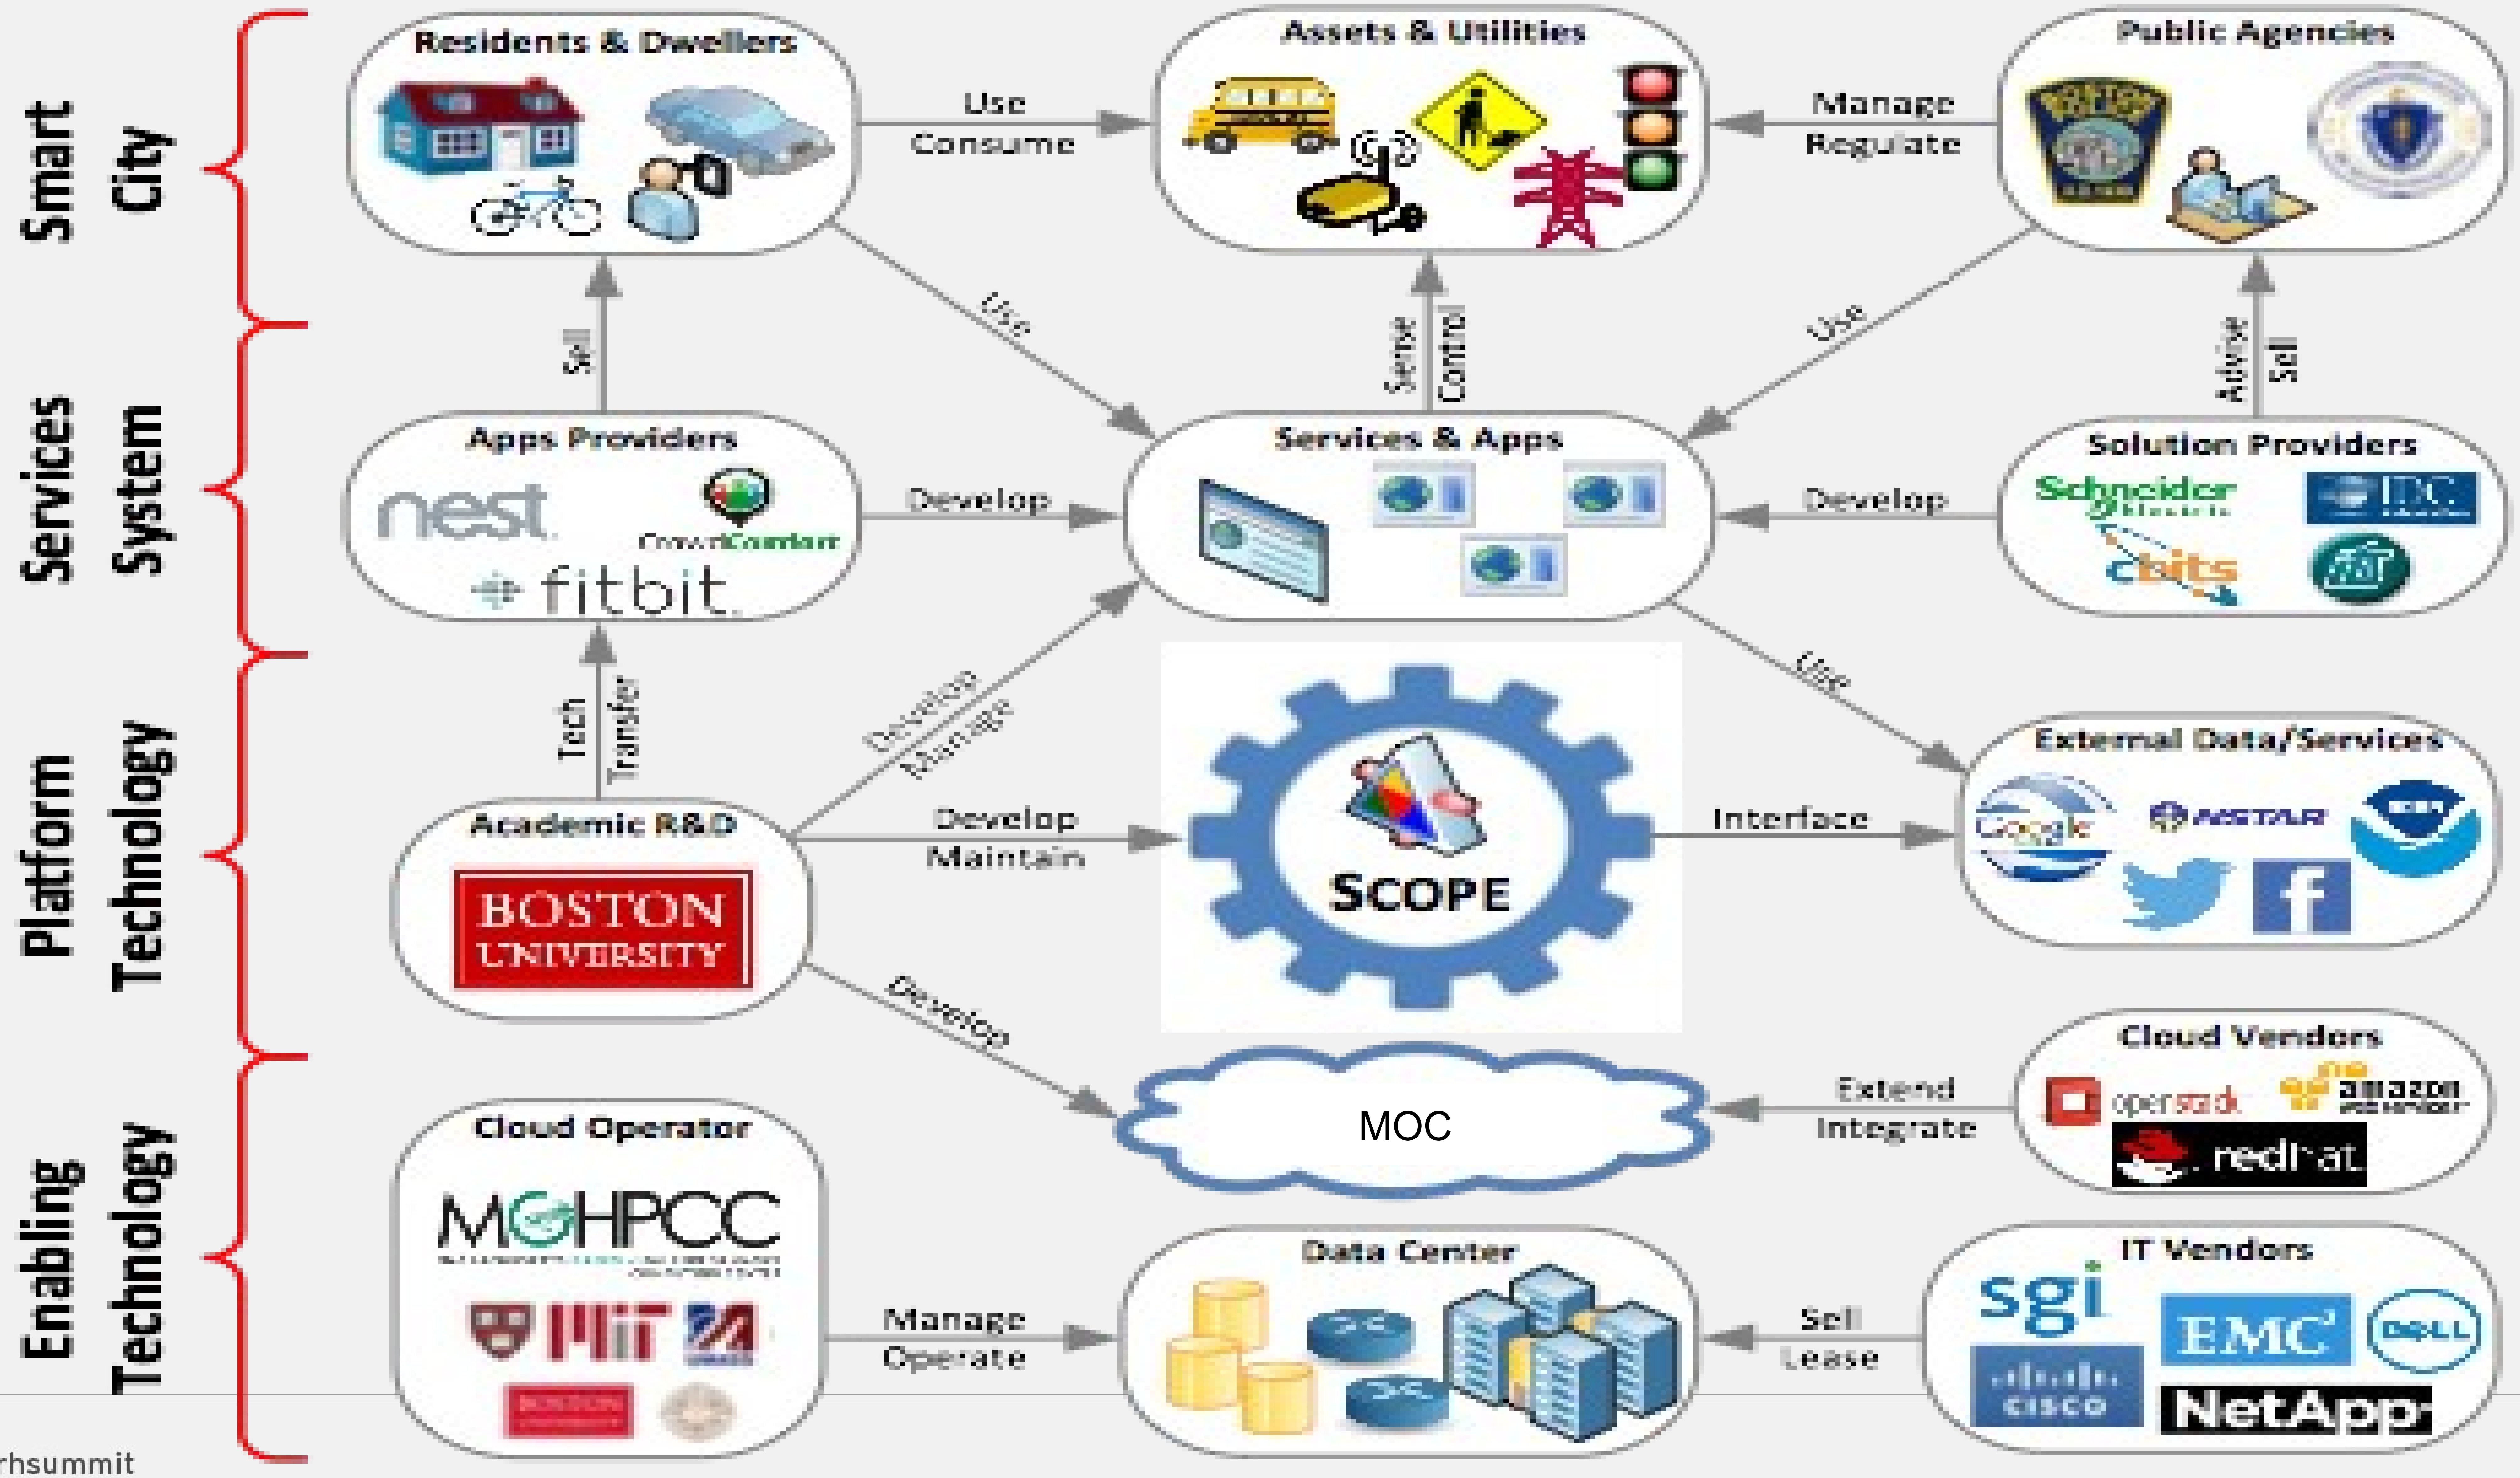

Examples


Spark
hadoop
VERTICA
Industry
redhat
sgi
GlusterFS

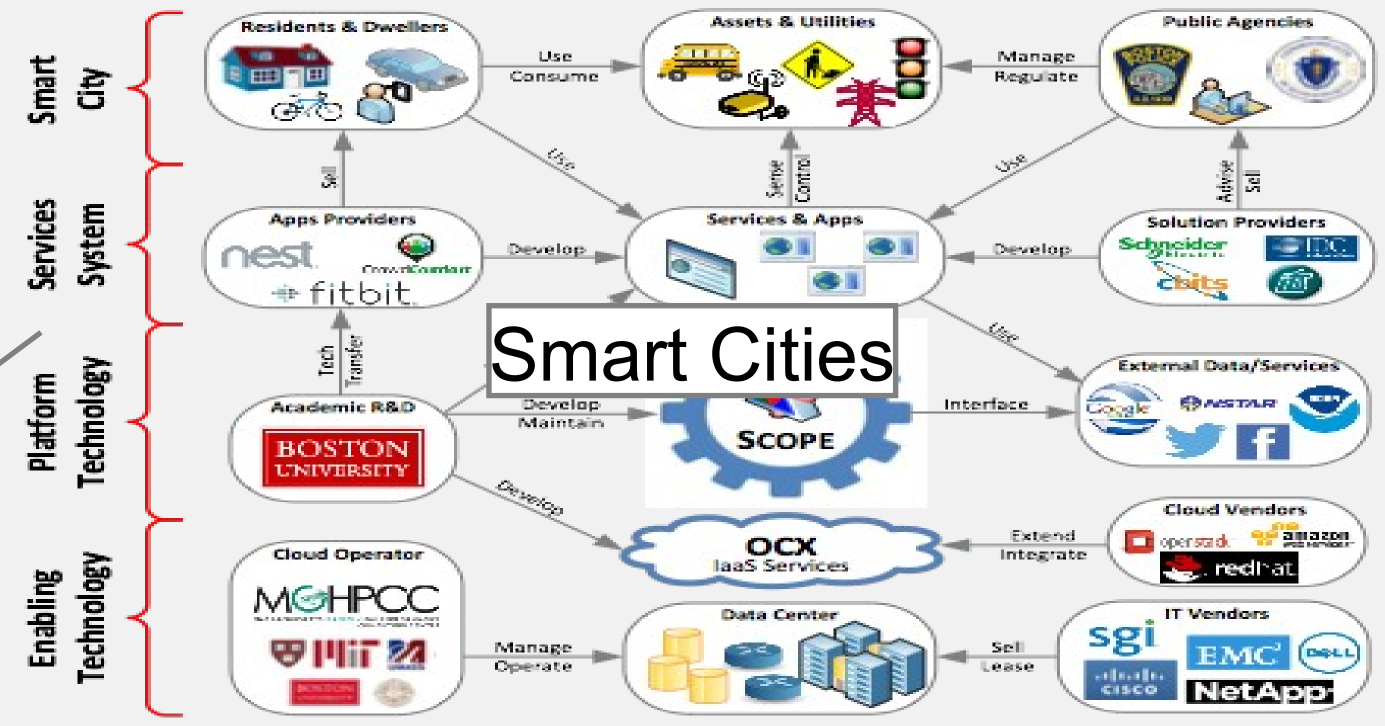
Smart City
Services System
Platform Technology
Enabling Technology
Smart Cities
SCOPE
OCX
IaaS Services
MGHPCC
BOSTON UNIVERSITY

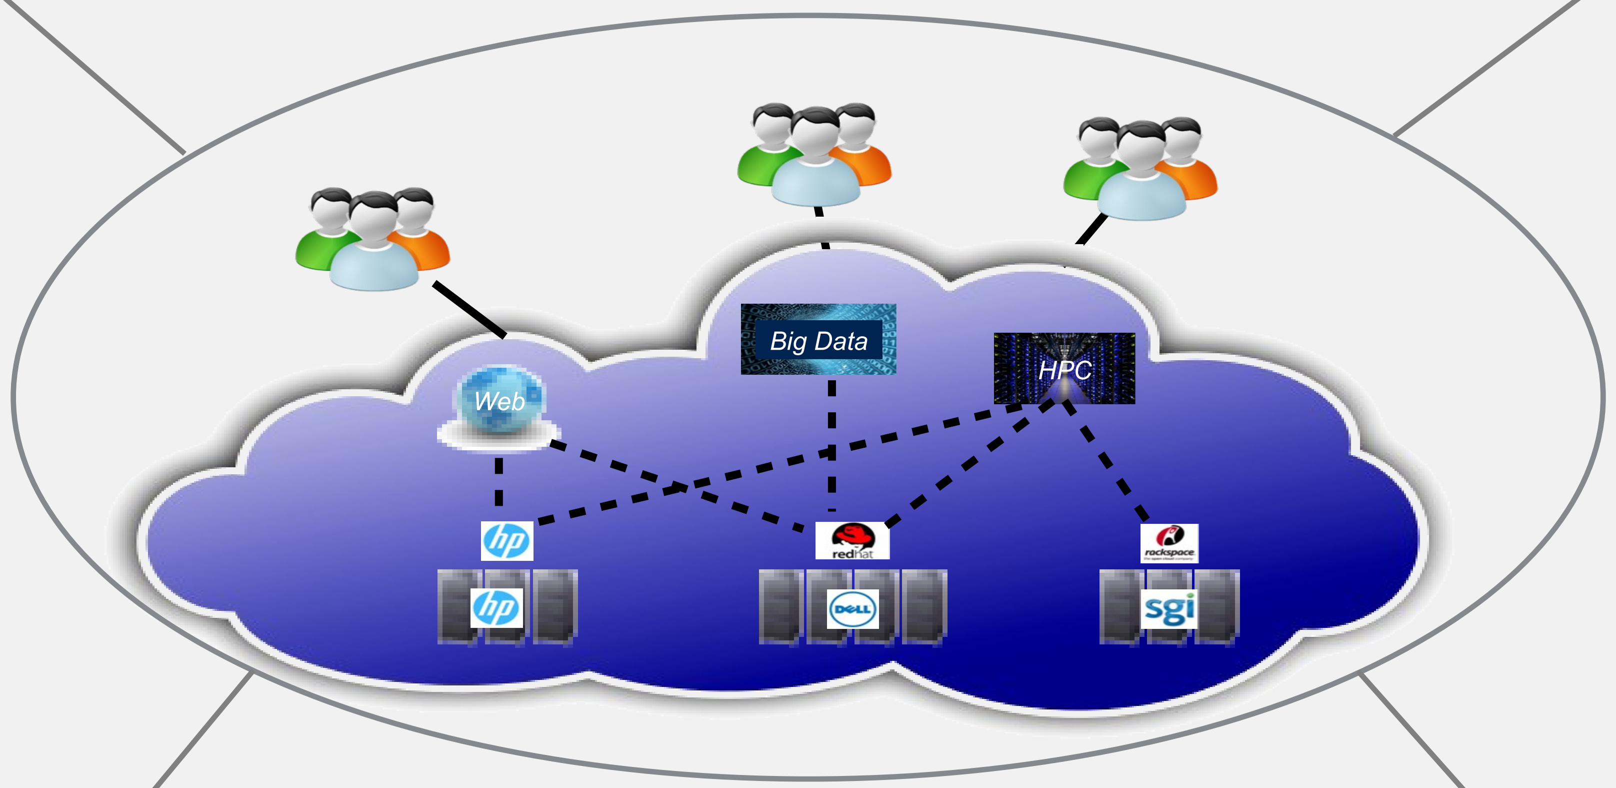
Web
Big Data
HPC
sgi

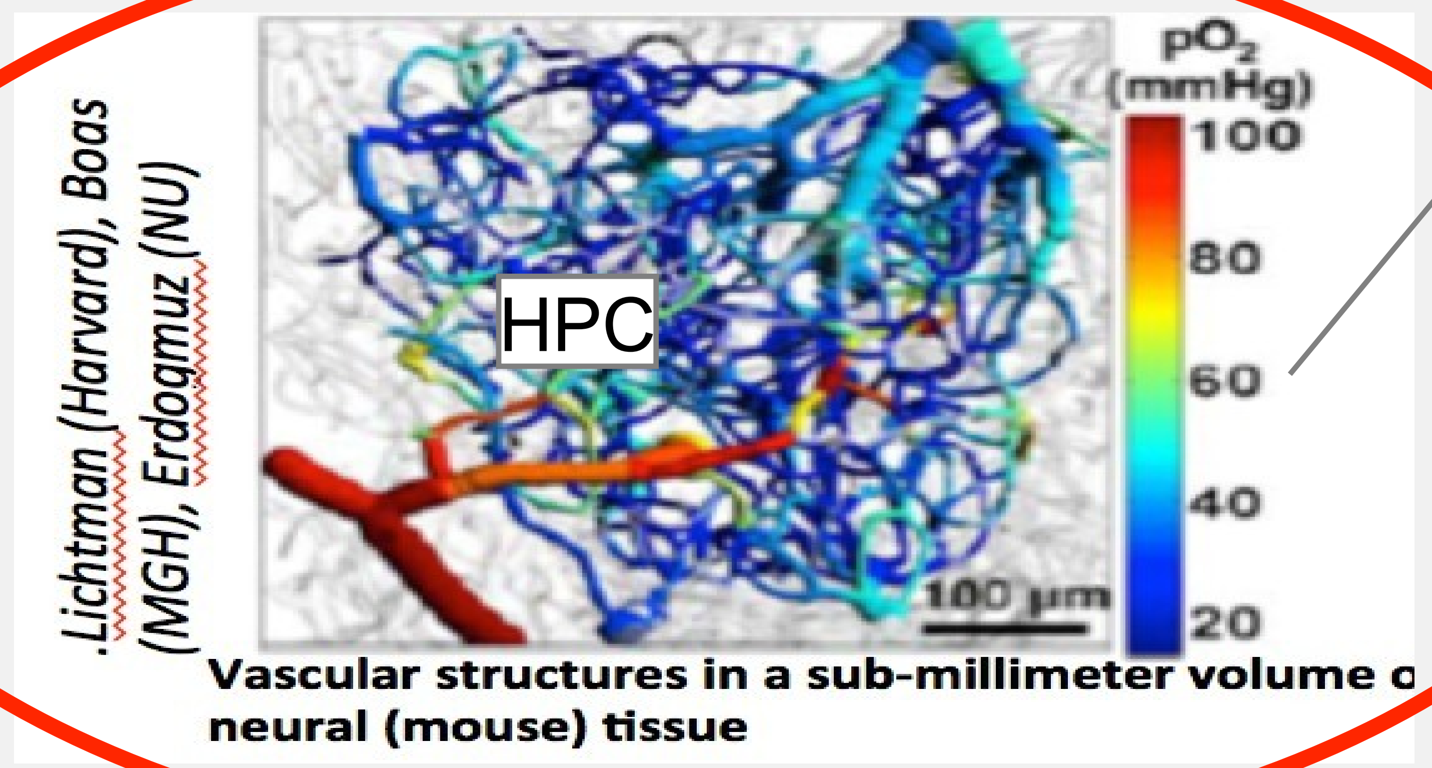
Lichtman (Harvard), Boas (MGH), Erdogmuz (NU)
HPC
pO_2 (mmHg)
100
80
60
40
20
100 µm
Vascular structures in a sub-millimeter volume of neural (mouse) tissue

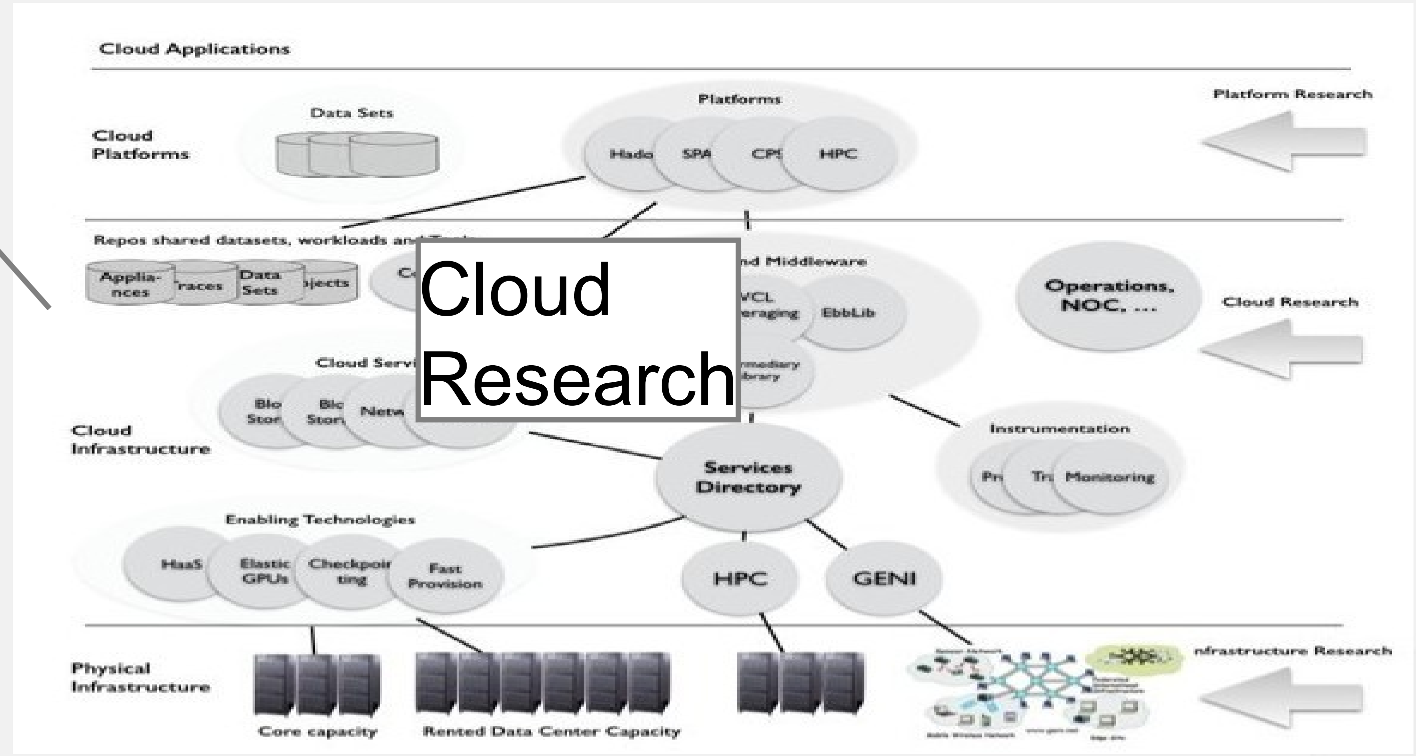
Cloud Applications
Cloud Platforms
Data Sets
Platforms
HPC
Platform Research
Cloud Research
Operations, NOC, ...
Cloud Infrastructure
Instrumentation
Monitoring
Services Directory
Enabling Technologies
HPC
GENI
Physical Infrastructure
Core capacity
Rented Data Center Capacity

HPC

- There is already an increasing convergence of cloud and HPC:
  - in 2011 and 2012, AWS and MSFT Azure fielded HPC Clusters based on their respective Cloud platforms that ranked 42nd and a 165th respectively.
  - full bi-sectional bandwidth networks to simplify scheduling and to enable new storage models.
  - GP-GPU, RDMA, SDN enabling bare metal provisioning
- An OCX could do much more:
  - mixing HPC and cloud applications drives up utilization enormously better economics
  - support whole new classes of interactive supercomputing applications

Supporting Interactive, Bursty HPC Applications

Elastic Software

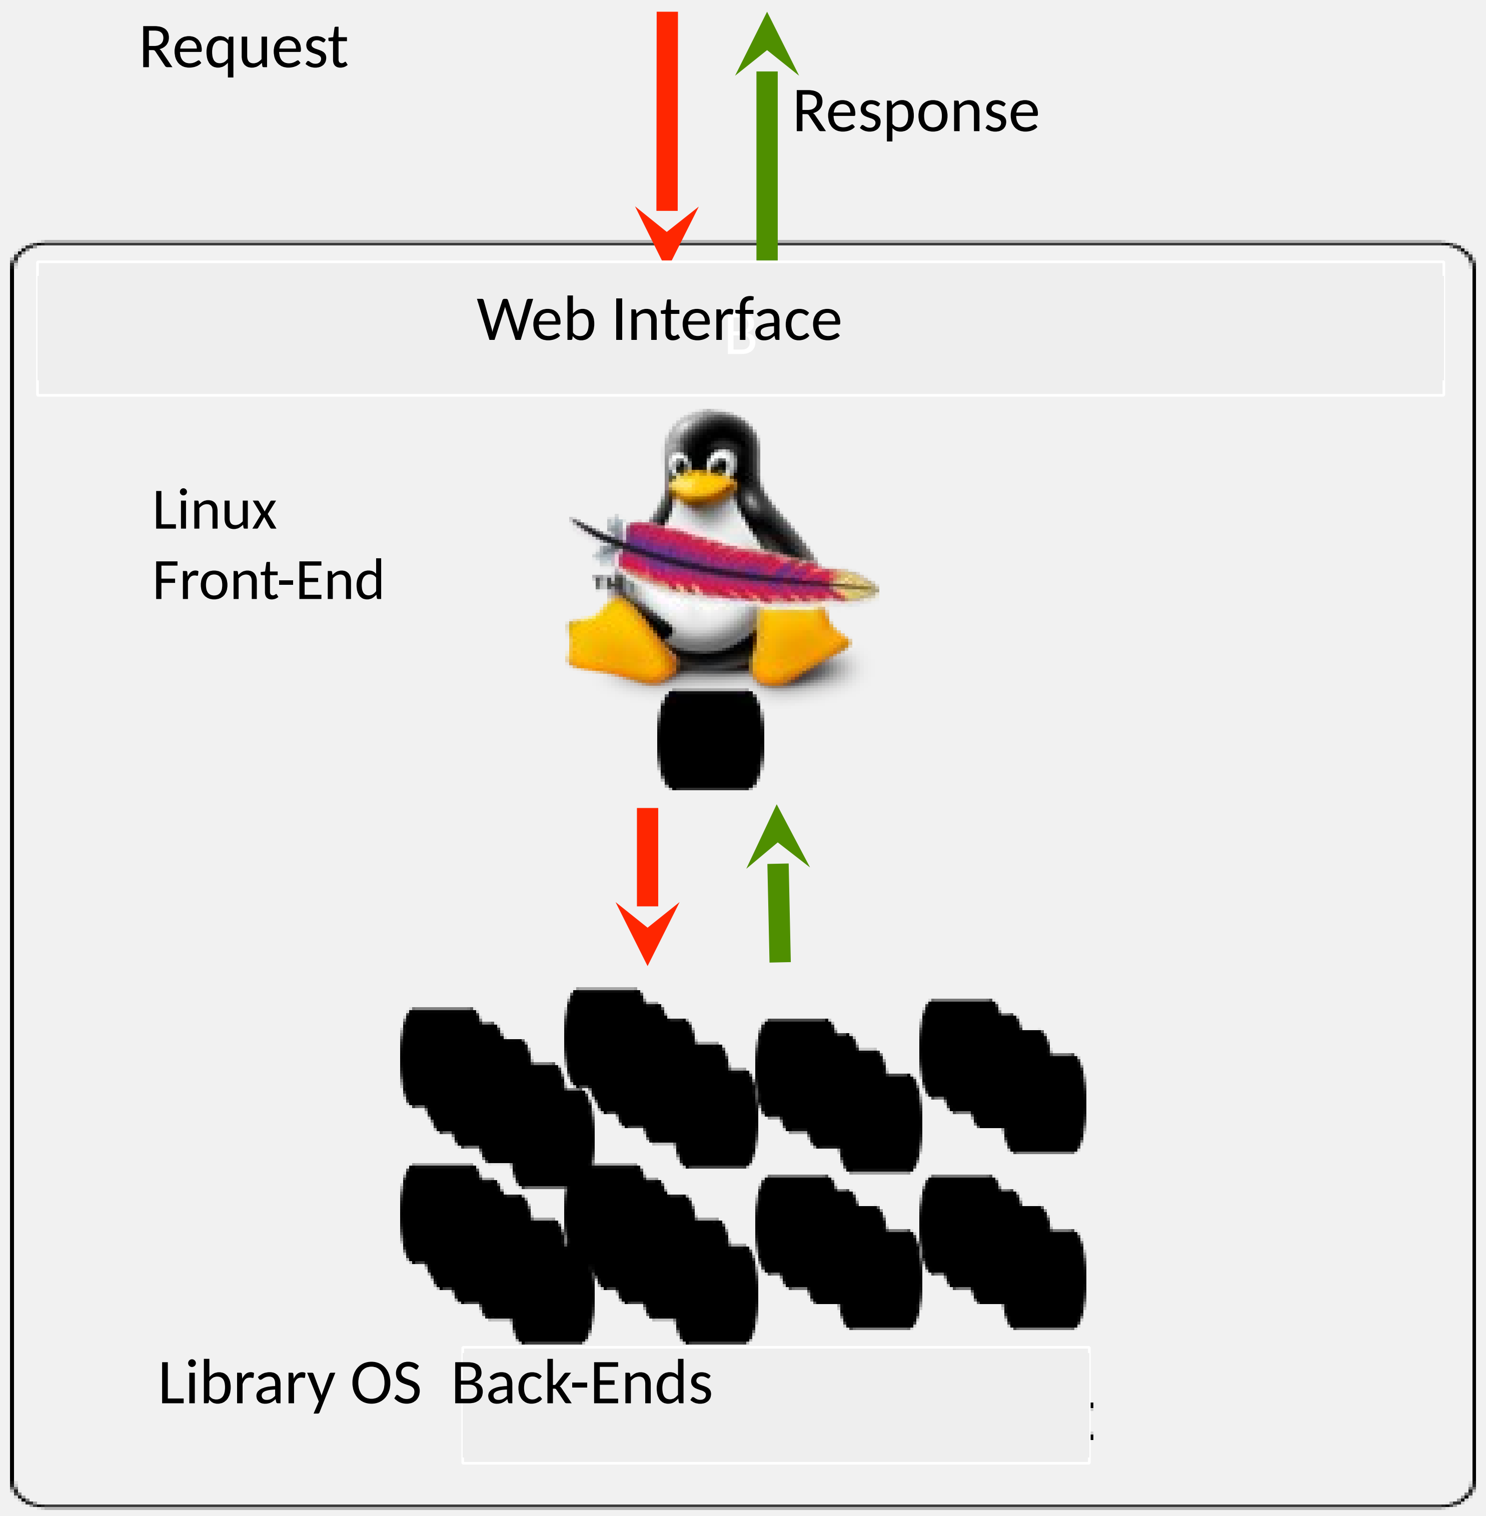

EbbRT distributed library OS [Appavoo BU]:

- Front-end Linux allocates bare-metal back-end nodes on demand
- Back-end nodes library OS customized to single application needs

Elastic Infrastructure

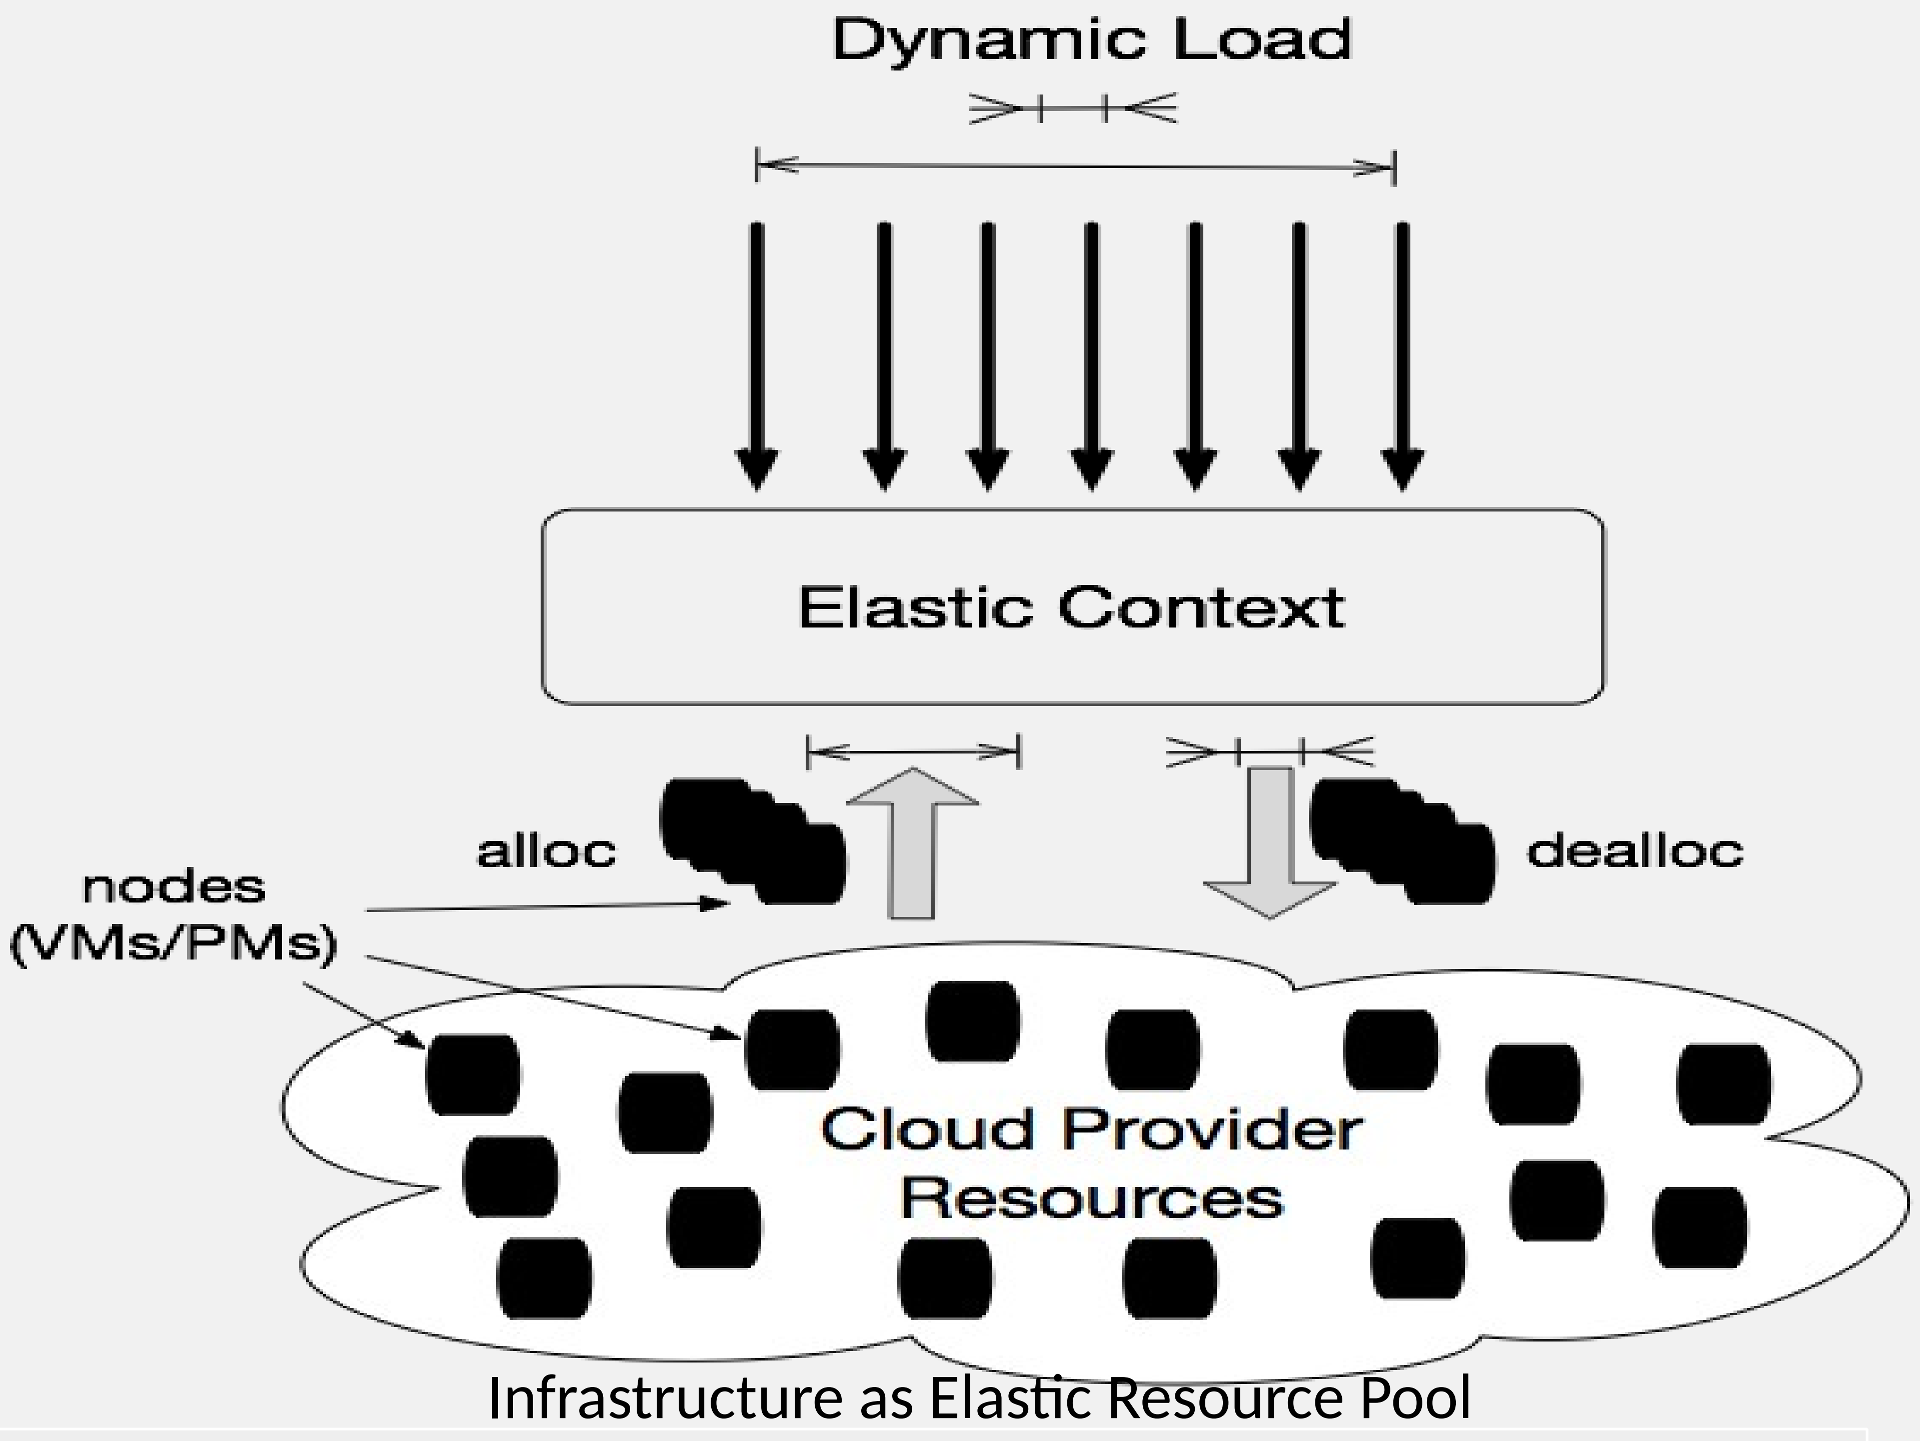

XSP compute service based on Kittyhawk [Appavoo IBM]

- Fast provisioning based on broadcast
- Hardware level based on HaaS
- IaaS level by pre-allocating VMs out of OpenStack

Potential Impacts

Health big-data
Quackenbush, Guttag

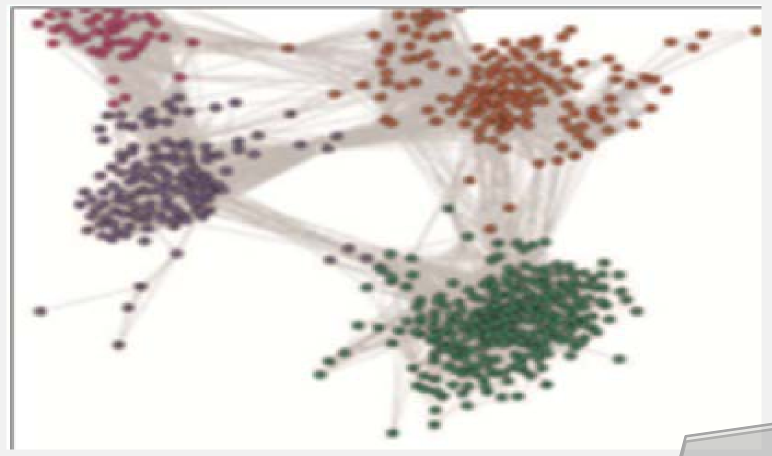

Neuro-science big-data
Lichtman, Boyden

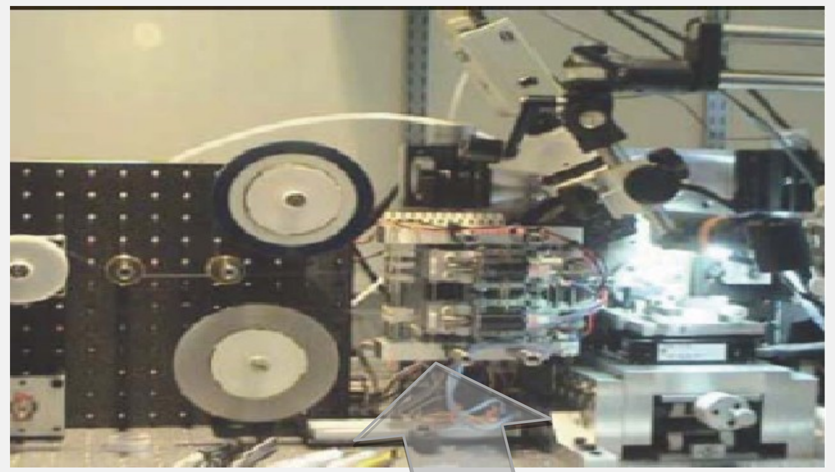

Clinical Imaging
Grant

SuperCloud D4M Analytics
Kepner

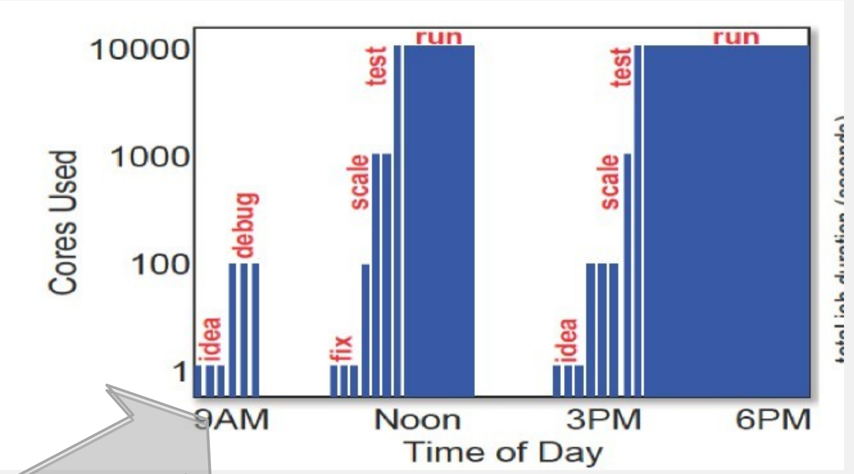

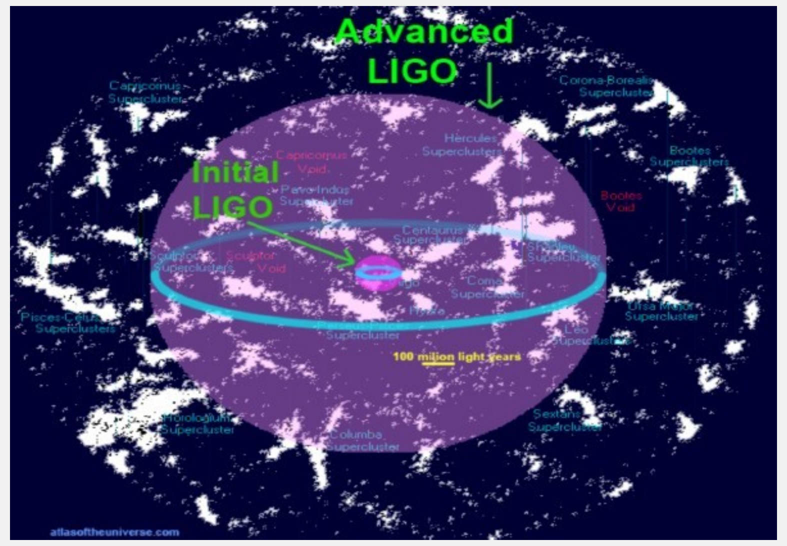

aLigo
Caltech/MIT

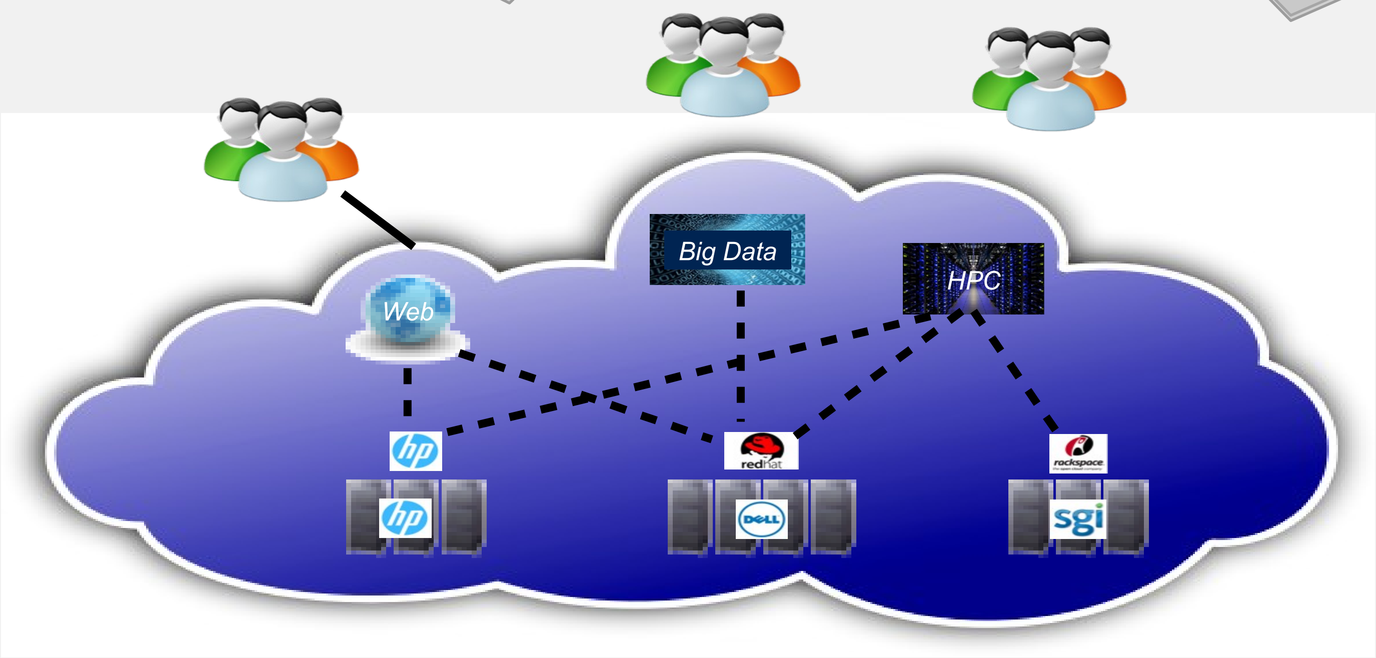

Authentic Model Interaction
Hill, Flierl

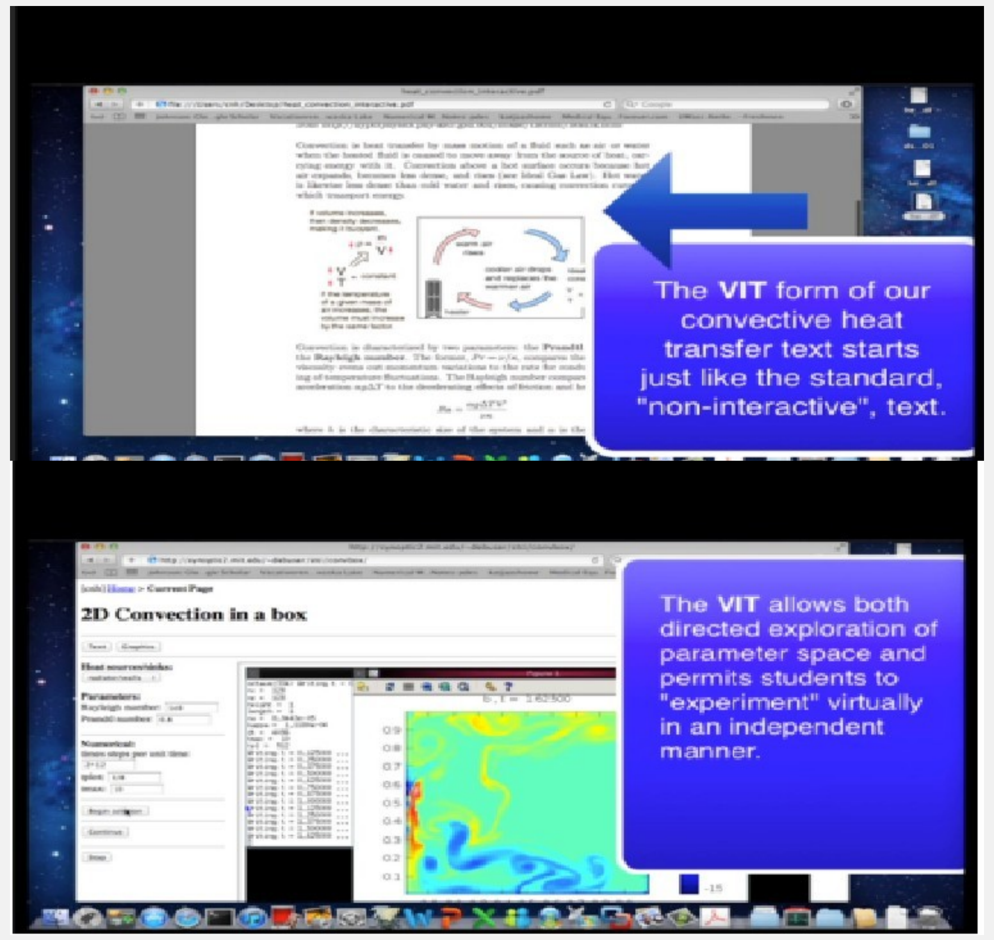

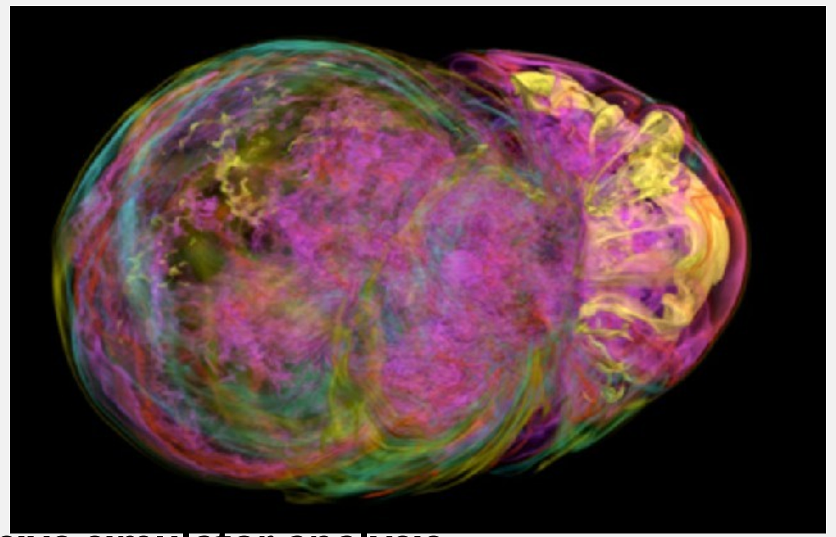

Massive simulator analysis
Mezzacappa

Virtual Materials Design
Many (Ceder, Aspuru-Guzik)

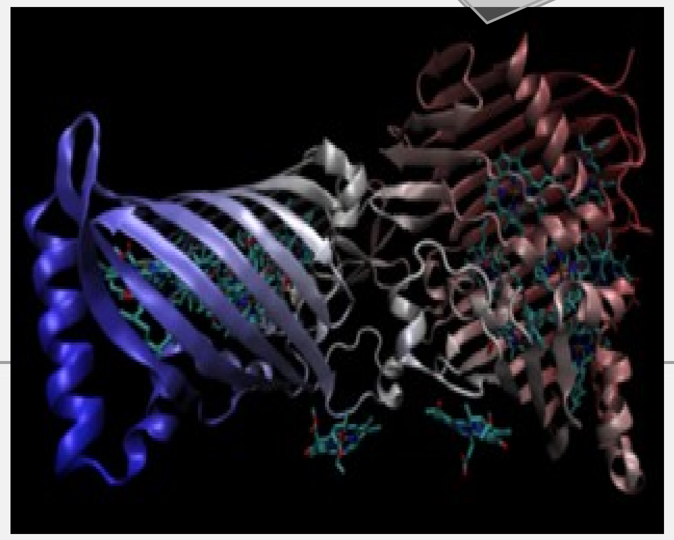

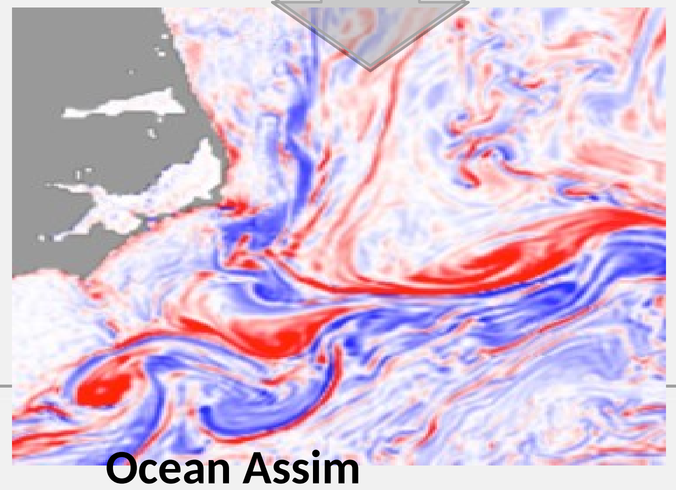

Ocean Assim
Wunsch et al

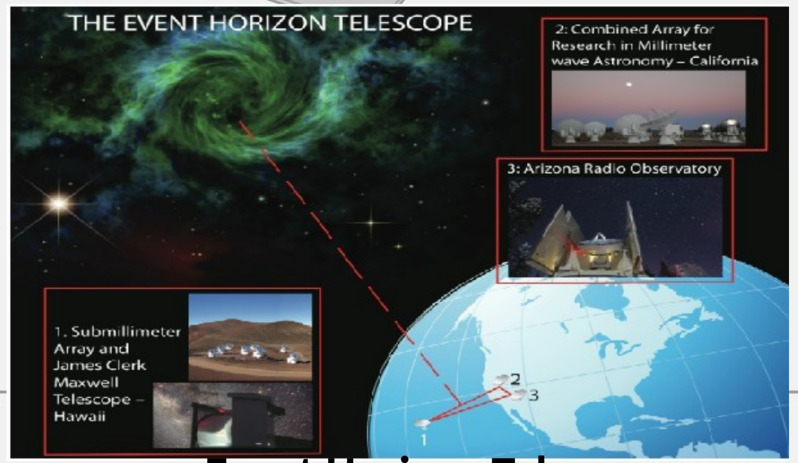

Event Horizon Telescope
Doleman

```
julia> r = @spawn rand(2,2)
RemoteRef(1,1,0)

julia> s = @spawn 1 .+ fetch(r)
RemoteRef(1,1,1)

julia> fetch(s)
1.10824216411304866 1.13798233877923116
1.12376292706355074 1.18750497916607167
```

Interactive Supercomputing Environments
Edelman

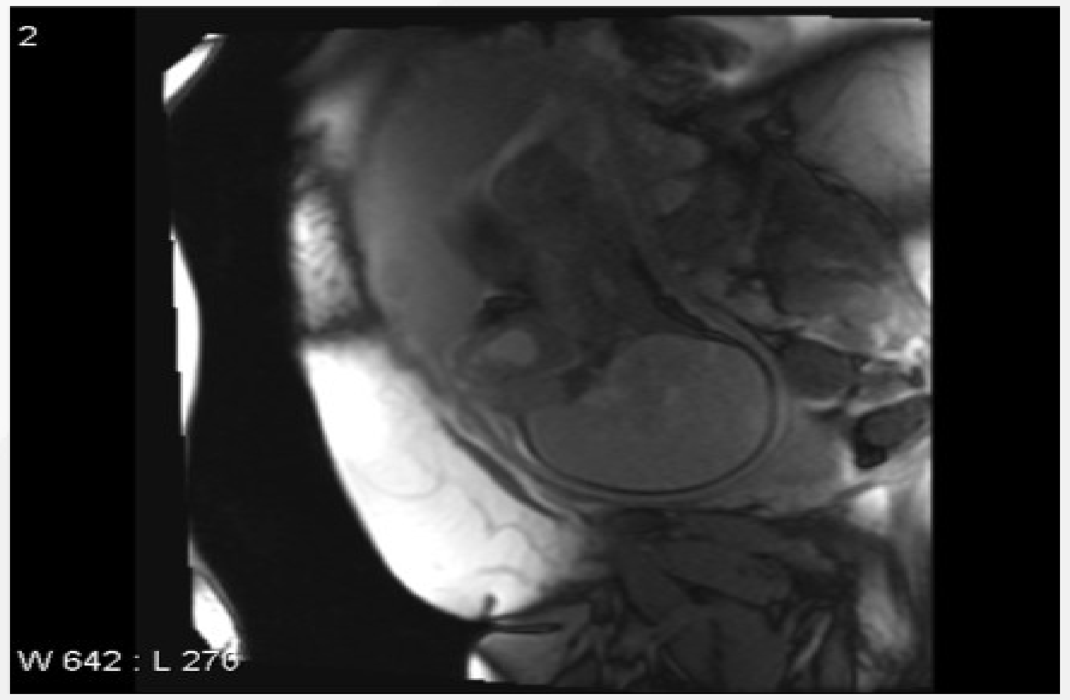

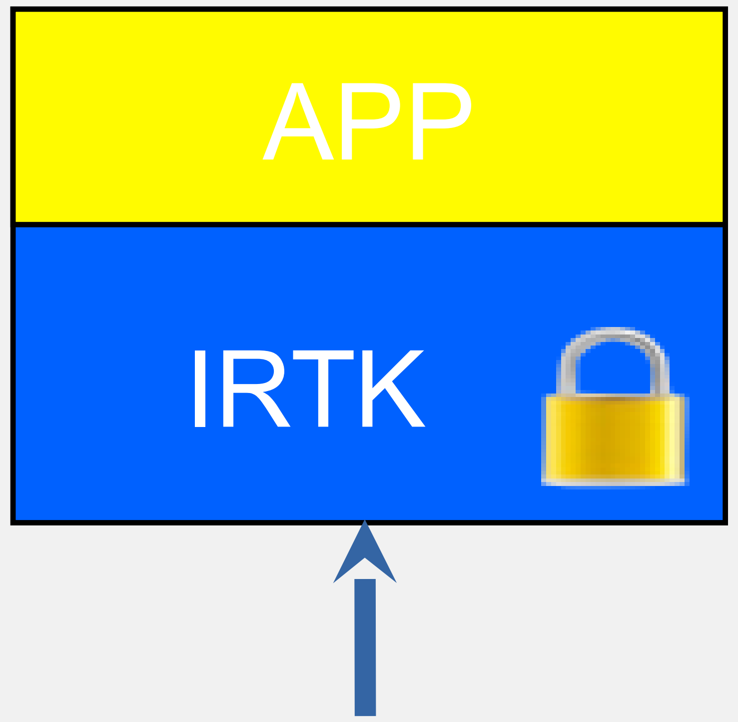

24hrs

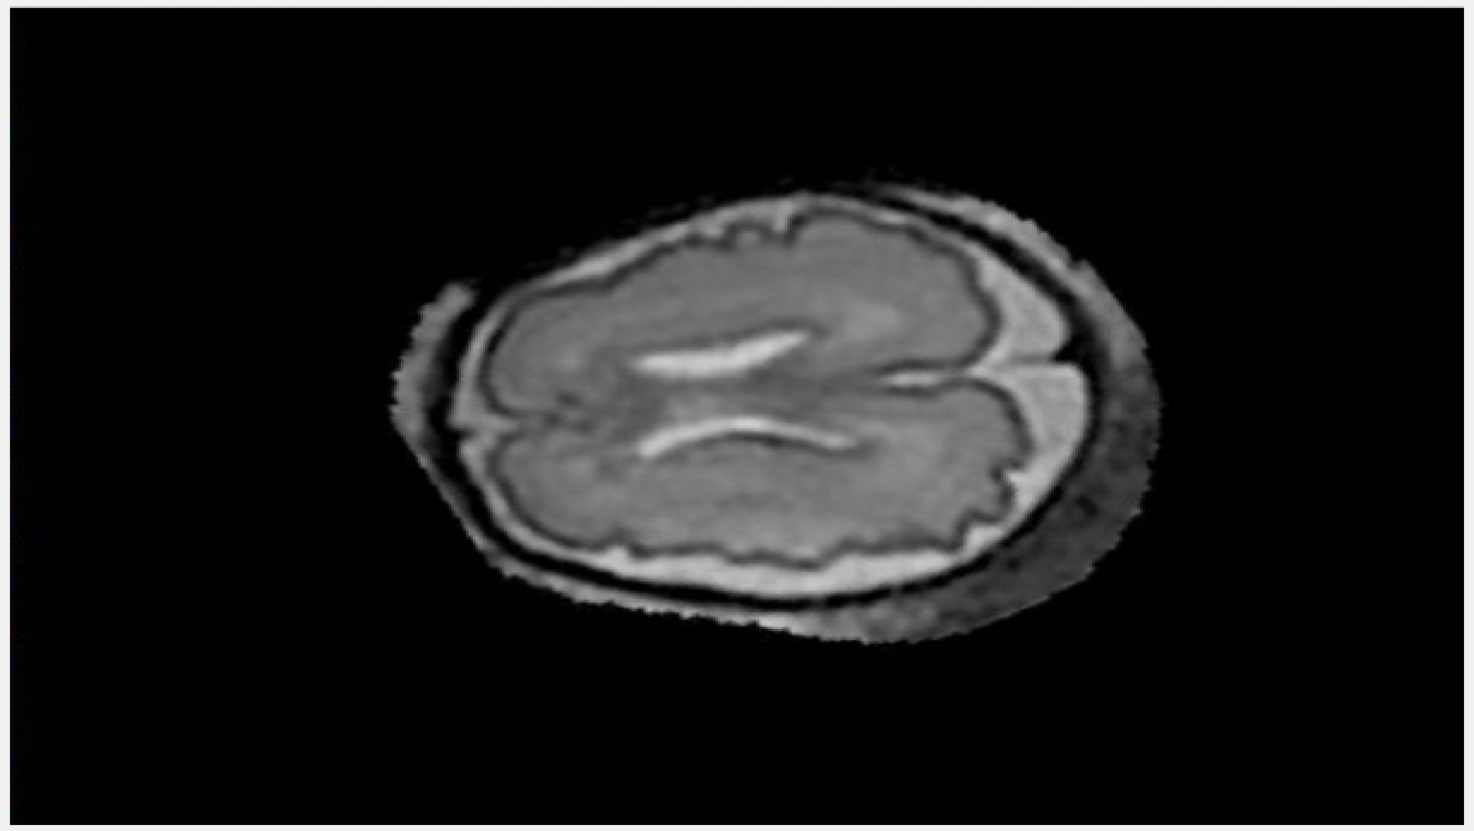

resized,
cropped
96x96
50 slices

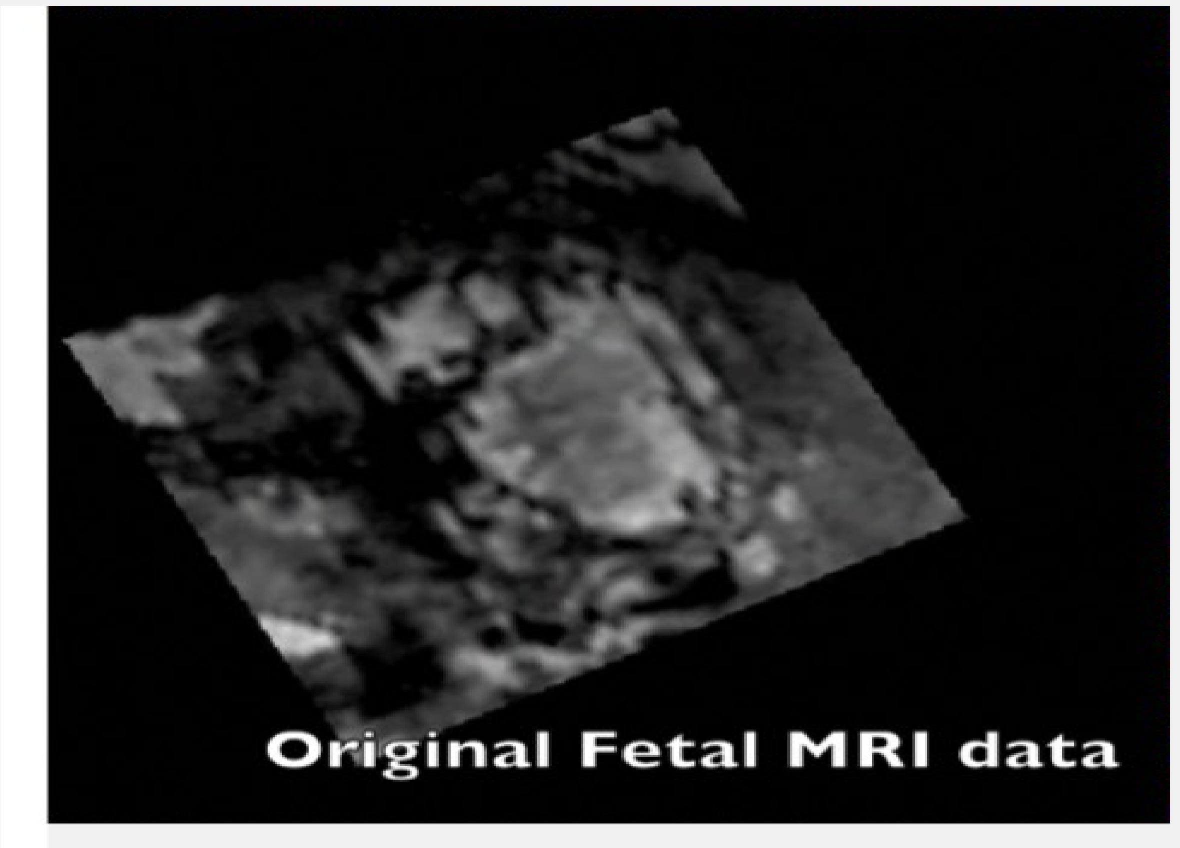

Fetal Image Reconstruction

synthetic
1024x1024
200 slices

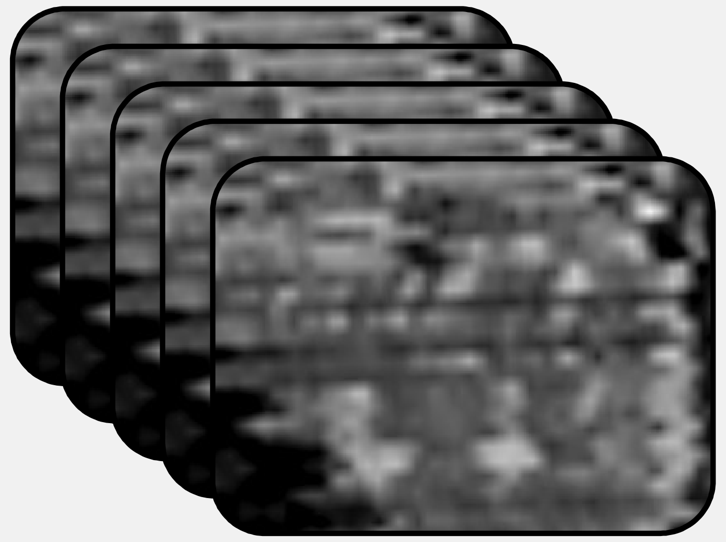

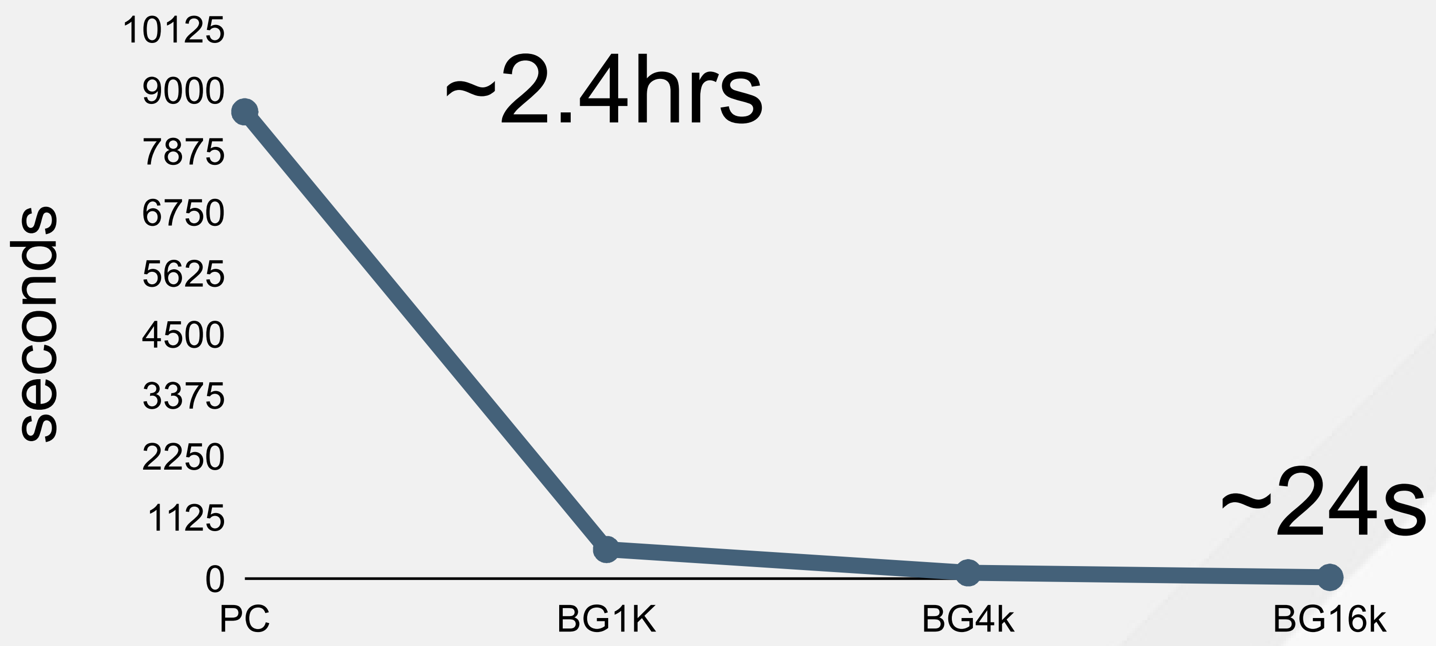

Examples


Spark
hadoop
VERTICA
redhat
Industry
sgi
GlusterFS

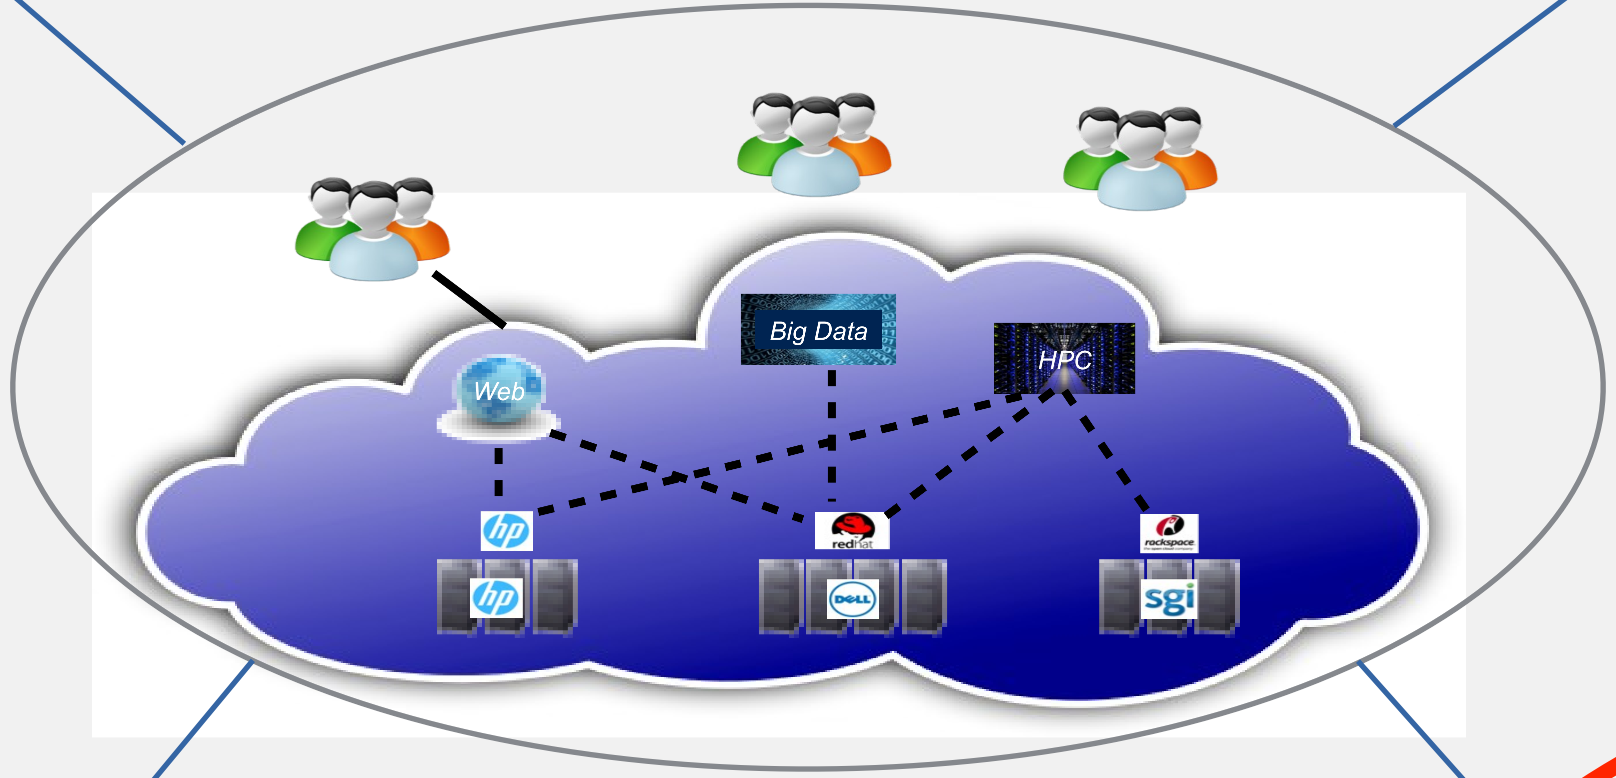
Web
Big Data
HPC
sgi

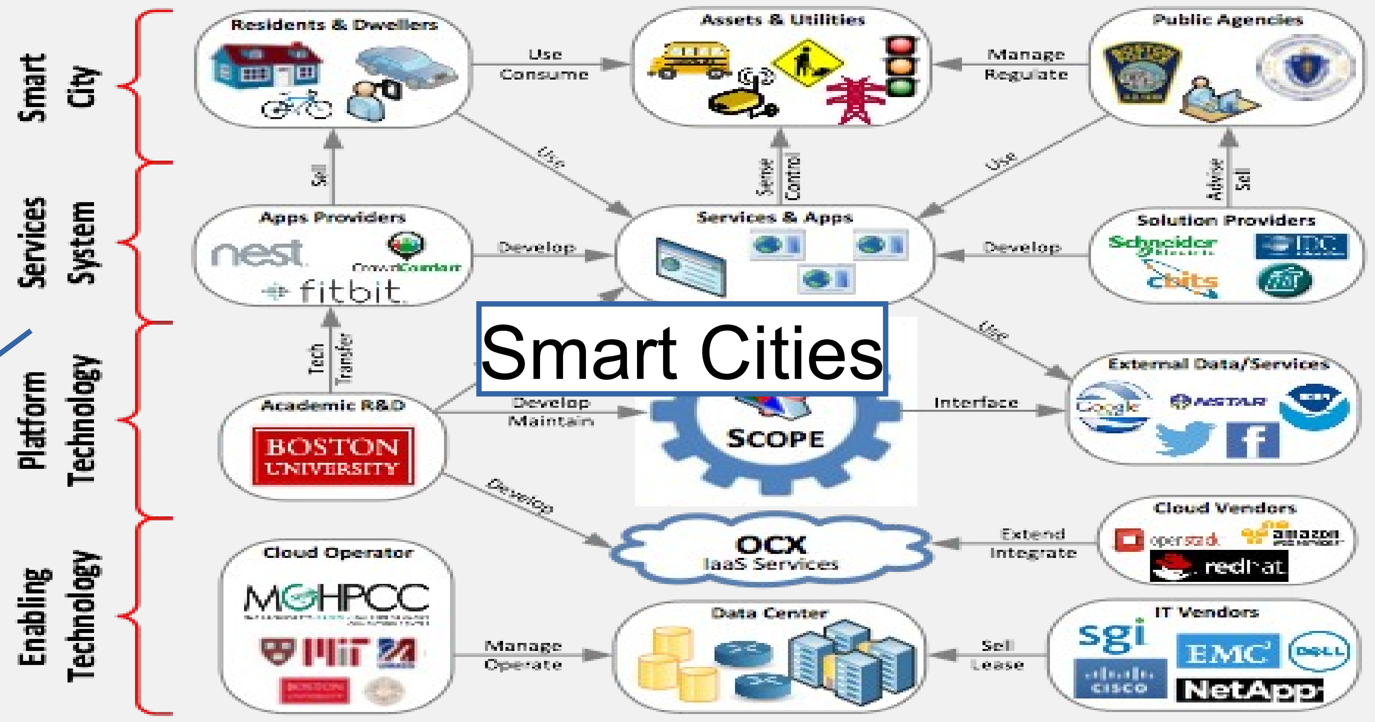
Smart City
Services System
Platform Technology
Enabling Technology
Residents & Dwellers
Assets & Utilities
Public Agencies
Apps Providers
Services & Apps
Solution Providers
Smart Cities
Academic R&D
BOSTON UNIVERSITY
SCOPE
External Data/Services
Cloud Operator
MGHPCC
OCX
IaaS Services
Cloud Vendors
Data Center
IT Vendors
sgi
EMC
NetApp

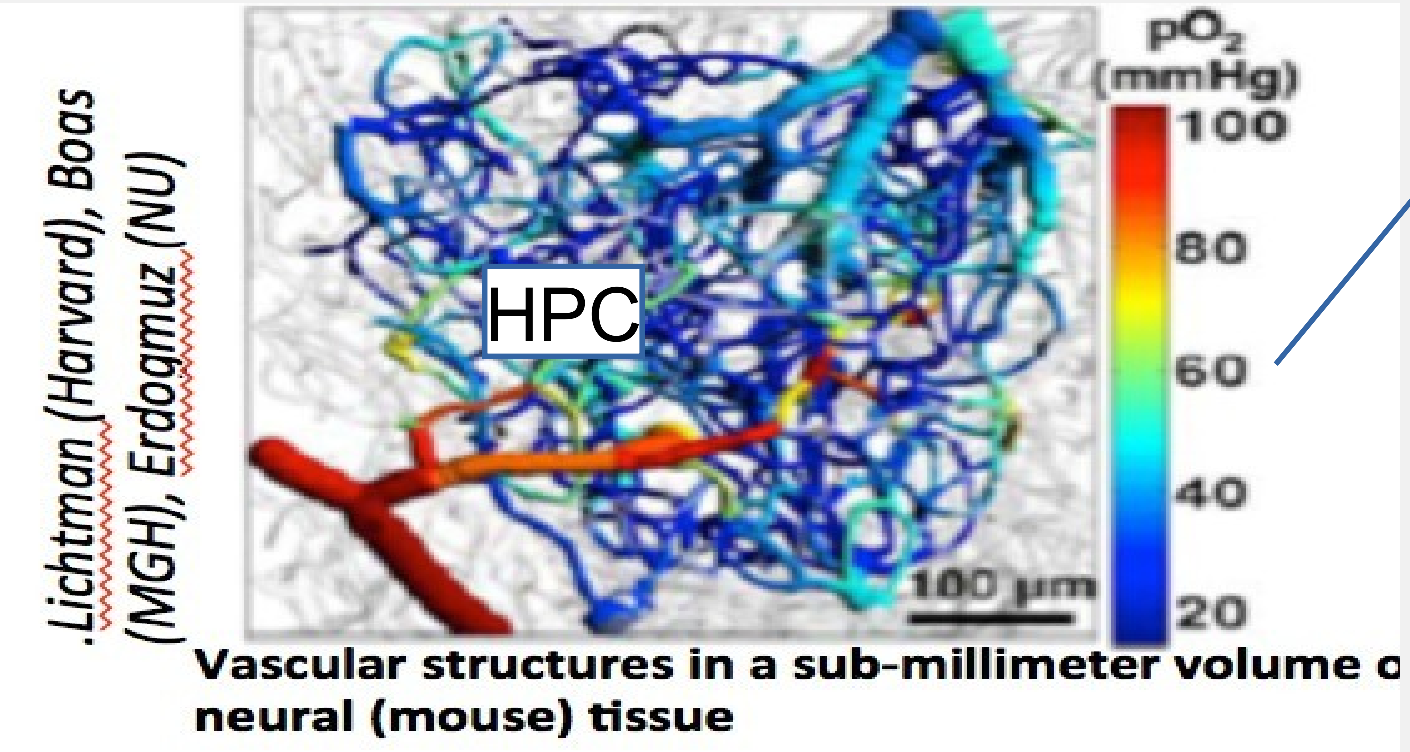
Lichtman (Harvard), Boas (MGH), Erdogmuz (NU)
HPC
pO_2 (mmHg)
100
80
60
40
20
100 µm
Vascular structures in a sub-millimeter volume o
neural (mouse) tissue

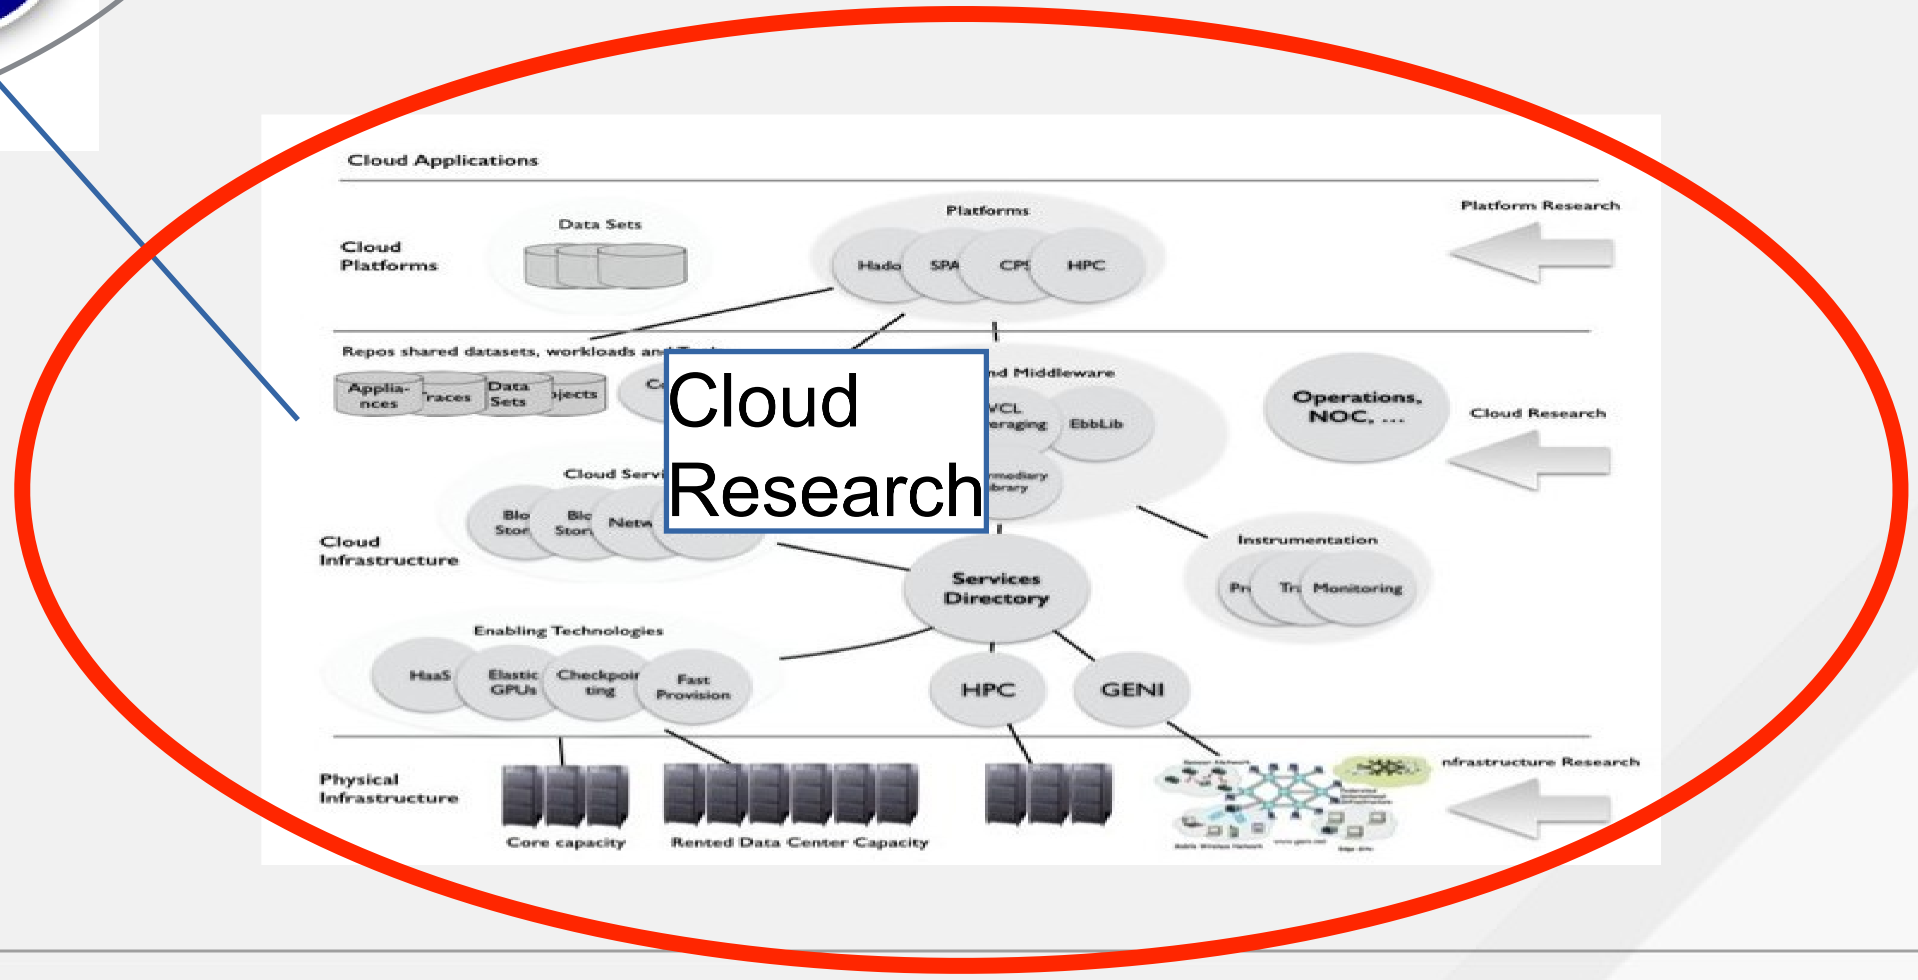
Cloud Applications
Cloud Platforms
Data Sets
Platforms
Platform Research
Cloud Research
Operations, NOC, ...
Cloud Infrastructure
Instrumentation
Services Directory
Enabling Technologies
HPC
GENI
Physical Infrastructure
Core capacity
Rented Data Center Capacity

Example: Modular Approach to Cloud Security $10M NSF Frontier Grant:

http://www.bu.edu/hic/research/macs/

Cloud computing

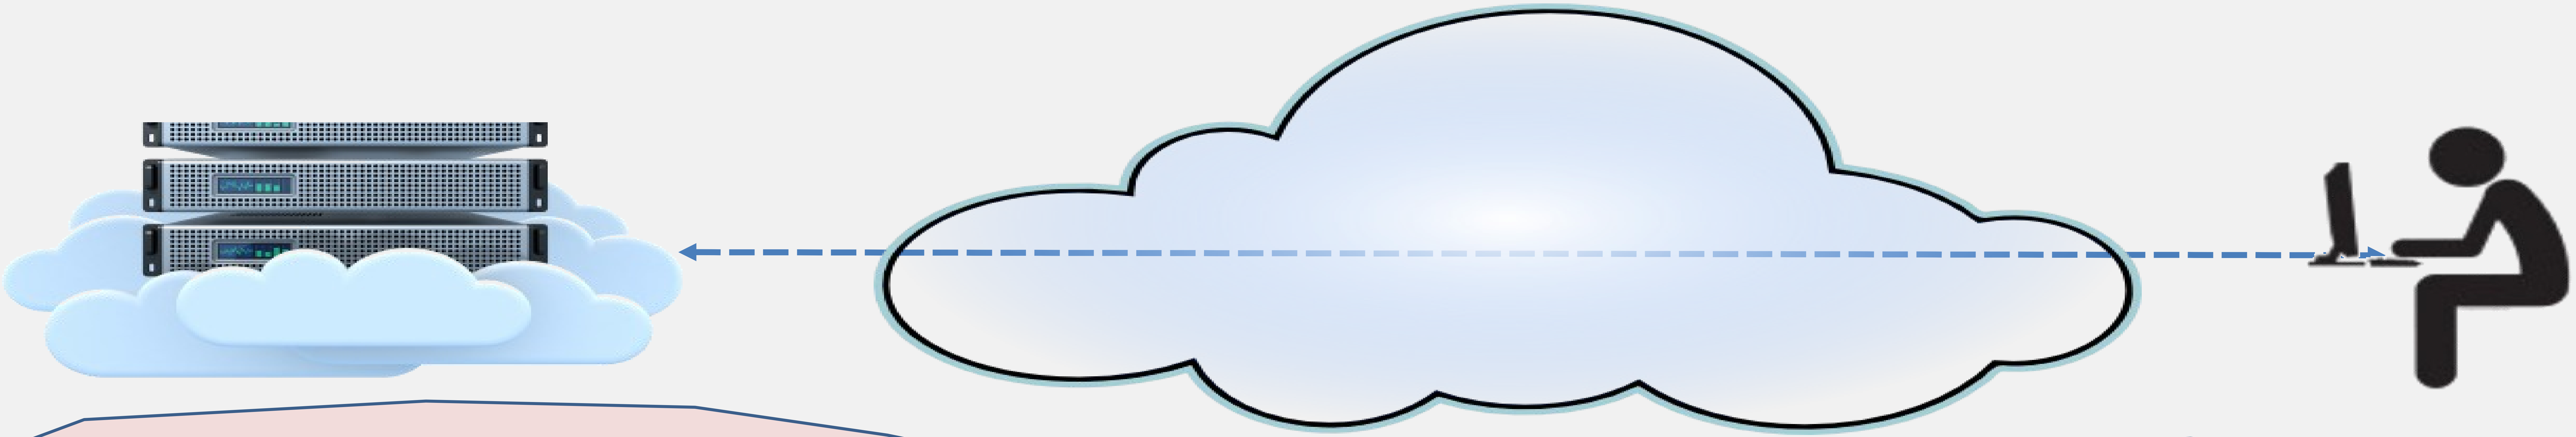

Cloud computing

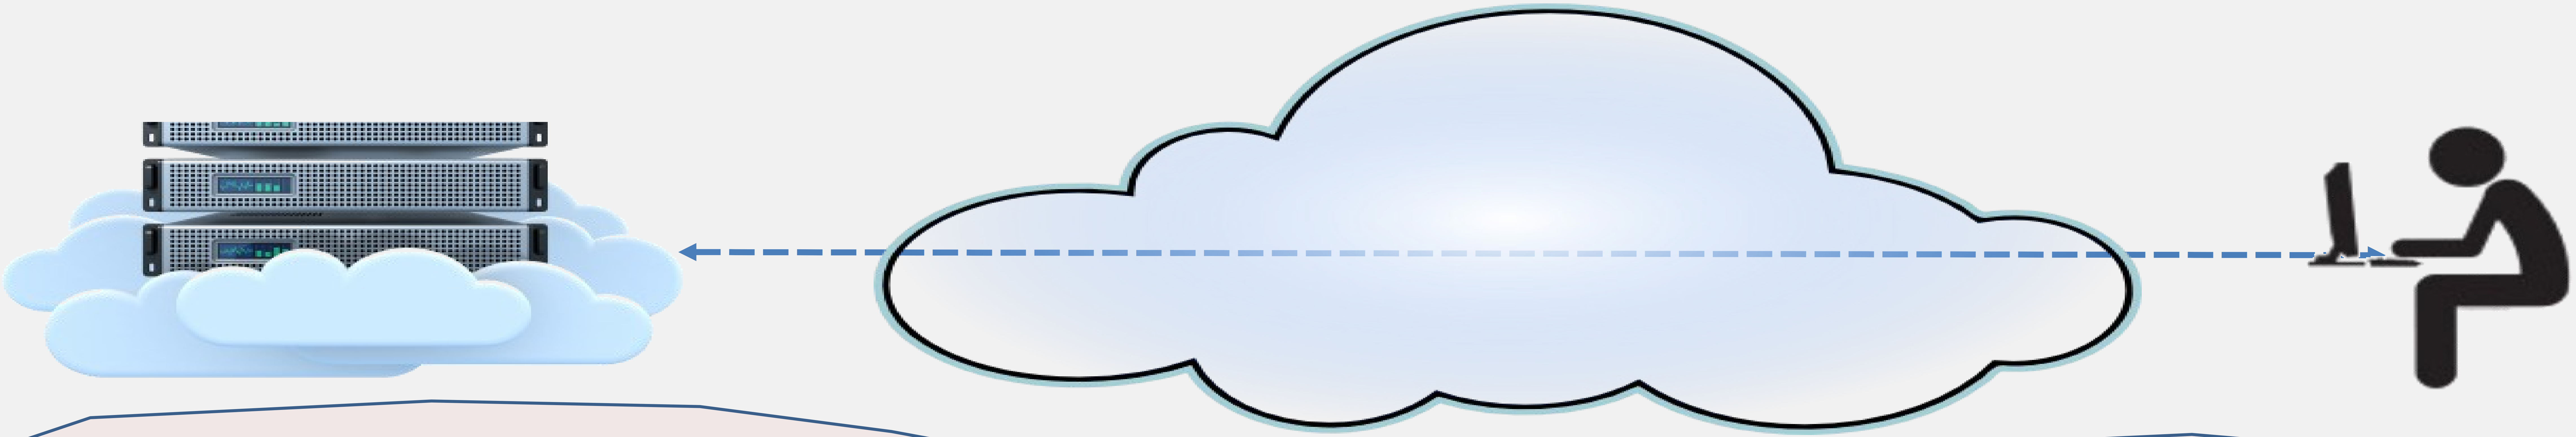

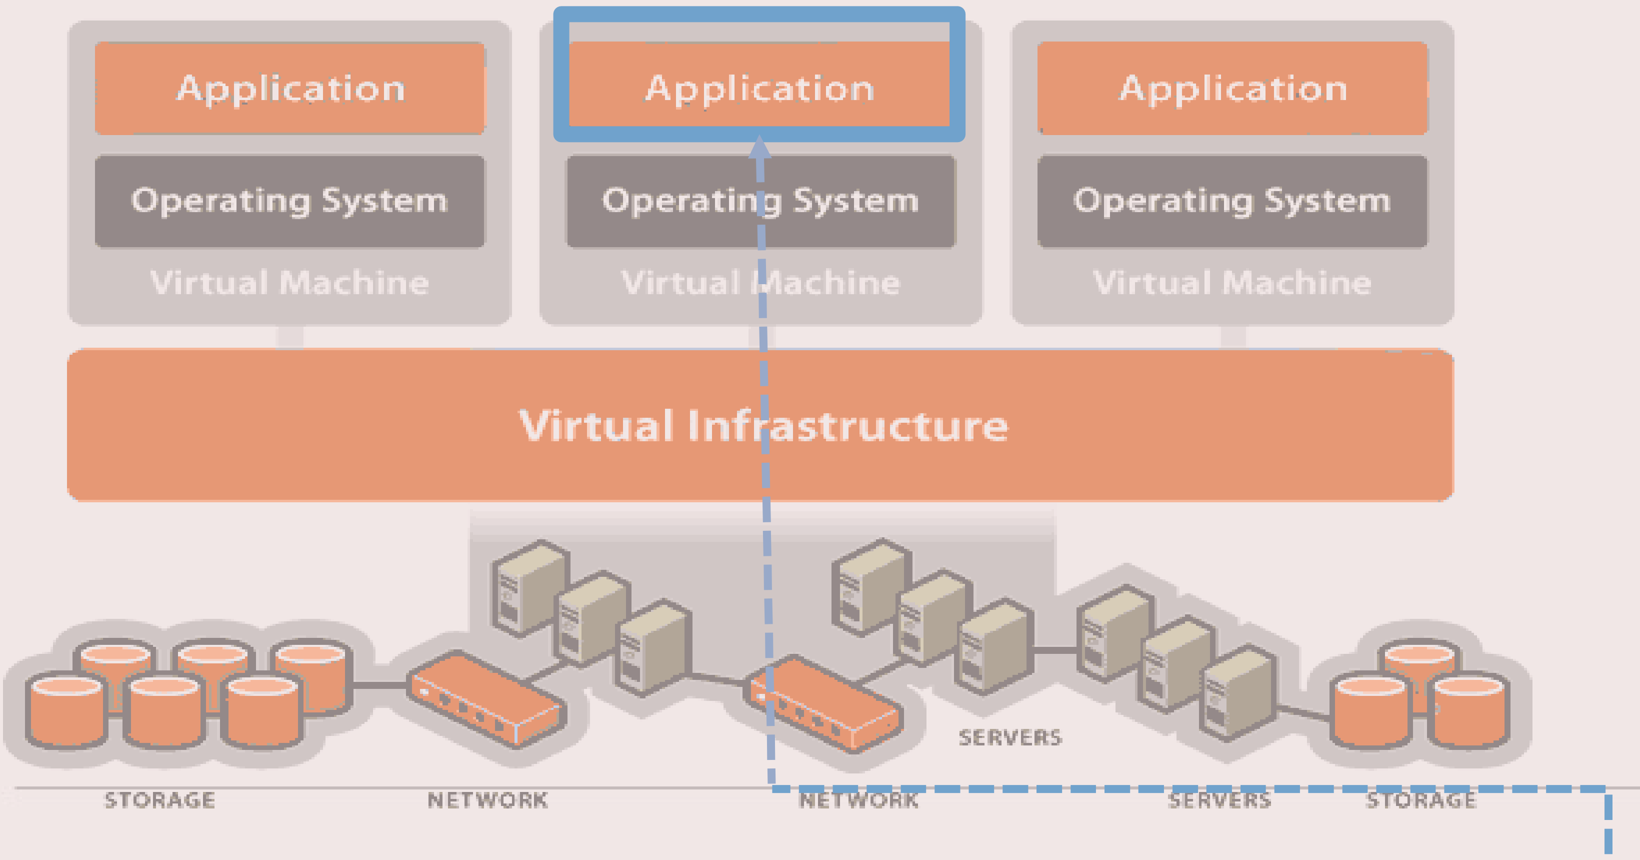

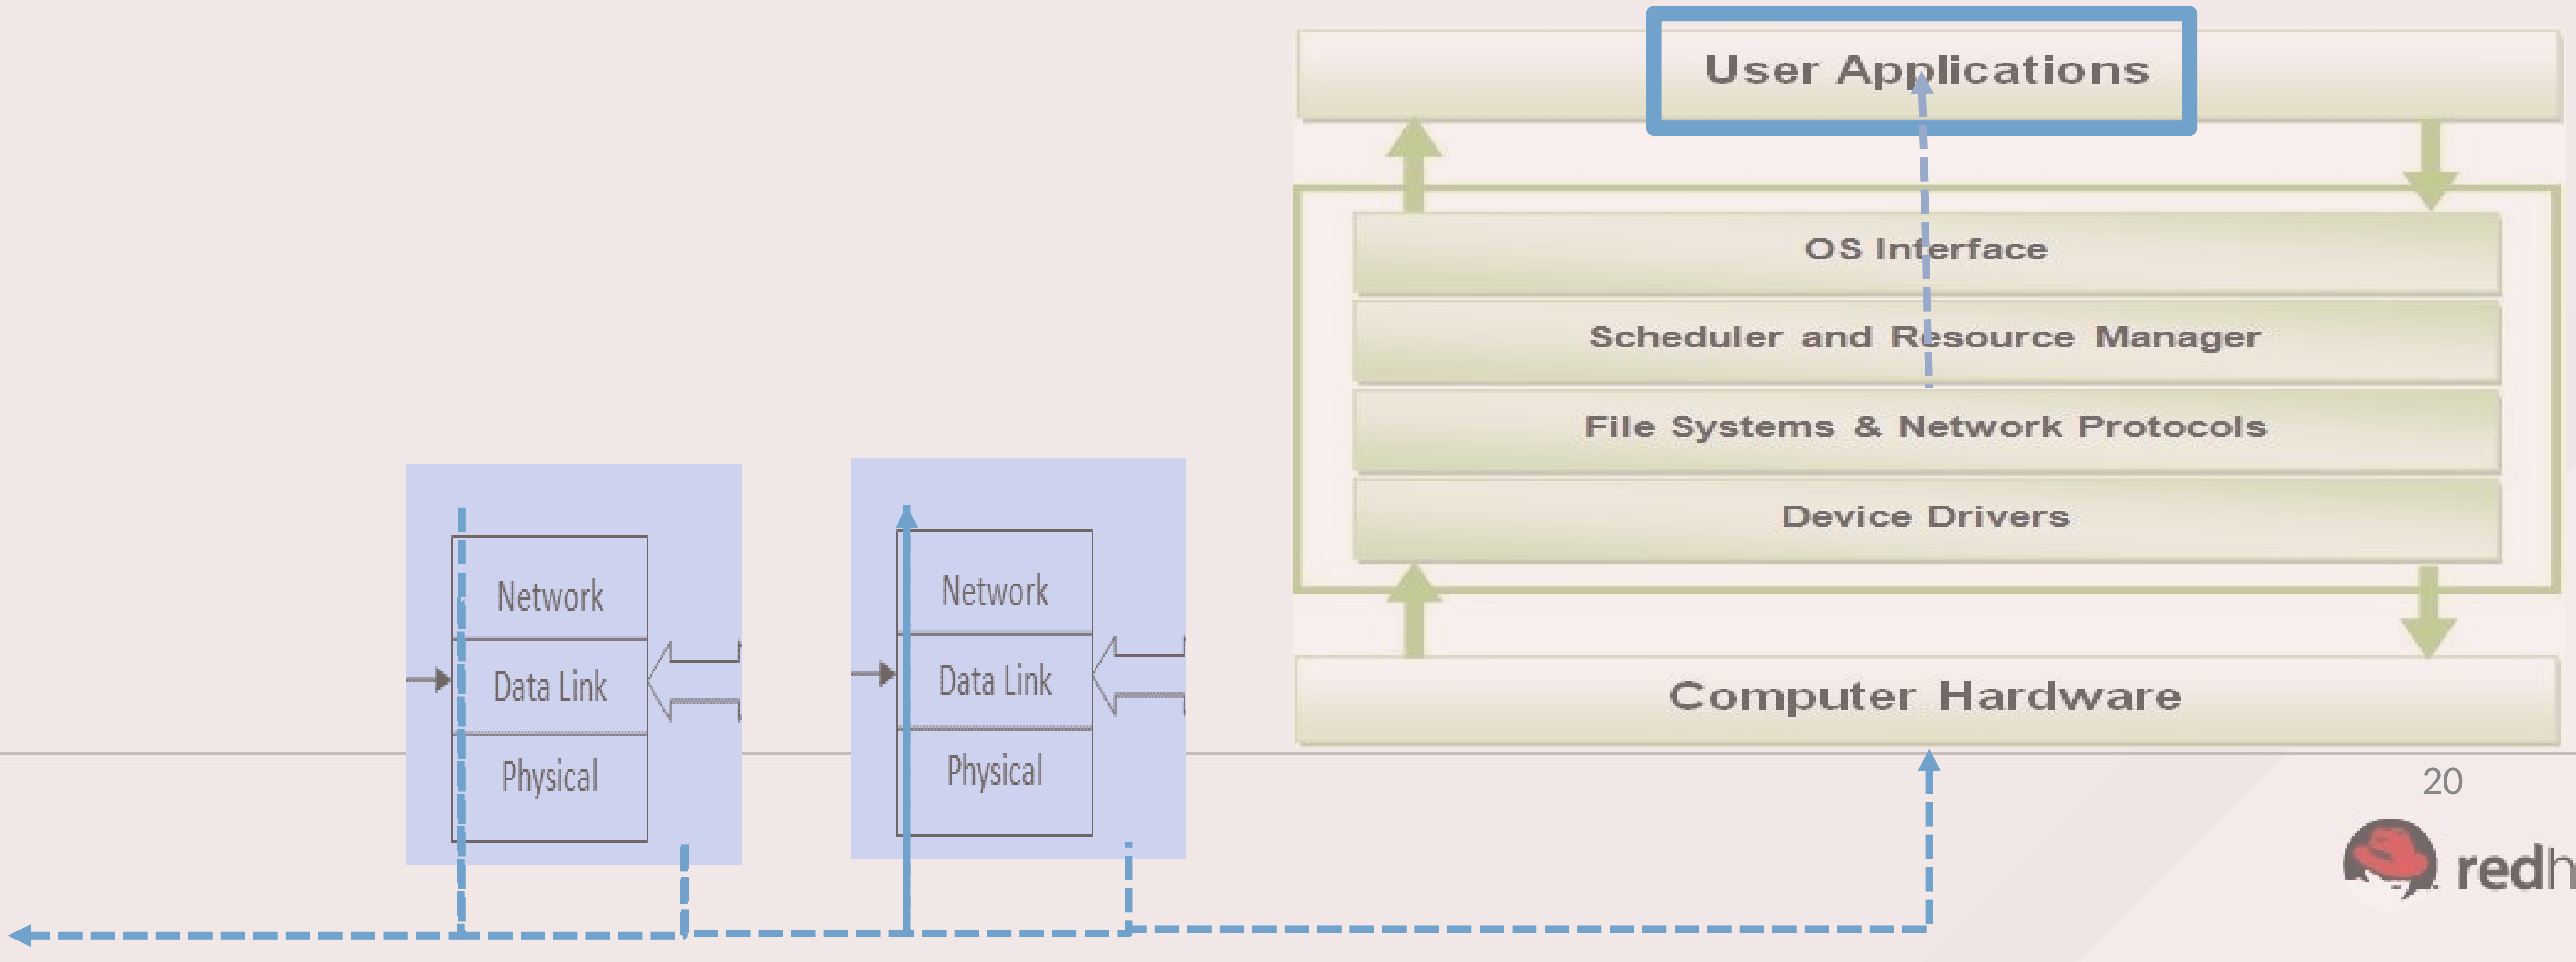

Cloud computing

Cloud computing

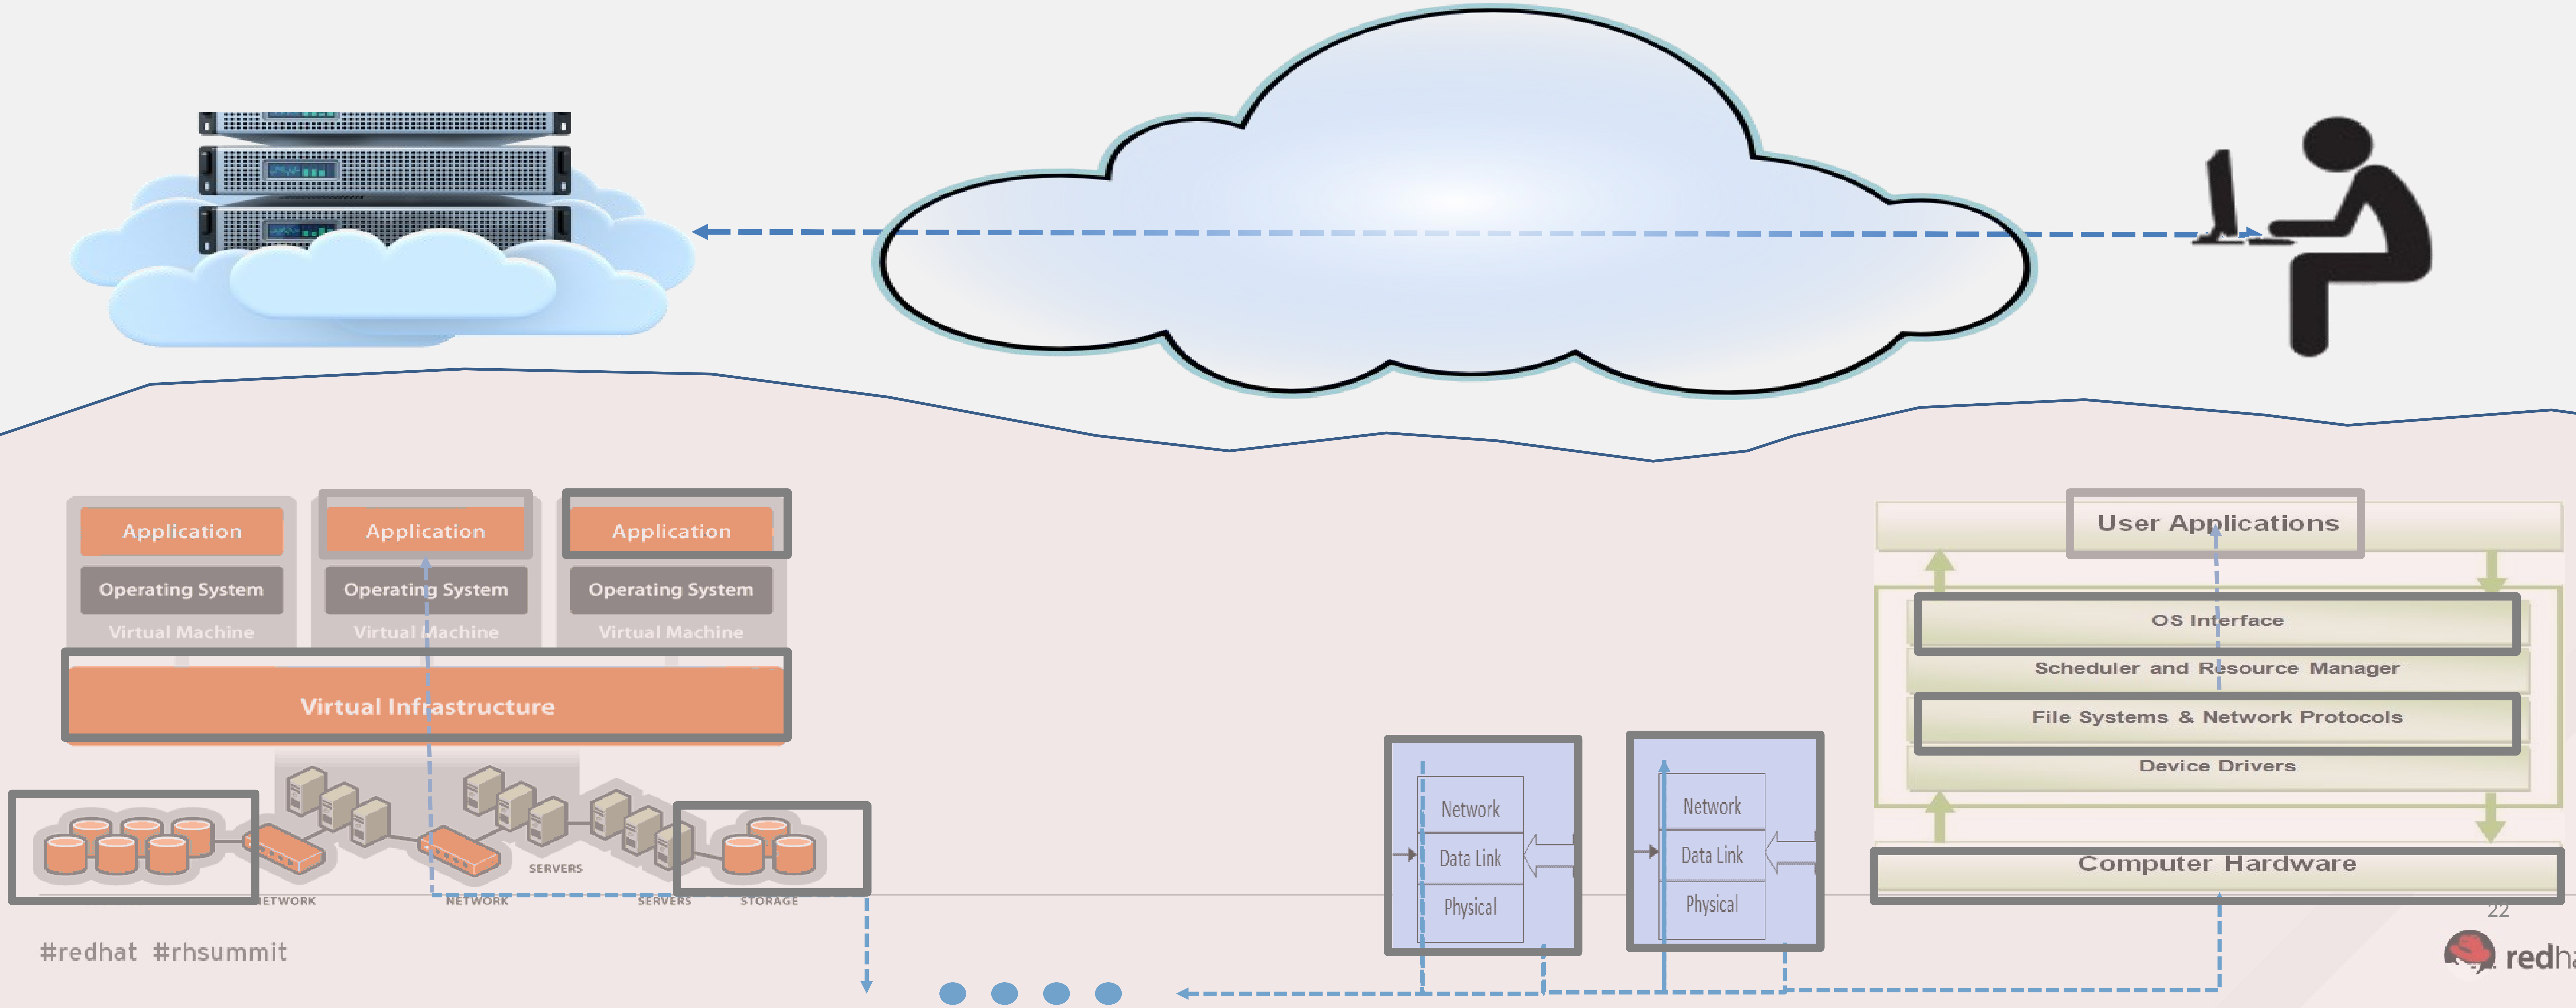

In security, the sum of the parts is often a *hole*.

Dave Safford, circa 2000

Key directions

- Modular security: extend model that has been successful in crypto community in a small scale
- Define architecture where security SLA can be expressed and realized by composition of different services, auditing
- Better system, network, and hardware design (CryptDB/Monomi,EbbRT, XIA, SDN, ASCEND…)
- Multiple solutions for each problem, e.g., oblivious computing:
  - secure HW & OS
  - homomorphic encryption
  - distributed secure computation
- Put together in common testbed, compare security, compare performance, offer to real users

Other Examples

- Marketplace model
- Cyberphysical systems:
- Energy optimization/matching power company signal
- Provenance aware storage
- Software defined networks
- New internet protocols, e.g., XIA
- FAWN like architectures
- New storage architectures
- Highly elastic software
- Intermediaries for Big Data, HPC, PaaS, SaaS
- Federation services

THE PROJECT

The Massachusetts Open Cloud

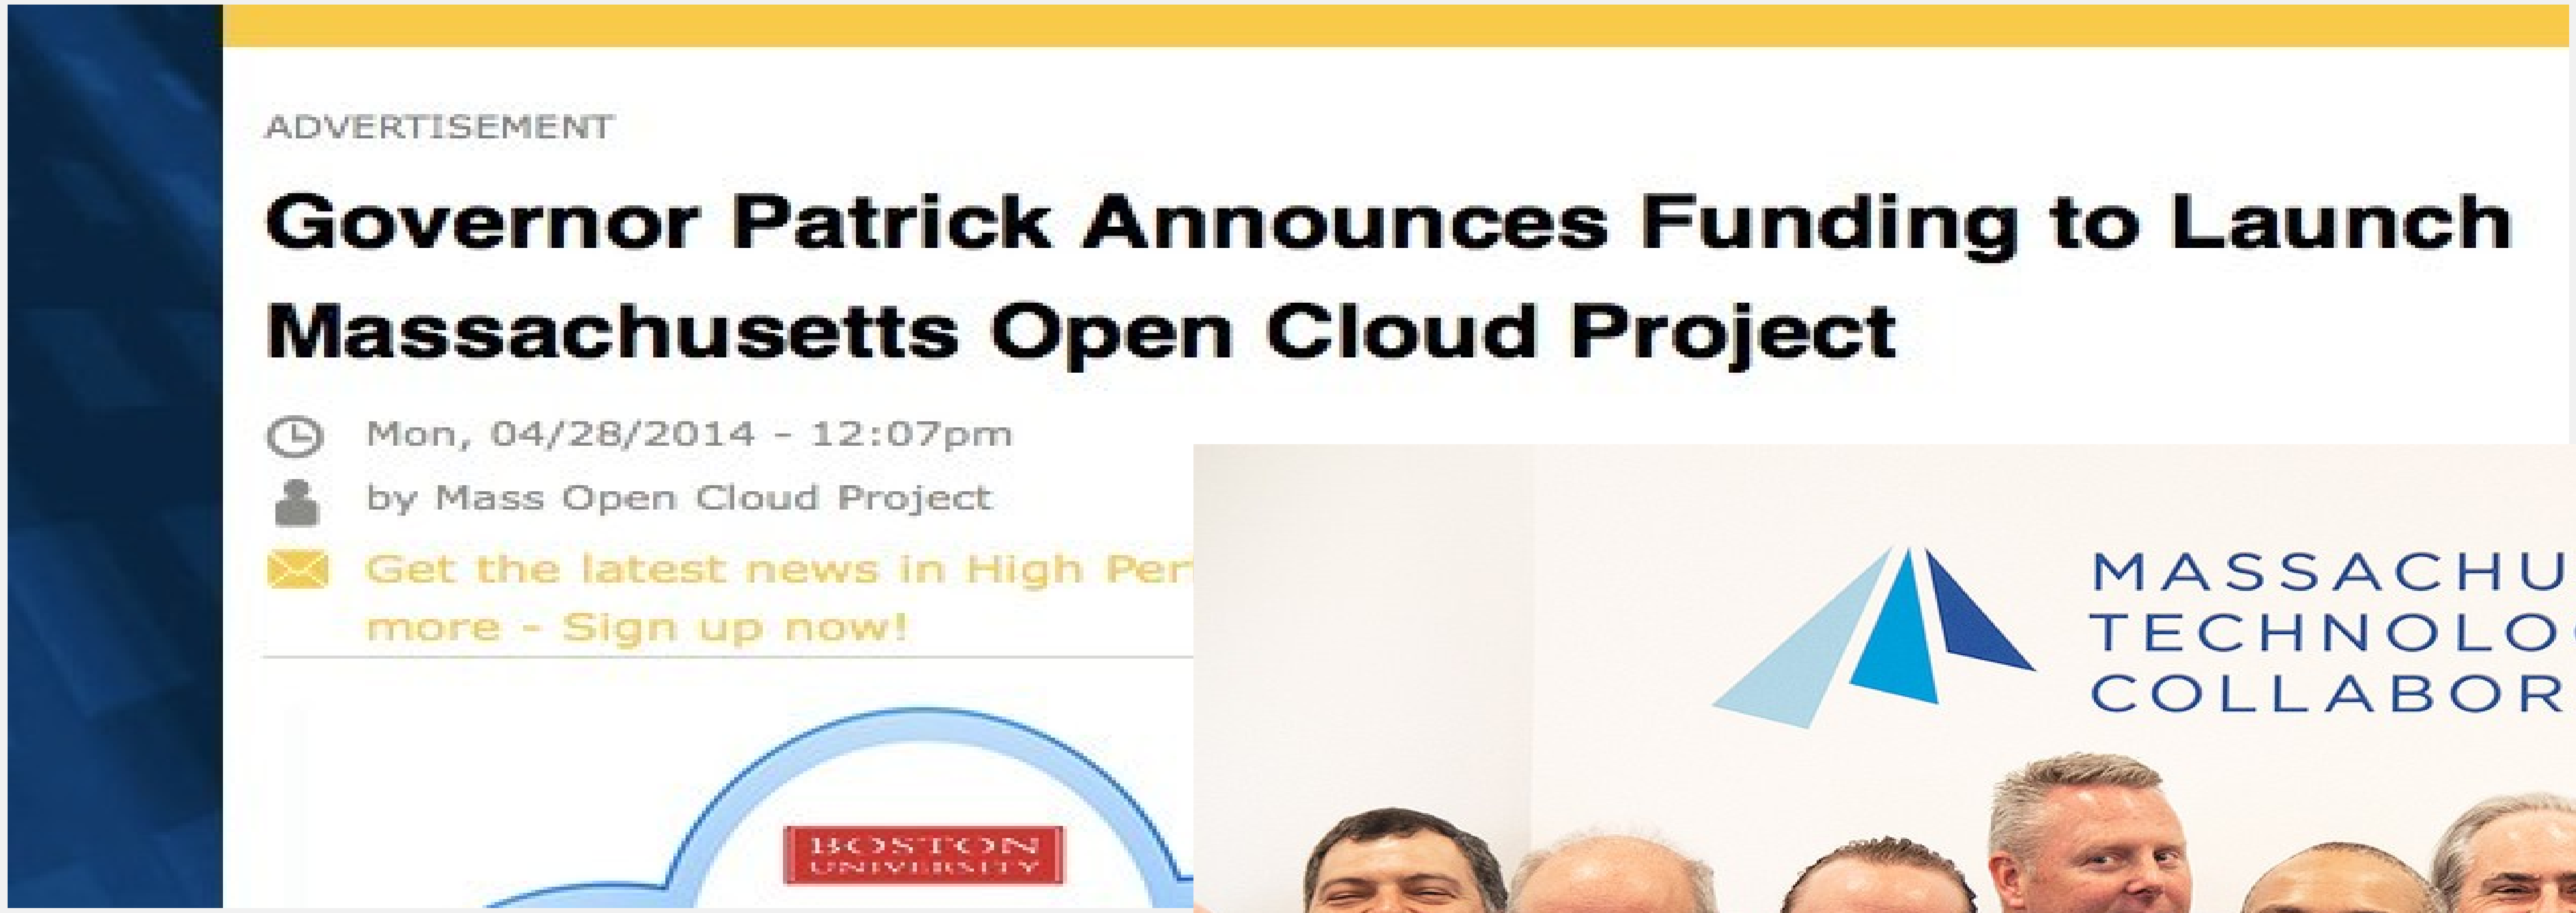
ADVERTISEMENT

Governor Patrick Announces Funding to Launch Massachusetts Open Cloud Project

Mon, 04/28/2014 - 12:07pm

by Mass Open Cloud Project

Get the latest news in High Per
more - Sign up now!

Facilities

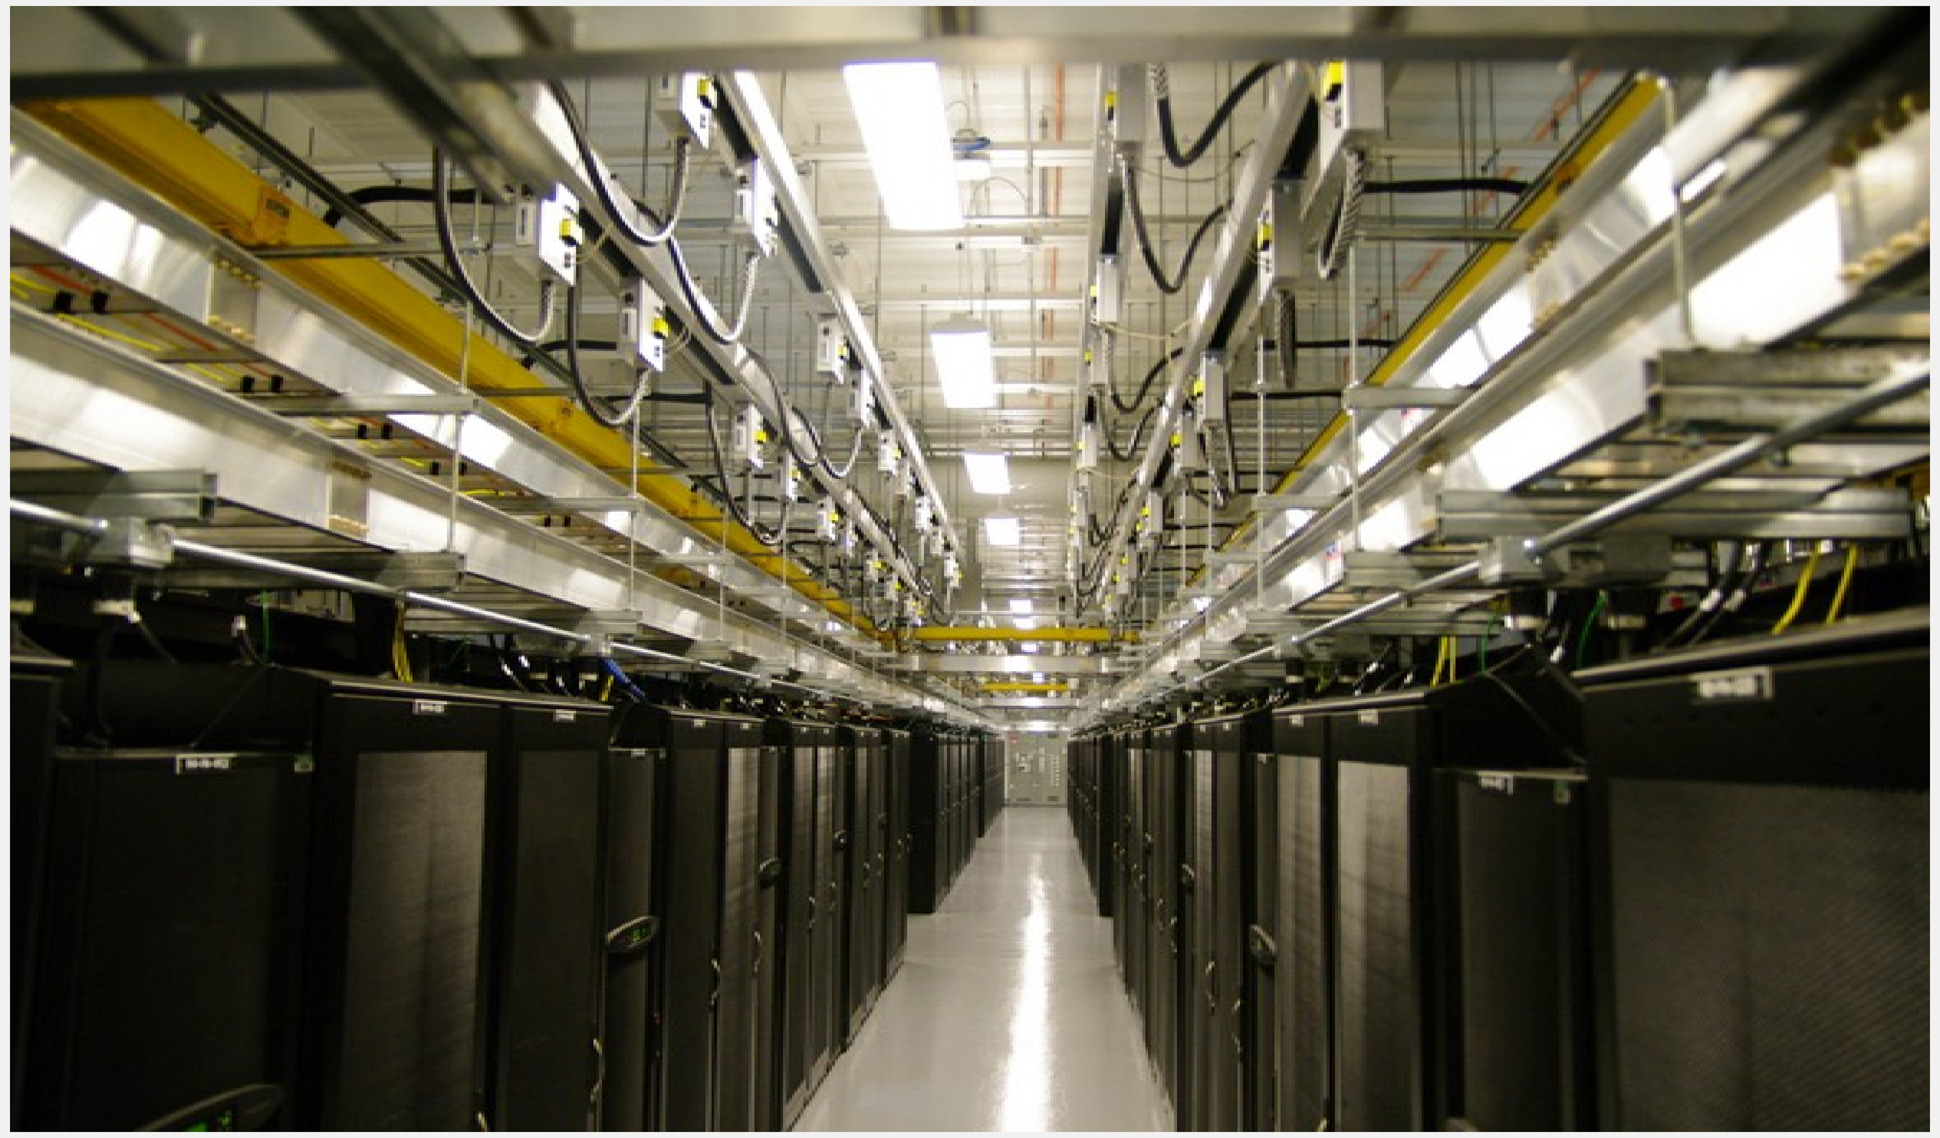

- MGHPCC: 15 MW, 90,000 square feet + can grow
- “Mass. Green High Performance Computing Center”
- Data center in Holyoke, Mass.
- Jointly run by BU, Northeastern, MIT, Harvard, UMass

MOC Ecosystem

Users/applications

Big Data, HPC, Life Sciences…

Education and Workforce

Students, industry

Data

Commonwealth, BU, HU, NU, MIT, UMass, MGHPCC, Foundations…

Industry

Red Hat, Cisco, Dell, Intel, Brocade, Lenovo, Mellanox, OpenStack, Linux, OpenDaylight

University Research IT

BU, HU, NU, UMass, MIT

Cloud Research

Operating Systems, Power, Security, Marketplace…

Core Dev/Ops

Multi-landlord support, HaaS, Billing, Intermediaries…

Industry partners

- Core Industry partners:
  - Intel: security, HPC/Cloud convergence
  - Lenovo: future technology, cloud, ecosystem
  - Cisco: Networking, OpenStack, Intercloud
  - Red Hat: future of cloud platform - OCX model
  - Brocade: OpenDaylight, networking challenges
- Industry Contributors: Fujitsu, Dell, Mellanox, Riverbed, MathWorks, SGI, DDN, EMC, Plexxi, Juniper, Cambridge Computing, Two Sigma

Project plan

Phase 1 (year 1):
Proof of concept

Phase 2 (year 2/3):
Technology partner ramp up

Phase 3 (year 3+):
Post project evolution and achieving self-sufficiency

	Name
1	OpenStack
2	Integration and feedback
3	phase 1:
4	Early users access (G1)
5	Automation, monitoring, metering
6	Services Directory & intermediaries
7	Big data and IaaS Intermediaries (G1, G2)
8	Evaluate 3rd party services
9	Initial dedicated capacity (G2)
10	Stream public data sets to MOC
11	Host data sets (G1)
12	Deploy/test first production offering
13	Trusted administrator HaaS
14	Production MOC (G1)
15	phase 2:
16	Multi-provider
17	SDL & Intermediary library
18	Compatibility tests
19	Multi-provider MOC (G2)
20	Enable differentiated services
21	Enable differentiated intermediaries
22	Partner controlled SPL
23	Highly differentiated services (G2)
24	Productize experimental User Services UI
25	General availability (G1)
26	Integrate new services
27	Self sufficiency (G1,G2)
28	Business model defniition

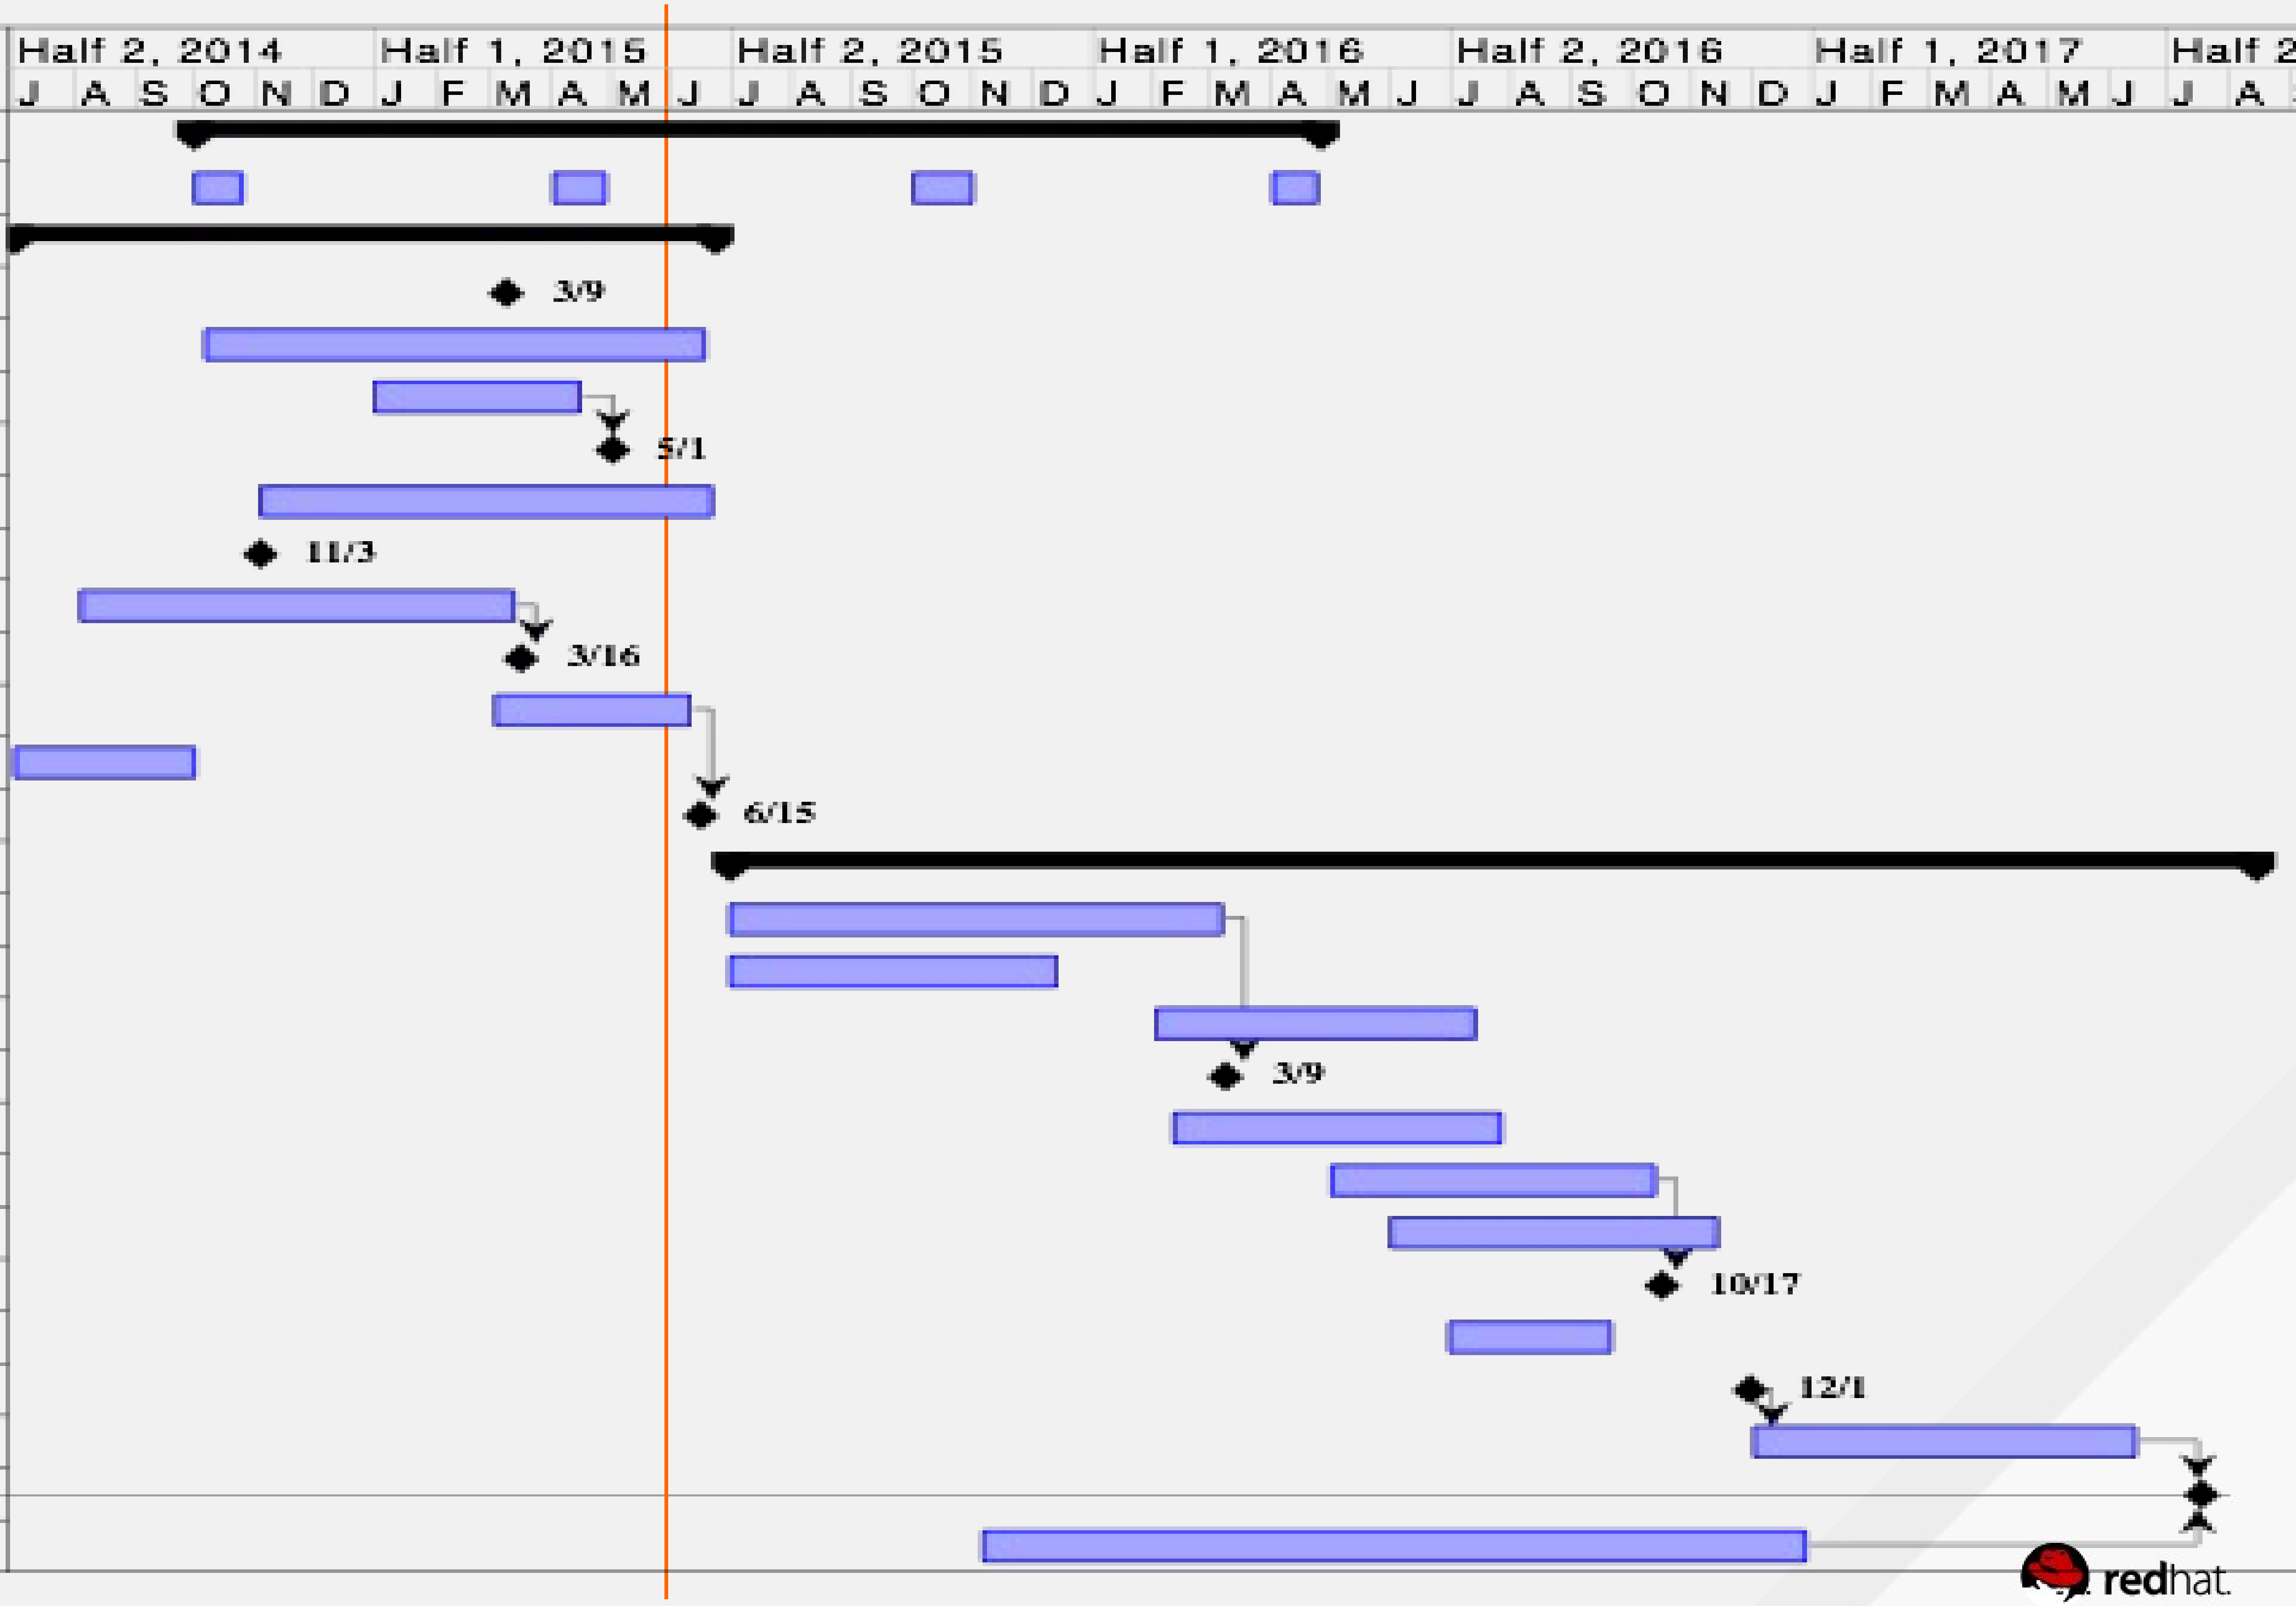

Example use cases

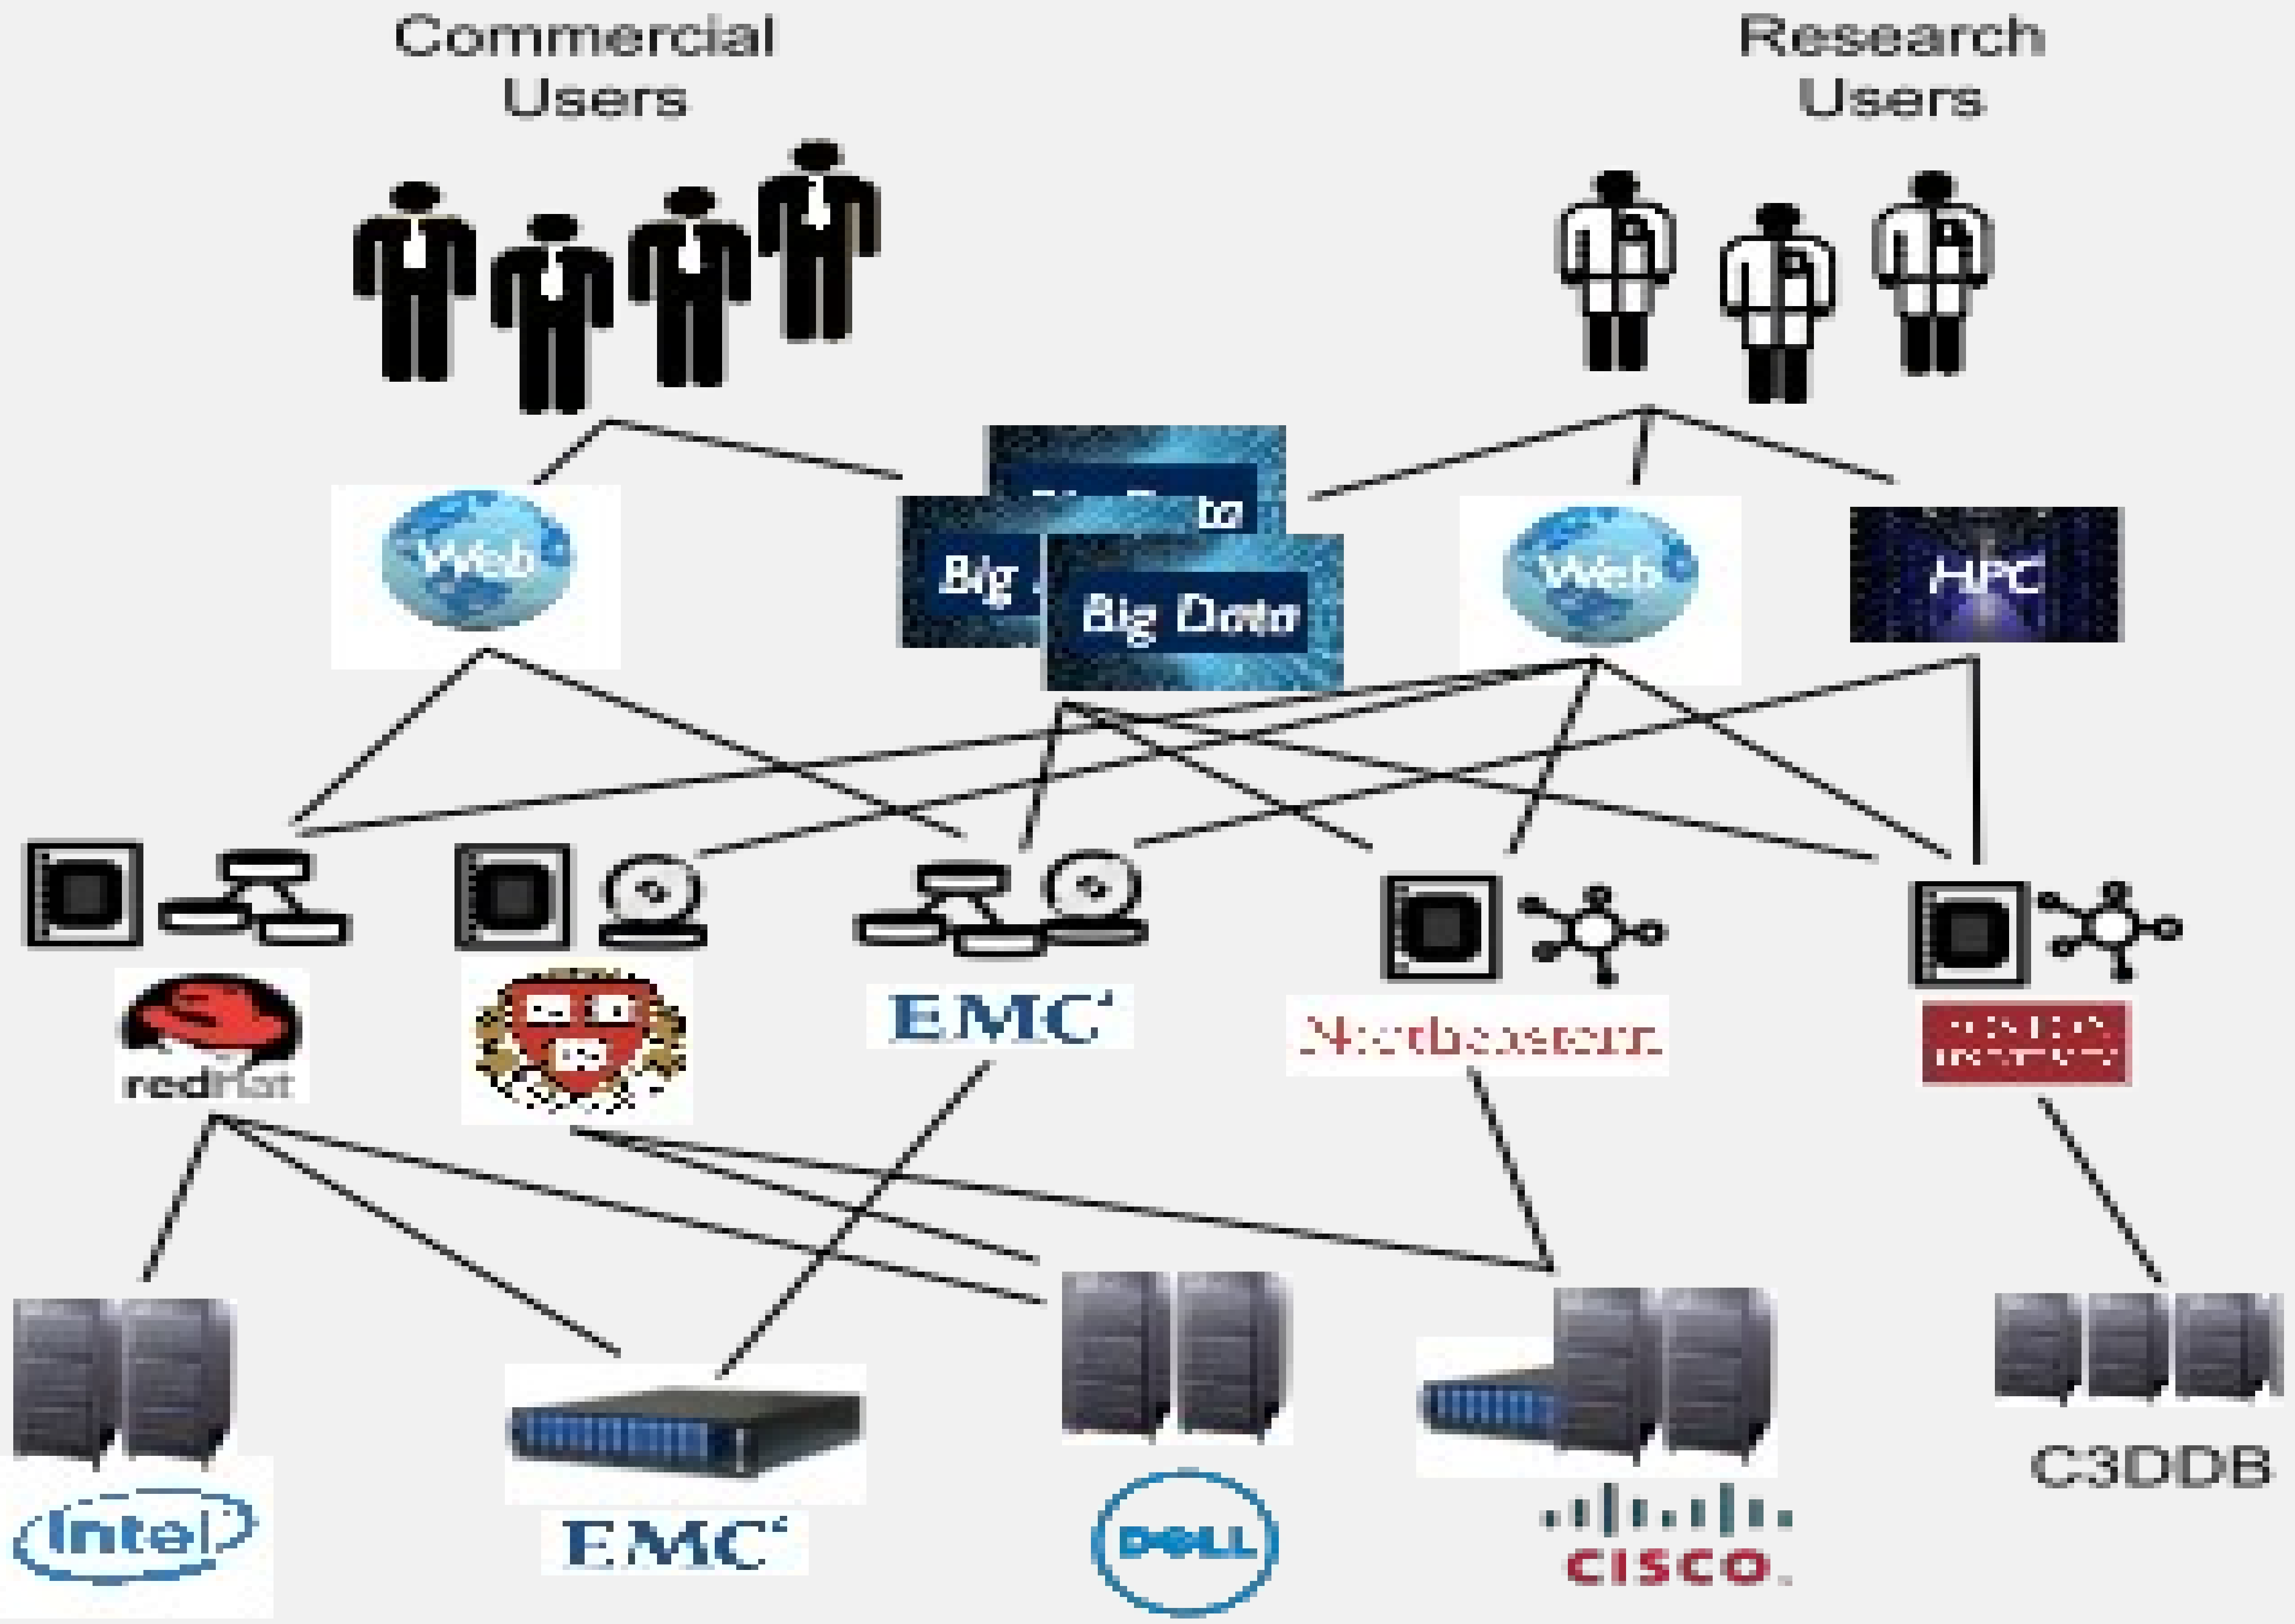

- Harvard & MIT private clouds provide resources MOC
- EMC stands up storage service, sets price, offers MOC and university clouds
- Company offers compute service with FPGAs, users can switch without moving data
- Research group stands up highly elastic experimental compute service
- Security sensitive company uses MACS SW to distribute across non-colluding providers

Why Academic involvement cloud?

Involvement Academia critical

- Education and workforce development
- Research:
  - Systems
  - Economic models...
- Challenging users of a cloud

The MOC is creating an at-scale testbed enabling research to get involved in developing, operating and using the cloud.

Cloud Computing Course

- Joint BU/NU (HU future) course on cloud computing for workforce development, early user access, industry engagement
  - 60 senior and graduate students first users of the MOC
  - Michael Daitzman (VECNA) teaching agile methodologies
  - Students work as agile teams mentored by industry partners

- 13 projects: Red Hat, CISCO, BCH, MBTA/MTC, Jackpine, Jisto, MOC
  - Create valuable artifacts around the MOC and exercise MOC e.g., for Big Data applications
  - Examples:
    - Enabling life sciences users: Radiology in the cloud
    - Public data sets: MBTA bus performance – 1 TB dataset
    - Startup: Network-aware container distribution
  - Four student projects enhancing the MOC: HaaS, OS intermediary, OCX UI, and OCX extensibility

“The engagement with the BU students has been beneficial for us and for them. They got first-hand exposure to cutting-edge IoT topics, we received two applications which will enhance greatly our Cisco Live PoC demo”, Lionel Florit, Principle Engineer, Cisco

Agile development importance

Quotes:

“Agile development techniques and sprint planning. This was incredibly valuable and will serve me well in my future internships and jobs.”

“How to manage agile sprints. How to work cohesively with others...”

“The importance of team work”

“Rather than any 1 topic, the experience of having a real project that we tried to approach like a real development team was the best part.”

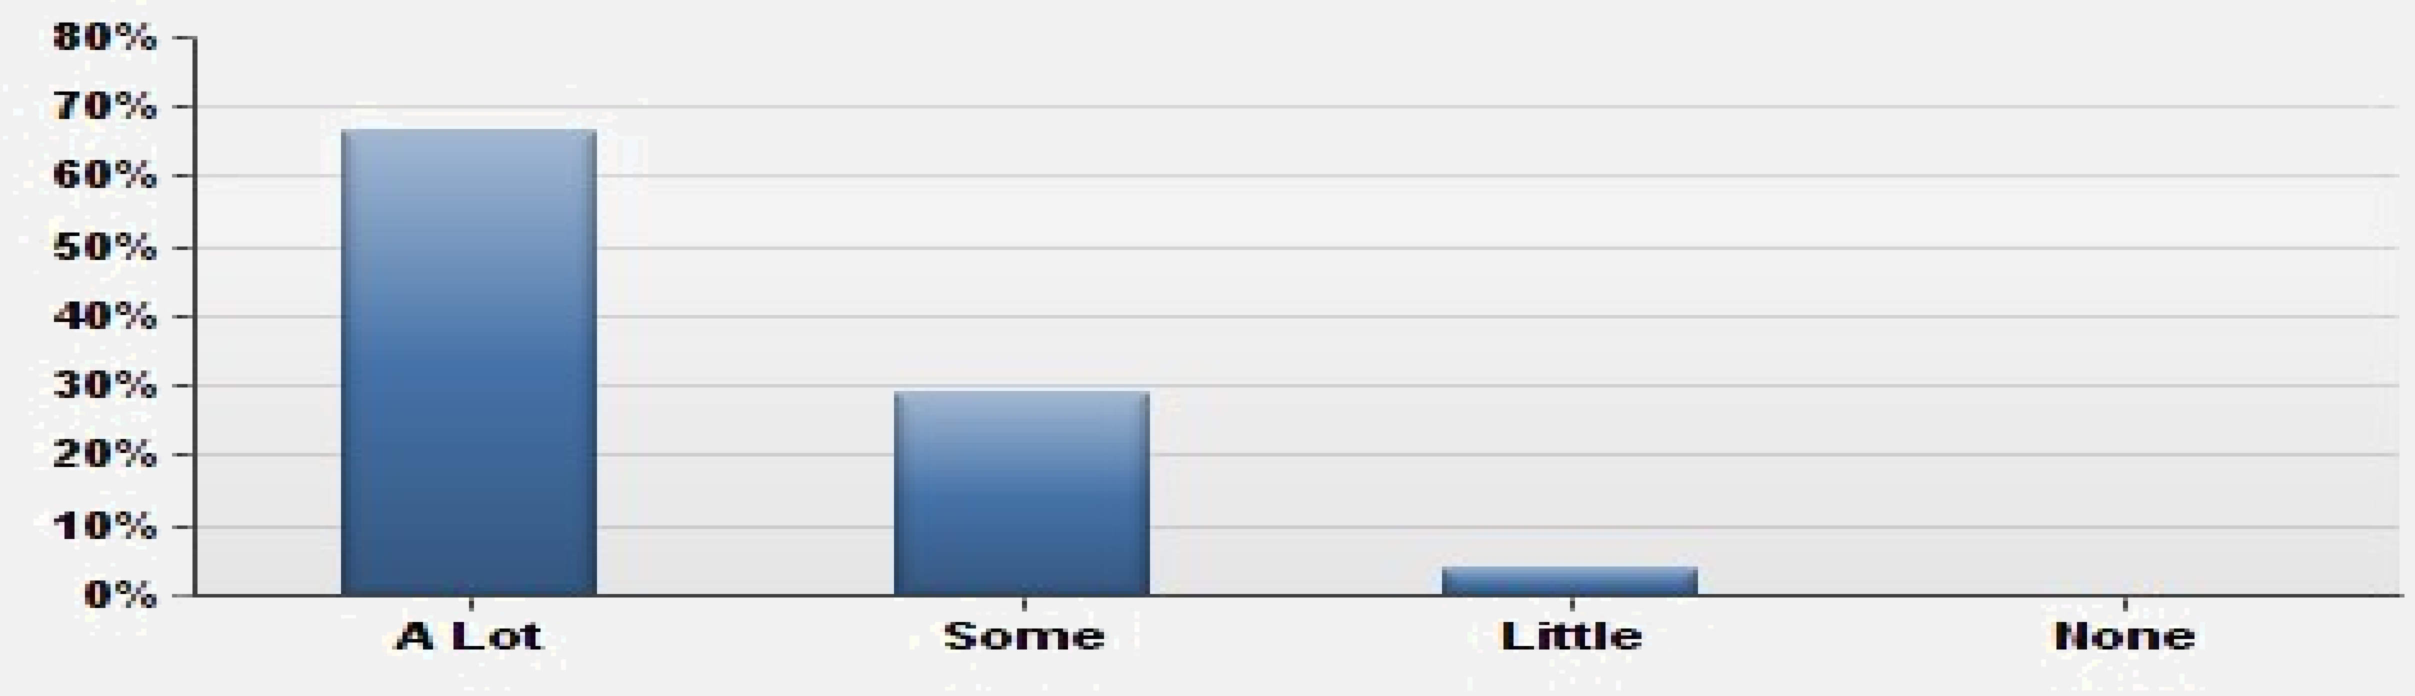

Preparedness Industry

- Gain
  - 30% prepared industry
  - 70% cloud concepts
  - 50% apply cloud position
- 6 students say course led to offers

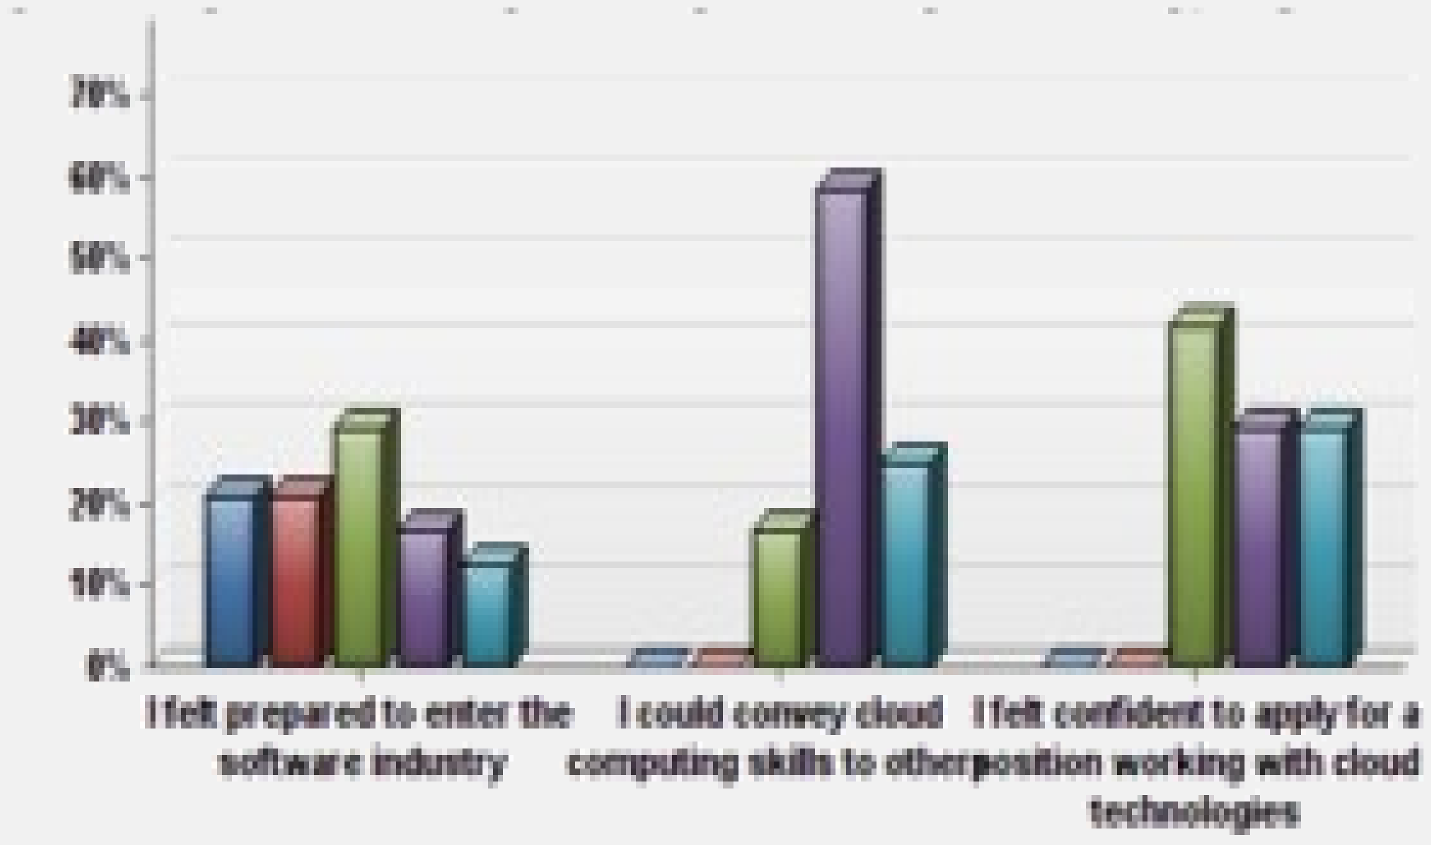

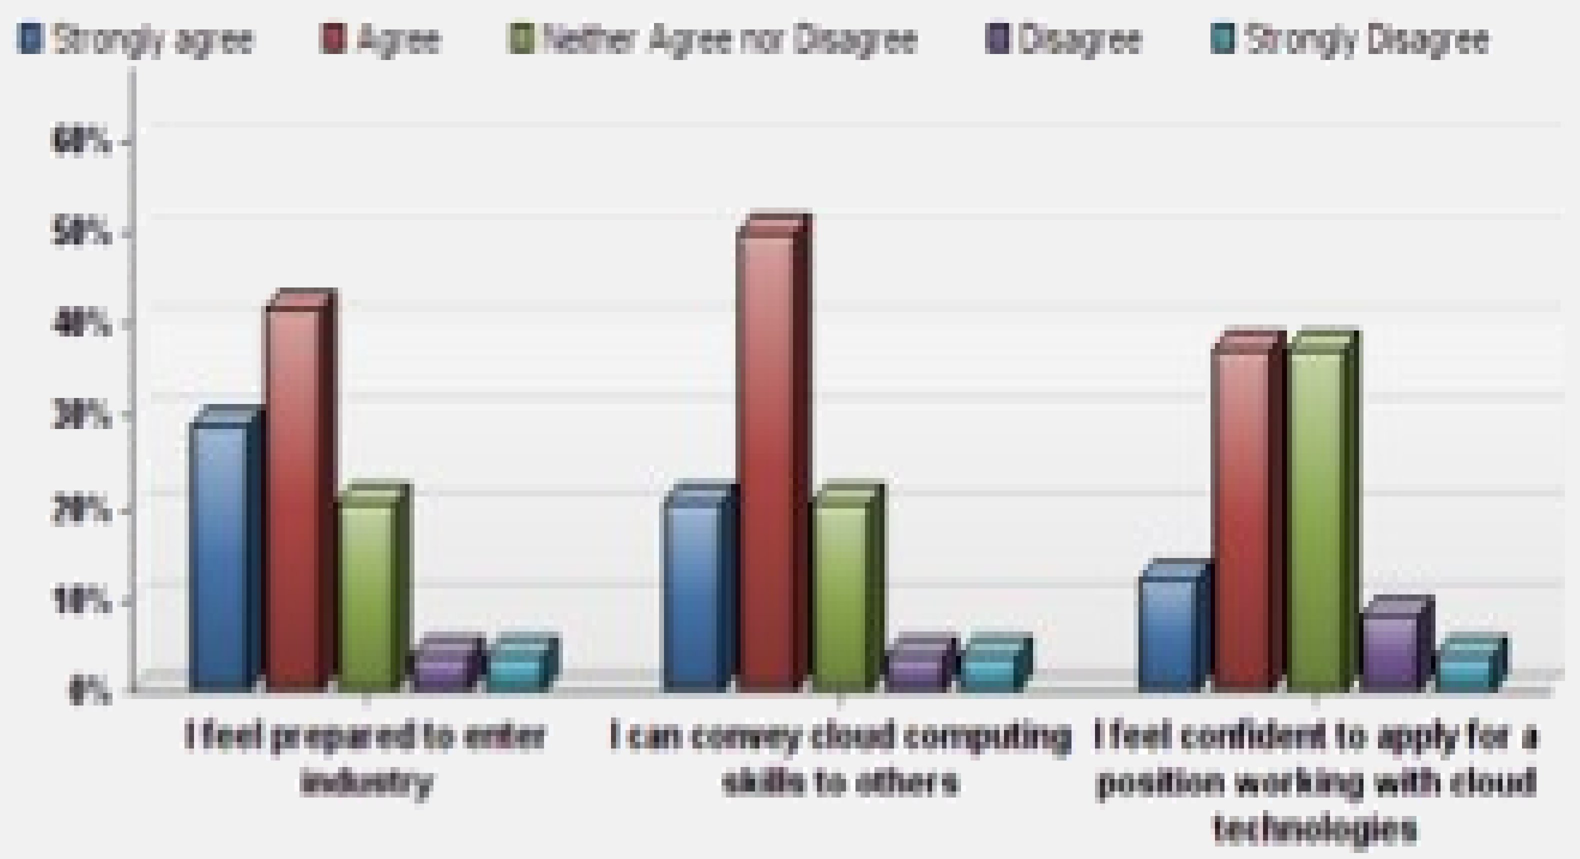

Federally Funded Research

November 2014 SCOPE Project Kickoff

- **SCOPE: A Smart-city Cloud-based Open Platform and Ecosystem**
  - $1M NSF Research Grant plus $1 M industry match
  - Involves cross-disciplinary team: BU, public and private sector partners
- **MACS: A Modular Approach to Cloud Security**
  - $10M NSF Frontier Research Grant
  - Involves cross-disciplinary team: Cryptography, Systems, DB, Cloud, Networking
- There are whole series of proposals being developed.

Technical Progress

Red Hat's interest

- Open Multi-vendor/landlord environment
  - Interoperate in a neutral multi-vendor environment
  - Build new/innovative services
  - Demonstrate value of services based on capabilities/cost

- Explore new use cases and approaches
  - Multi-vendor/application/service billing and services catalog
  - Per-service/capabilities choice

- Work collaboratively to develop a community of communities

- Open and collaborative environment not controlled by a single technology company

Start with OpenStack

- Open Source
- Most of needed functionality
- Broad acceptance:
  - large developers community
  - industry support
- Highly modular design
- Standard APIs allow flexibility in underlying implementation
- Rich plugin architecture

- However, assumes single provider and one instance of each service
  - e.g. one SDN controller in entire data center.

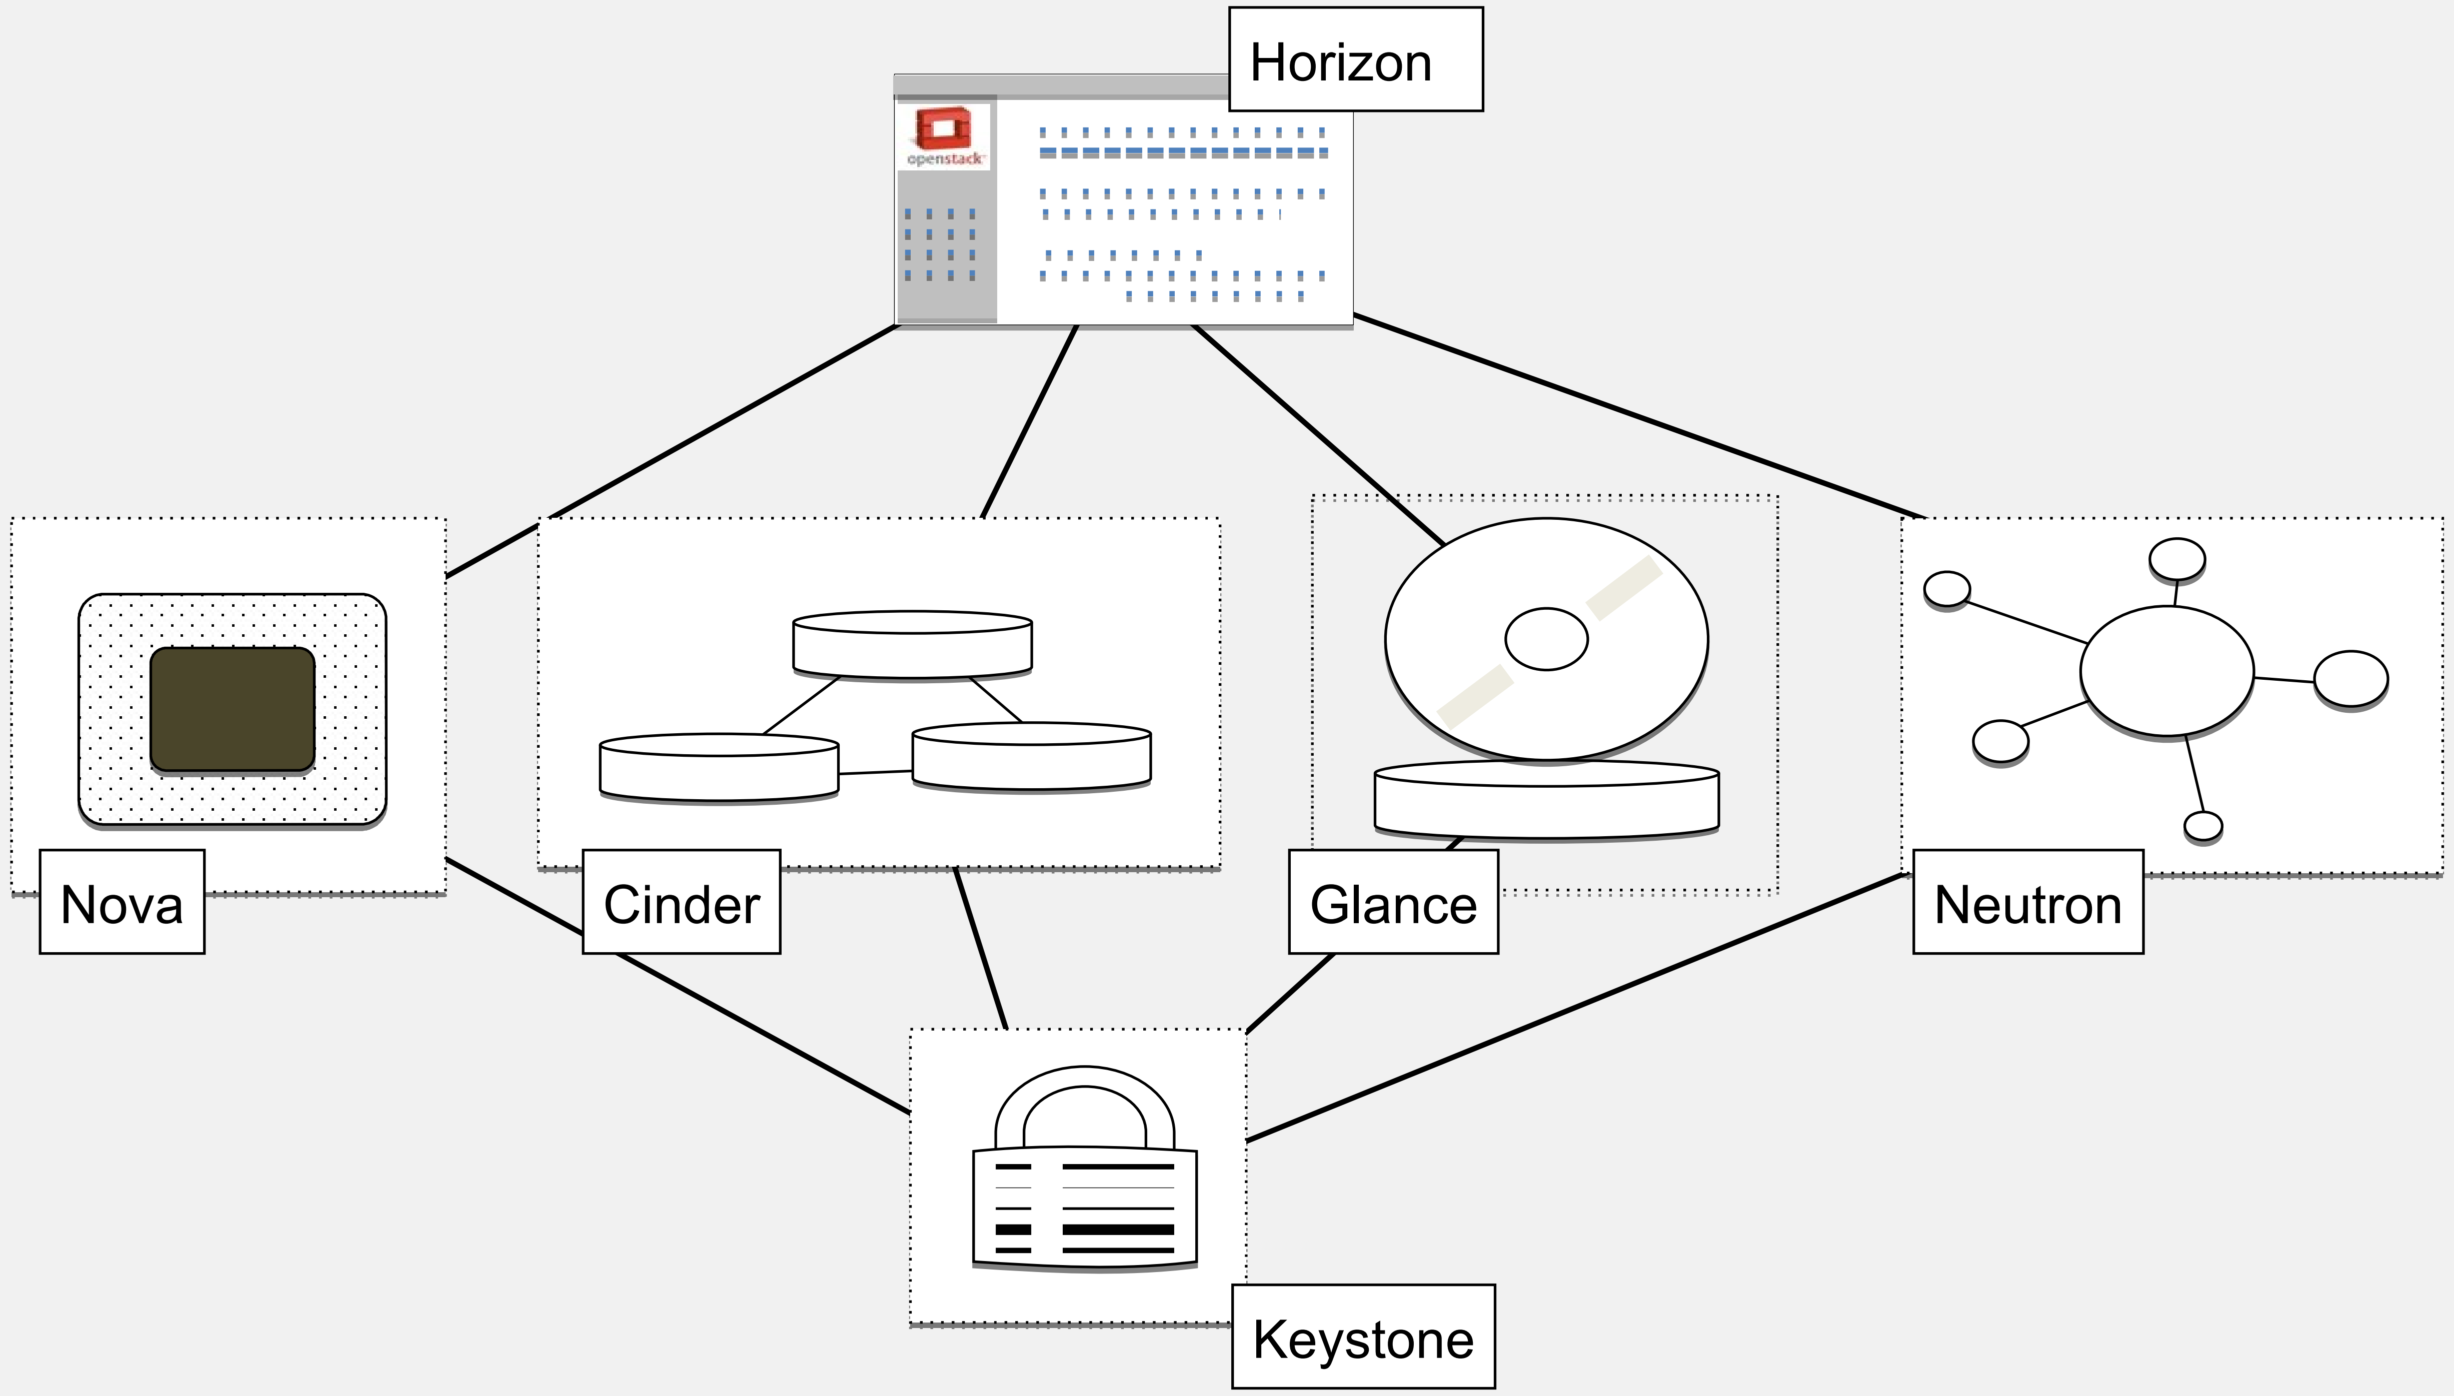

We want OpenStack to handle this

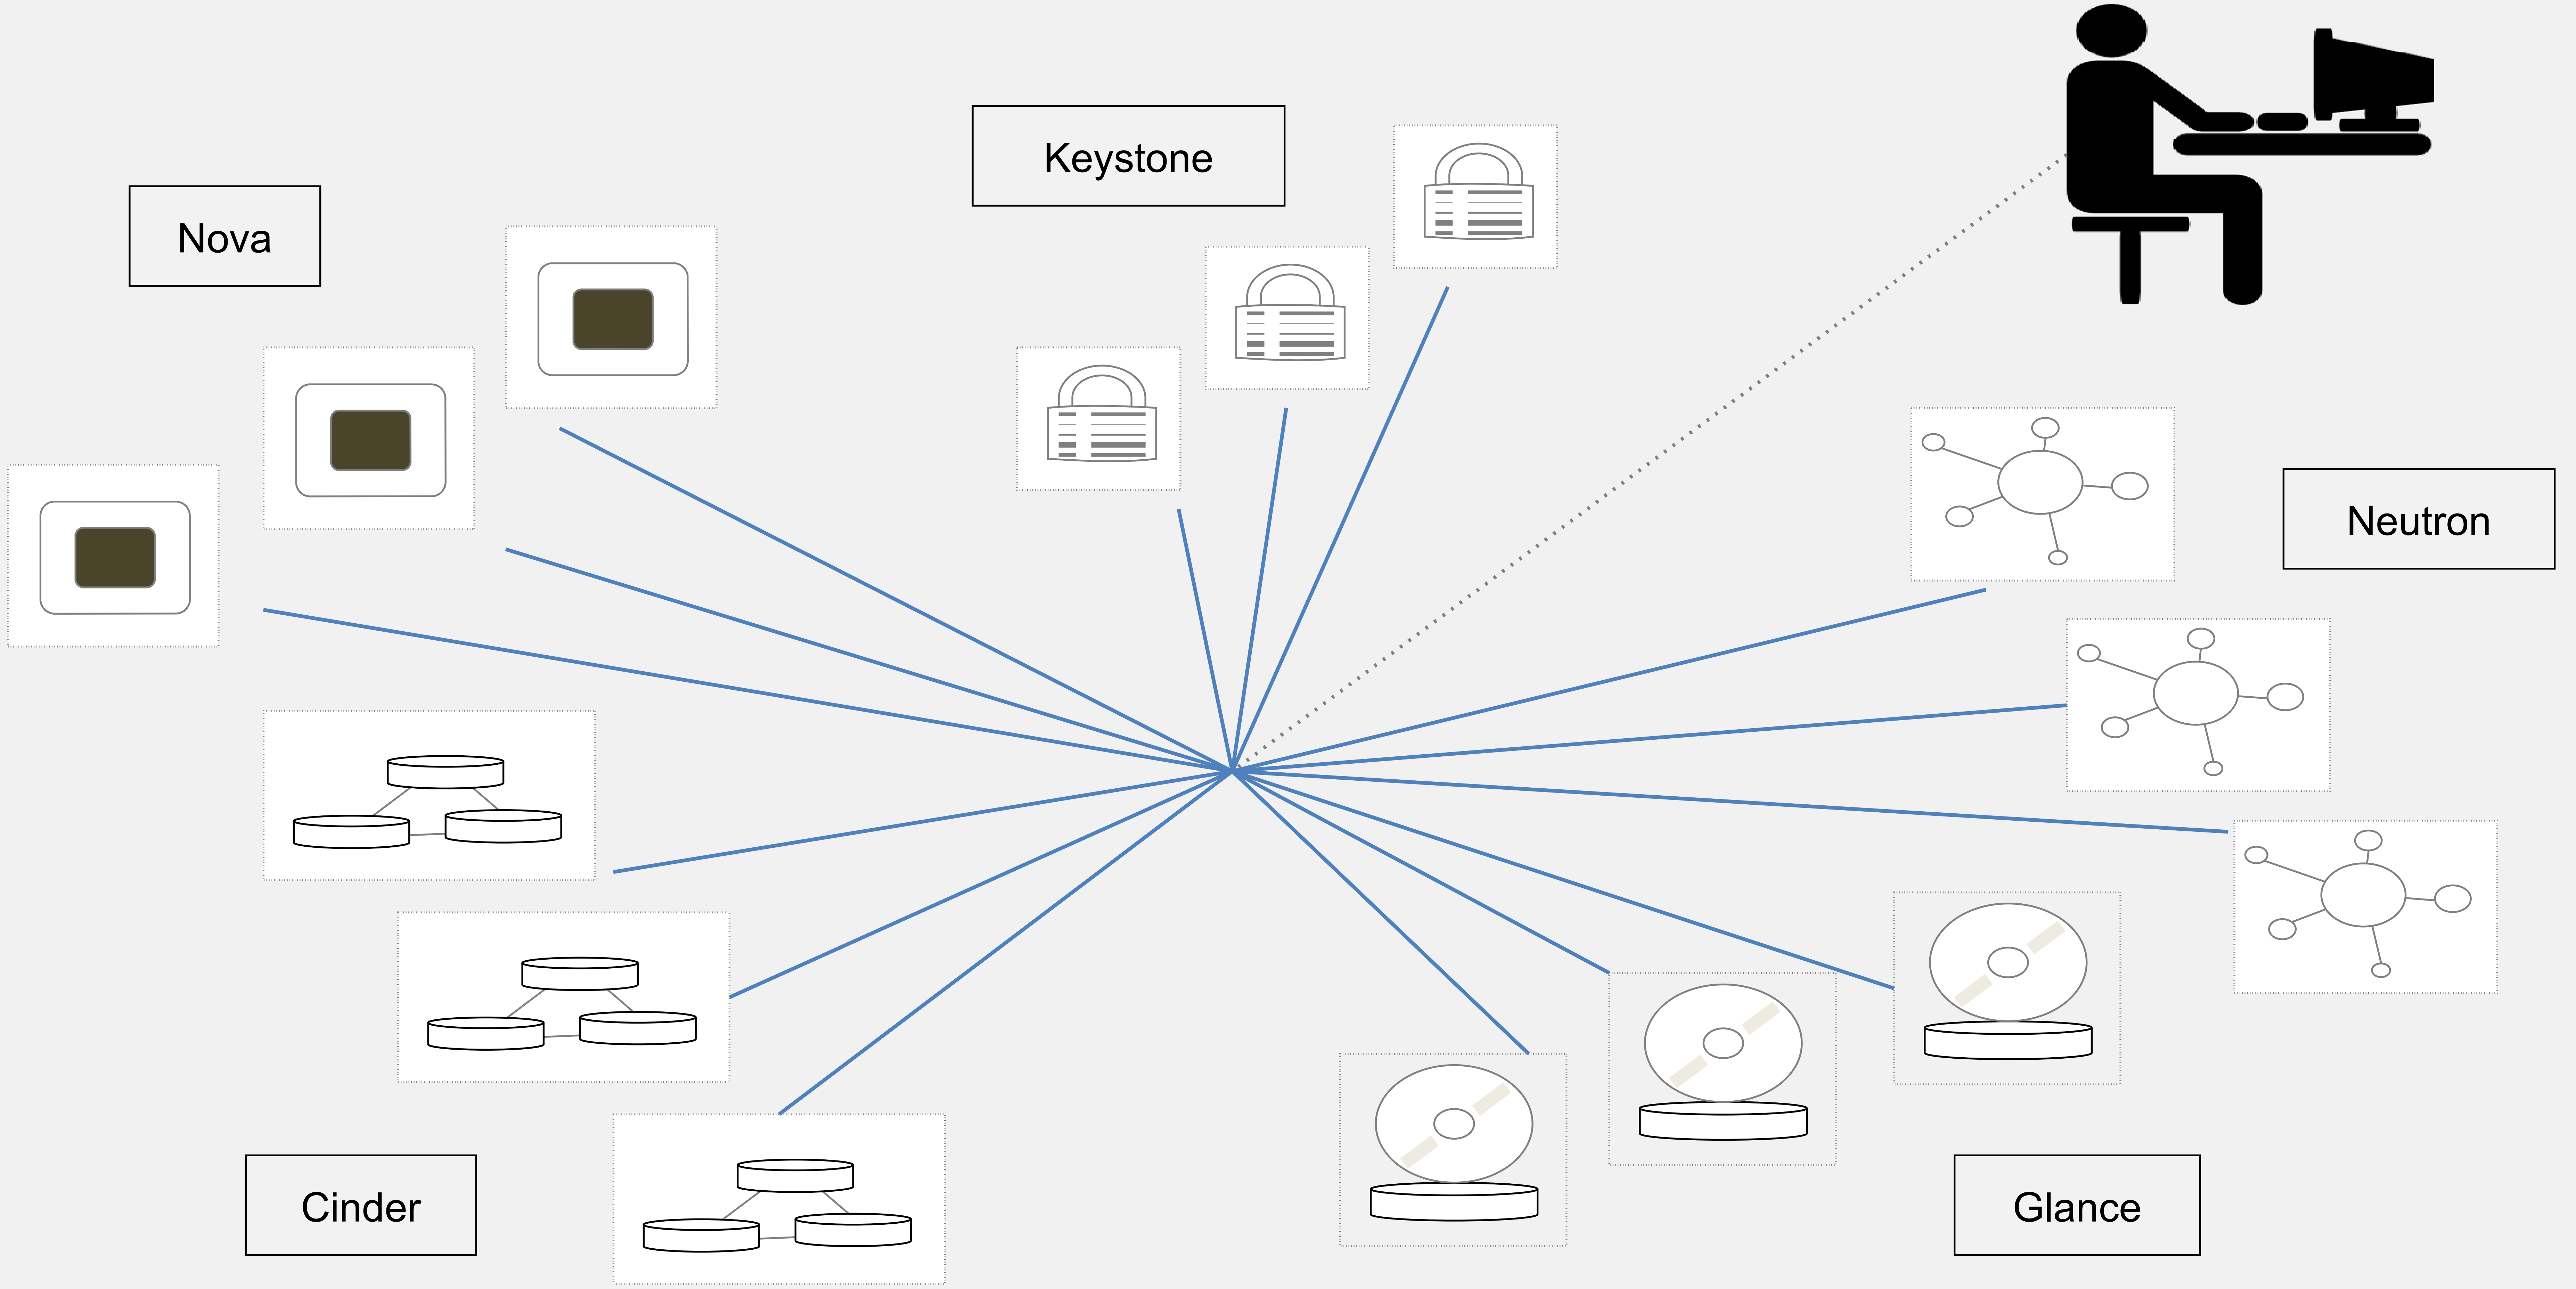

Required changes

- Changes to core OpenStack: limiting trust and enabling mix &match
  - Federated authorization - new plan central keystone
  - Fully qualified UUIDs - blueprint; limits mix & match per project until have solution
  - Per-object access control - long term critical, may be able to exploit hierarchical multi-tenancy
- Extensions: client drives choice
  - OCX Service Directory - marketplace longer term
  - OCX Library - development in progress
  - OCX GUI - moving to Horizon
  - Hardware as a Service - initial development complete
- Constraints: network spanning, storage protocol

Alignment Mercador project

https://wiki.openstack.org/wiki/Mercador

- Initiated by Geoff Arnold (Cisco) in support of Intercloud federation vision
- Builds on Hierarchical multi-tenancy and federated authentication: federation with quotas
- Observation:
  - Orthogonal and complementary to OCX changes
  - Provides functionality required in long run for OCX model
  - OCX functionality valuable for some Intercloud use cases
- Strategy: Join forces, demonstrate both sets of use cases & how work together by Tokyo summit

Hardware as a Service (HaaS)

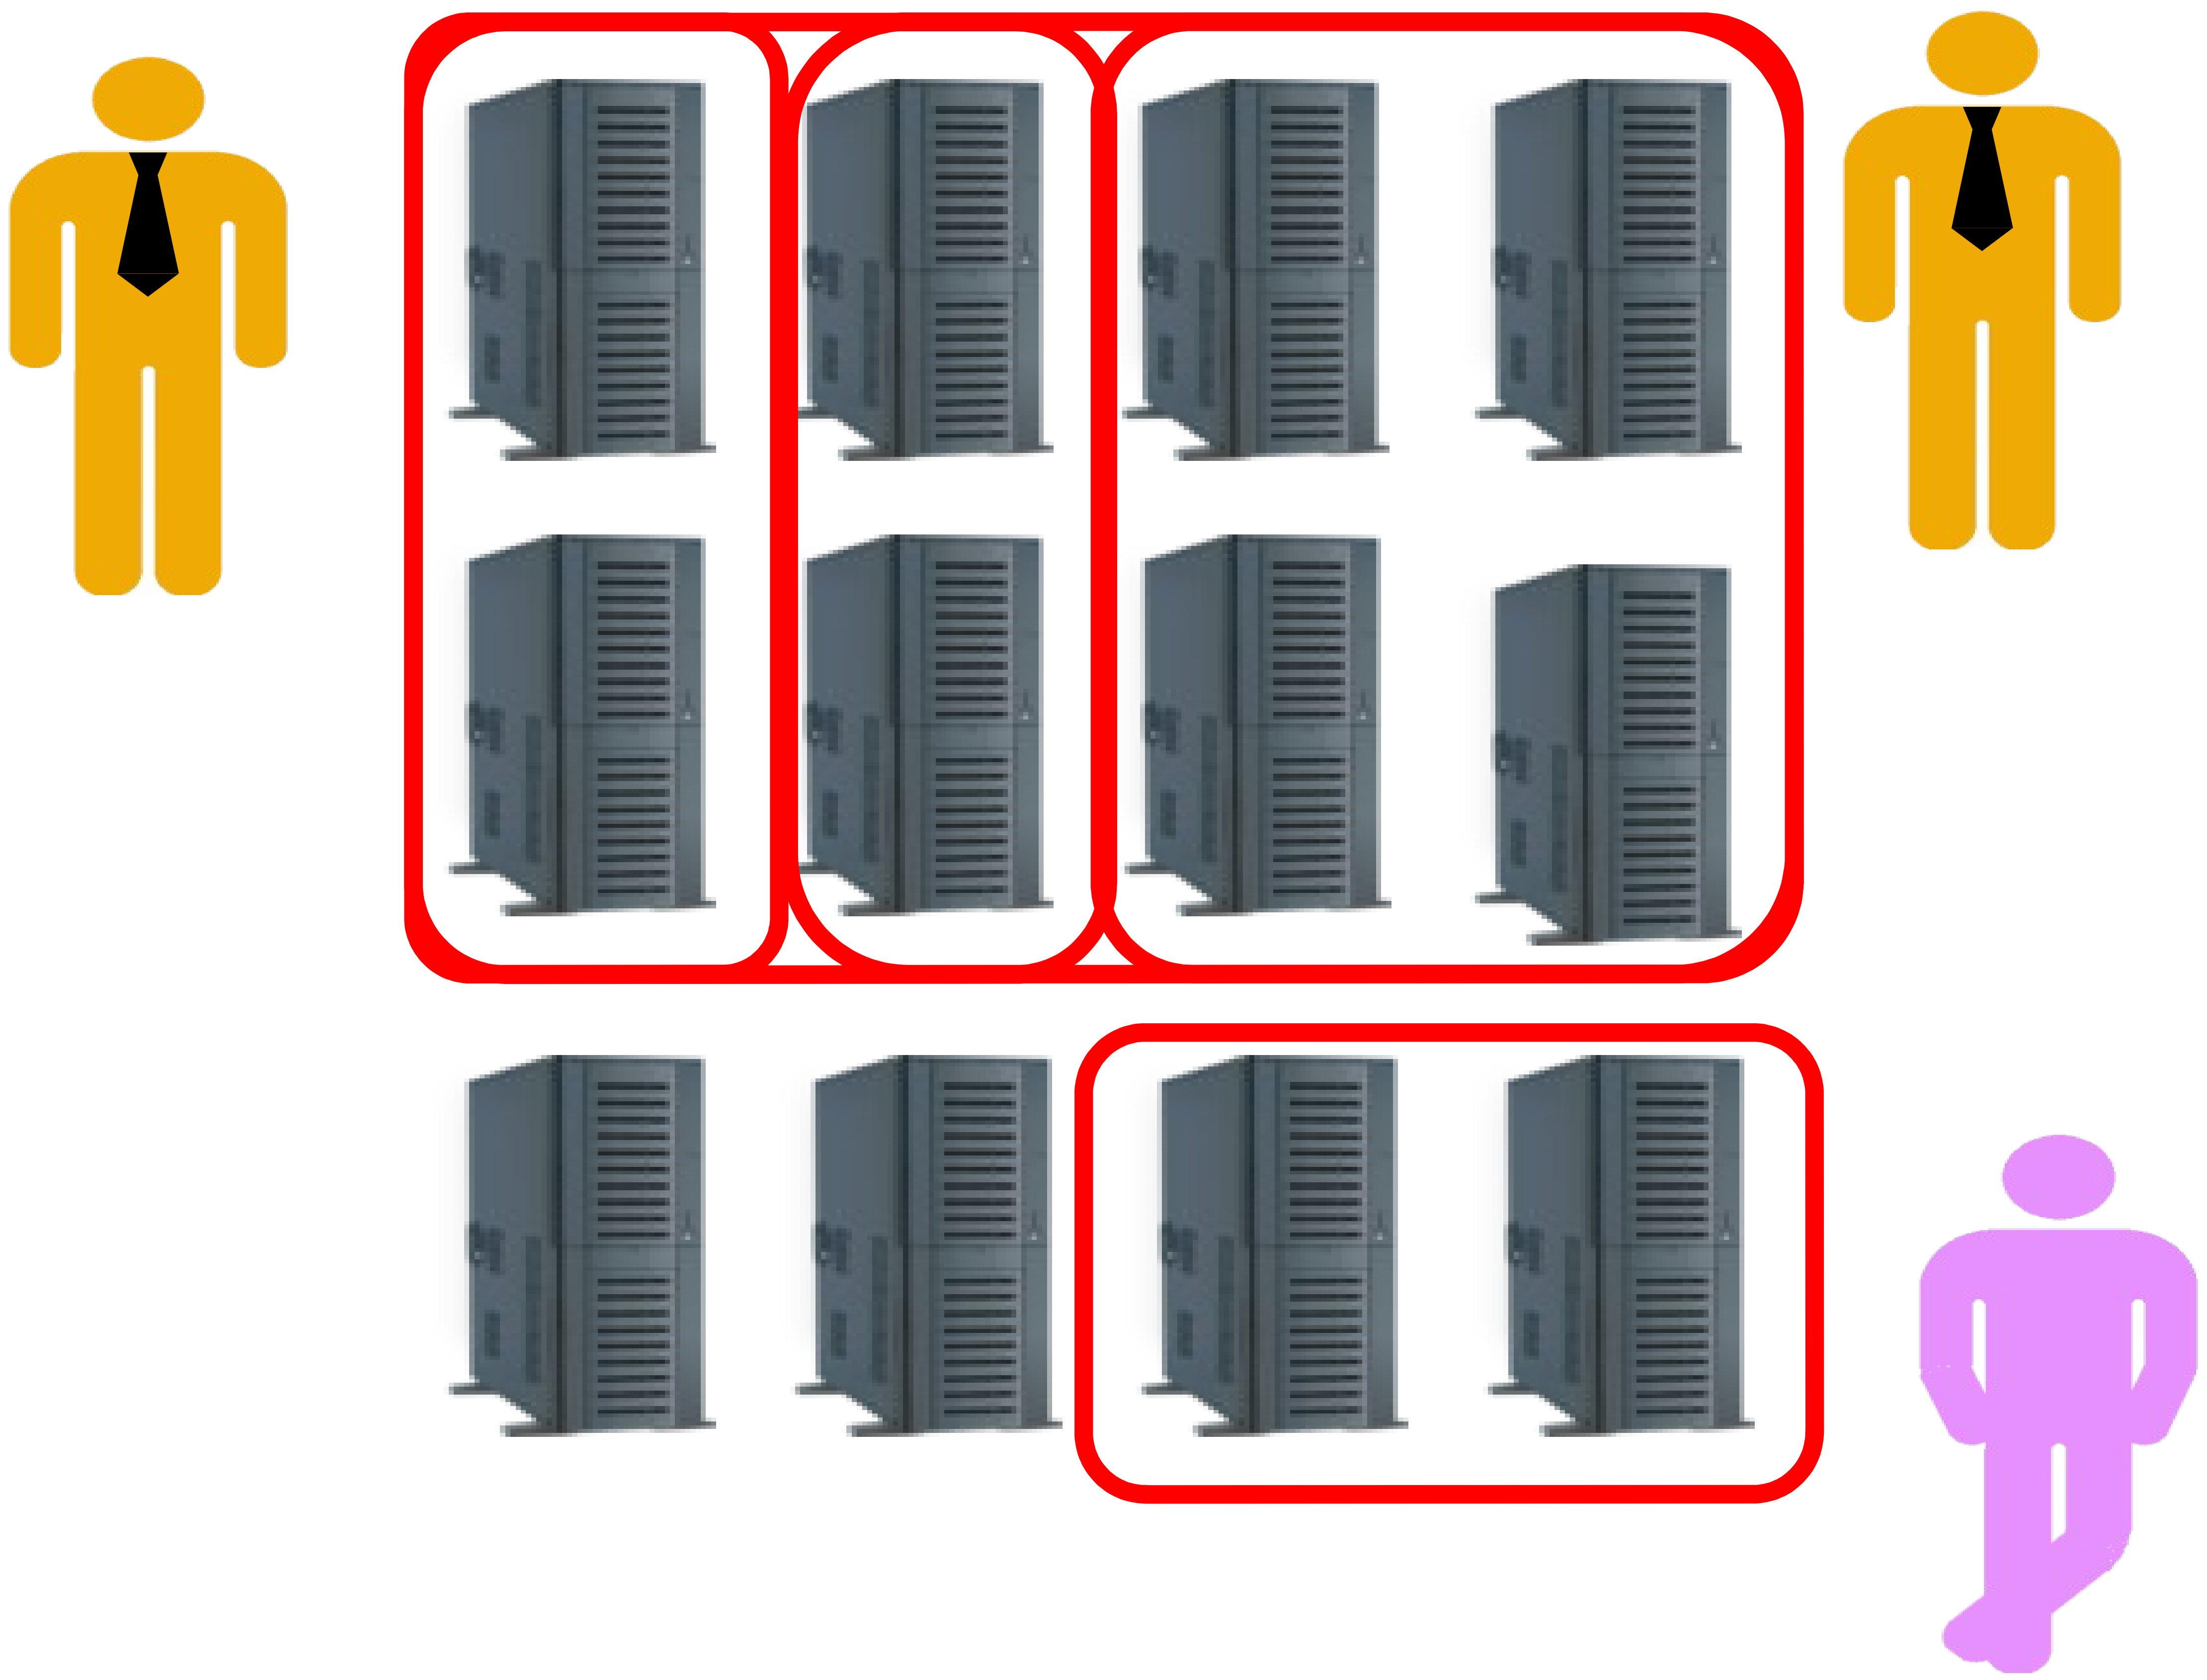

- Existing data center clusters can be incorporated with no changes
- Allows resources to be moved depending on demand to improve system utilization
  - Extending marketplace model all the way down to the hardware
- New services can be deployed without trusting administrator, enabling:
  - Virtualization research
  - New cloud services
  - Applications requiring access to specialized HW
  - Reproducibility

HaaS & Security

- Security implications: simple/minimum TCB, tenant controls provisioning,
- Secure tenant data center in the middle of a public cloud data center:
  - full bi-sectional BW to public data sets
  - wherever resources needed: e.g., content distribution and ingestion
  - requires: secure FW, allocate TOR switch, different network topologies, HaaS admin
- In emergency, e.g., Boston Marathon Bombing, warp out secure DC to consume most resources.
- Broad community (e.g., financial, life sciences, enterprise, technology companies) share hardware and common technologies
- Compartmentalize different financial modeling teams, defense teams…
- Involving USAF, academic and financial partners in evolution of technology

Technical Progress

- Devops Team
  - Senior technical manager
  - Three junior programmers
  - Two Postdocs, USAF, CISCO, Red Hat and growing team of students (12 this summer)
  - Partnerships with IT organizations at Boston University, Harvard, Northeastern, and MIT/CSAIL
- Development progress:
  - Hardware as a Service (HaaS)
  - Submitted Blueprints for key changes to OpenStack to enable OCX model
    - Federated Authorization, Fully qualified UUIDs, Per-object auth
  - Puppet manifests
  - UI Builds for the marketplace (integrated and out of stack)
  - End to end demo for Tokyo Summit
  - Changes to Ceph to enable Spark and big data integration

Operational Progress

- HW infrastructure
  - current: Pool of 16 Dell servers in HU portion of MGHPCC deployed RH (RDO/Foreman) OpenStack
  - 48 CISCO servers racked and tested in NU part MGHPCC
  - Fujitsu Ceph appliance for initial production storage
  - Extend to 200 research servers + Big Data platform collaboration between Intel, RH, Brocade, Lenovo
  - Incorporating CSAIL environment (80+ servers, 11k VMs)
- Monitoring infrastructure on Sensu
- Bring your own provisioning infrastructure
  - Foreman and puppet included in Satellite
  - Importing custom developed manifests
- Logging using Logstash and Elasticsearch – do users get access?
- HA and multi-node – follow the docs
- Neutron tenant networking model with VPNaaS → Wouldn't a standard be nice between SDNs?

Red Hat contributions

- “Red Hat has been the anchor partner since the start of the project, and was the critical partner in the original proposal to the commonwealth” -Orran

- Deep technical engineering engagement:
  - Implementation changes for OCX model, enhancements to Ceph on Intel/Brocade/Lenovo Big Data platform, OpenStack changes, Partner engineering
- Operations involvement helping stand up and operate the OpenStack envioronmnet
  - Red Hat Cloud Innovation Practice Team
  - Cloud Technical Account Manager
  - Subscriptions for RHEL-OSP6 and RHEL-OSP7 Beta, Ceph, Cloudforms, Satellite
- Business model guidance with senior level technical engagement

Lessons learned

- OpenStack is like Linux in 2003, massive maturation, operational challenges
  - Moving to more enterprise level management: Foreman, puppet
  - Hugely valuable to have partner for support
- Technologies of cloud fail at small scale:
  - Initial prototype used Ceph on 3 VMs and only 12 disks; **huge** mistake
- Challenges/benefits open source community
  - Federated authorization using SAML
  - Abandoning OCX UI to integrate with new Angular JS Horizon strategy
- Educational aspect more critical than realized
- Fundamentally new Business model ongoing challenge

RED HAT
SUMMIT

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.