

A marketplace driven direction for SDI: Towards an “Open Cloud Exchange”

Orran Krieger (BU),
Peter Desnoyers (NU), John Goodhue (MGHPCC), Azer Bestavros (BU), Johnathan
Appavoo (BU), Margo Seltzer (Harvard), Larry Rudolph (TwoSigma), Rodrigo Fonseca
(Brown)...


Definition of SDI

Software-defined infrastructure (SDI) is the definition of technical computing infrastructure entirely under the control of software with no operator or human intervention. It operates independent of any hardware-specific dependencies and are programmatically extensible

That's great, but...


- The assumption is that one entity controls the SDI.
e.g., one controller for SDN, one entity stands up each cloud, ...
- Have Contrail, ODL, Midonet, Big Switch, Nicera, Pexxi... all saying solve problems, but they are the solution

For example, consider today's IaaS clouds



For example, consider today's IaaS clouds

- One company responsible for implementing and operating the cloud
- Typically highly secretive about operational practices
- Exposes limited information to enable optimizations



What's the problem

- Lots of innovation above the IaaS level... but
 - consider EnterpriseDB, or Akamai
- Lots of different providers... but
 - bandwidth between providers limited
 - offerings incompatible; switching a problem
 - price challenges to moving
- No visibility/auditing internal processes
- Price is terrible for computers run 24x7x365


More challenges

- Provider incentive not aligned with efficient marketplace:
 - stickiness in price, in differentiation
 - advantage other services
 - homogeneity for efficiency
- Hard for large provider to efficiently support niche markets, radically different economic models...
- Niche providers probably can't support rich ecosystem


More challenges

- Provider incentive not aligned with efficient marketplaces
 - sticking
 - advanced
 - homogenized
 - Hard for large provider to efficiently support niche markets, radically different economic models...
 - Niche providers probably can't support rich ecosystem
- We are in the equivalent of the pre-Internet world, where AOL and CompuServe dominated on-line access


Is a different model possible? An “Open Cloud eXchange (OCX)”




Is a different model possible? An “Open Cloud eXchange (OCX)”



Is a different model possible? An “Open Cloud eXchange (OCX)”

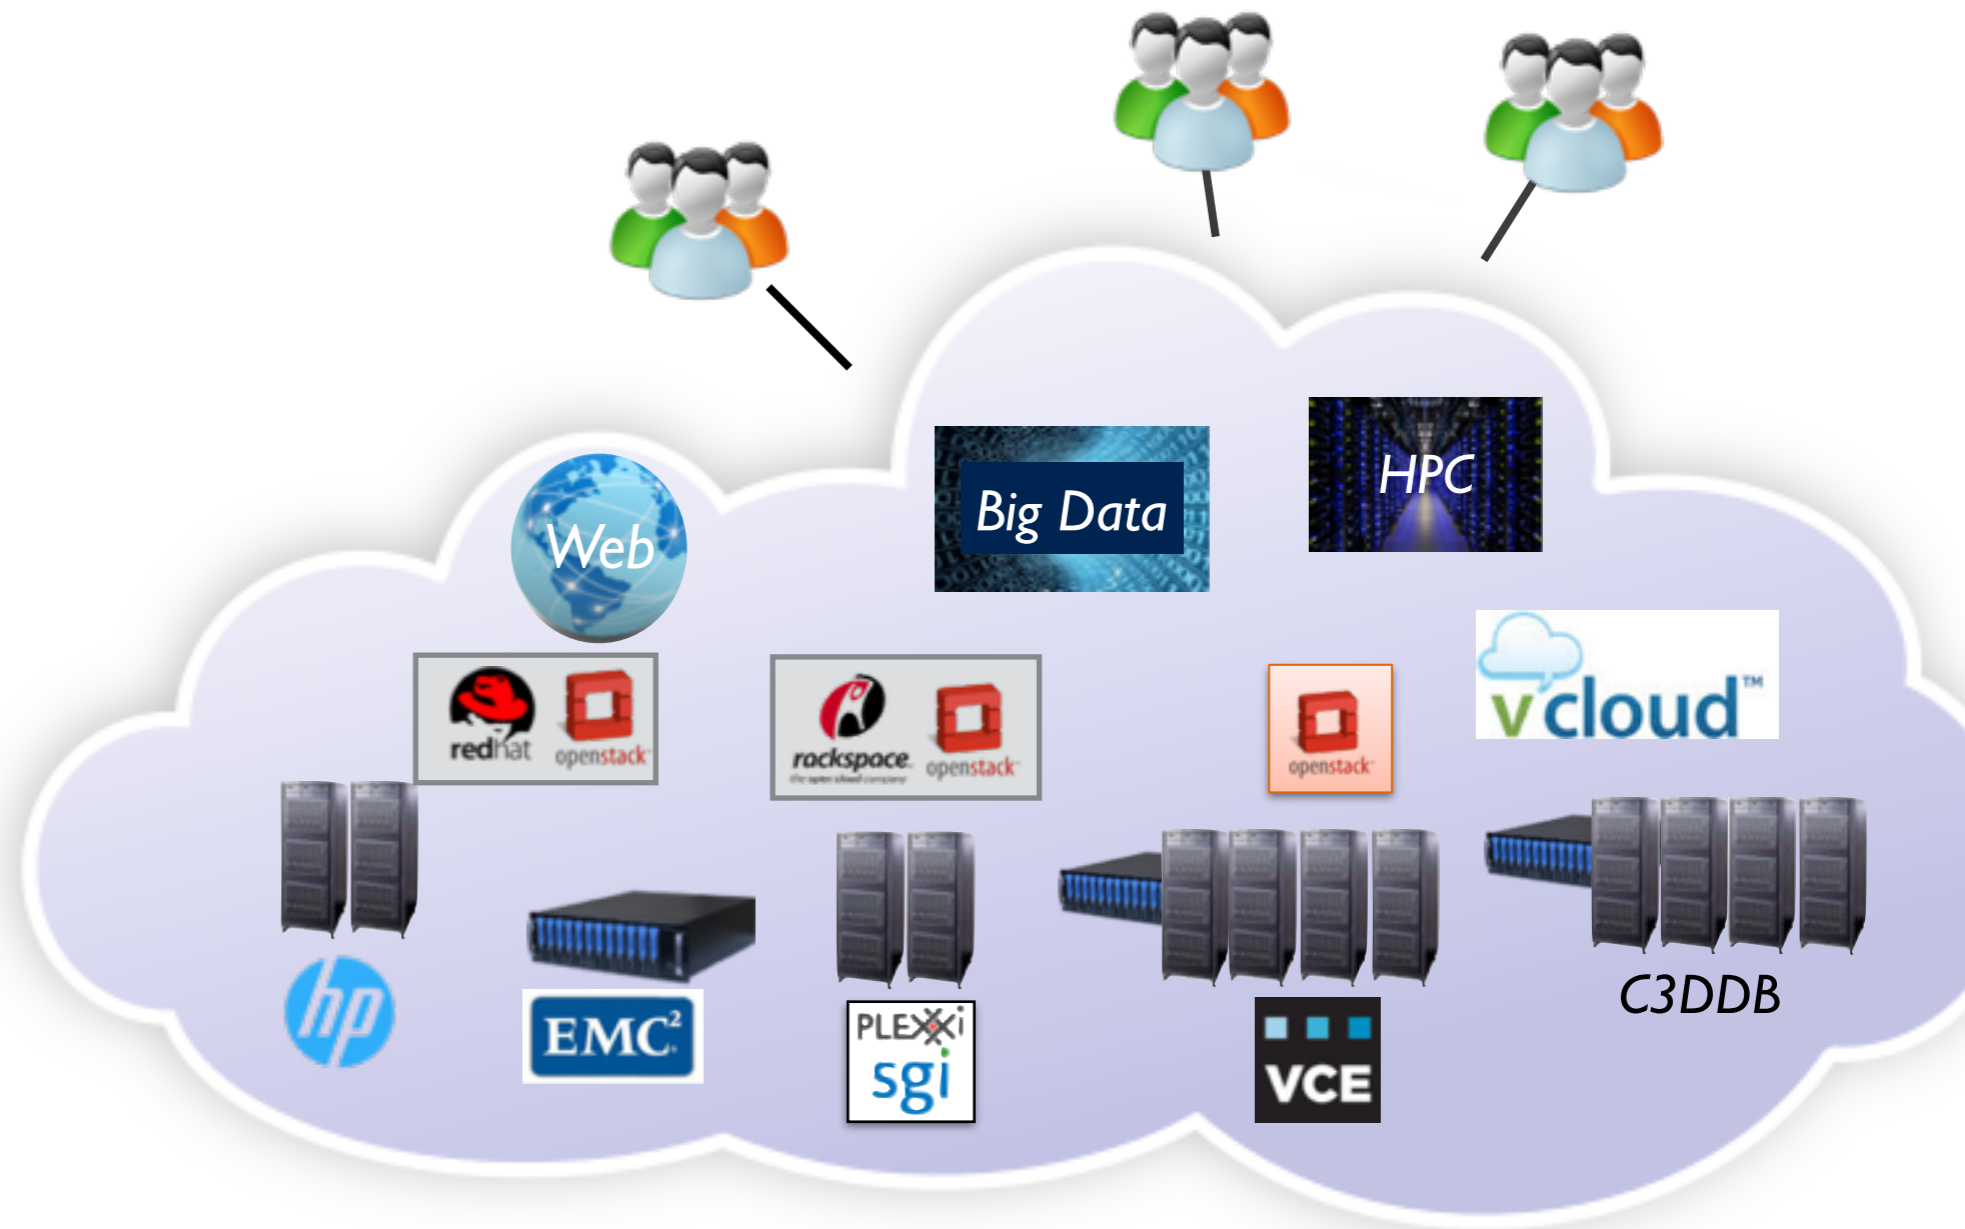


Is a different model possible? An “Open Cloud eXchange (OCX)”

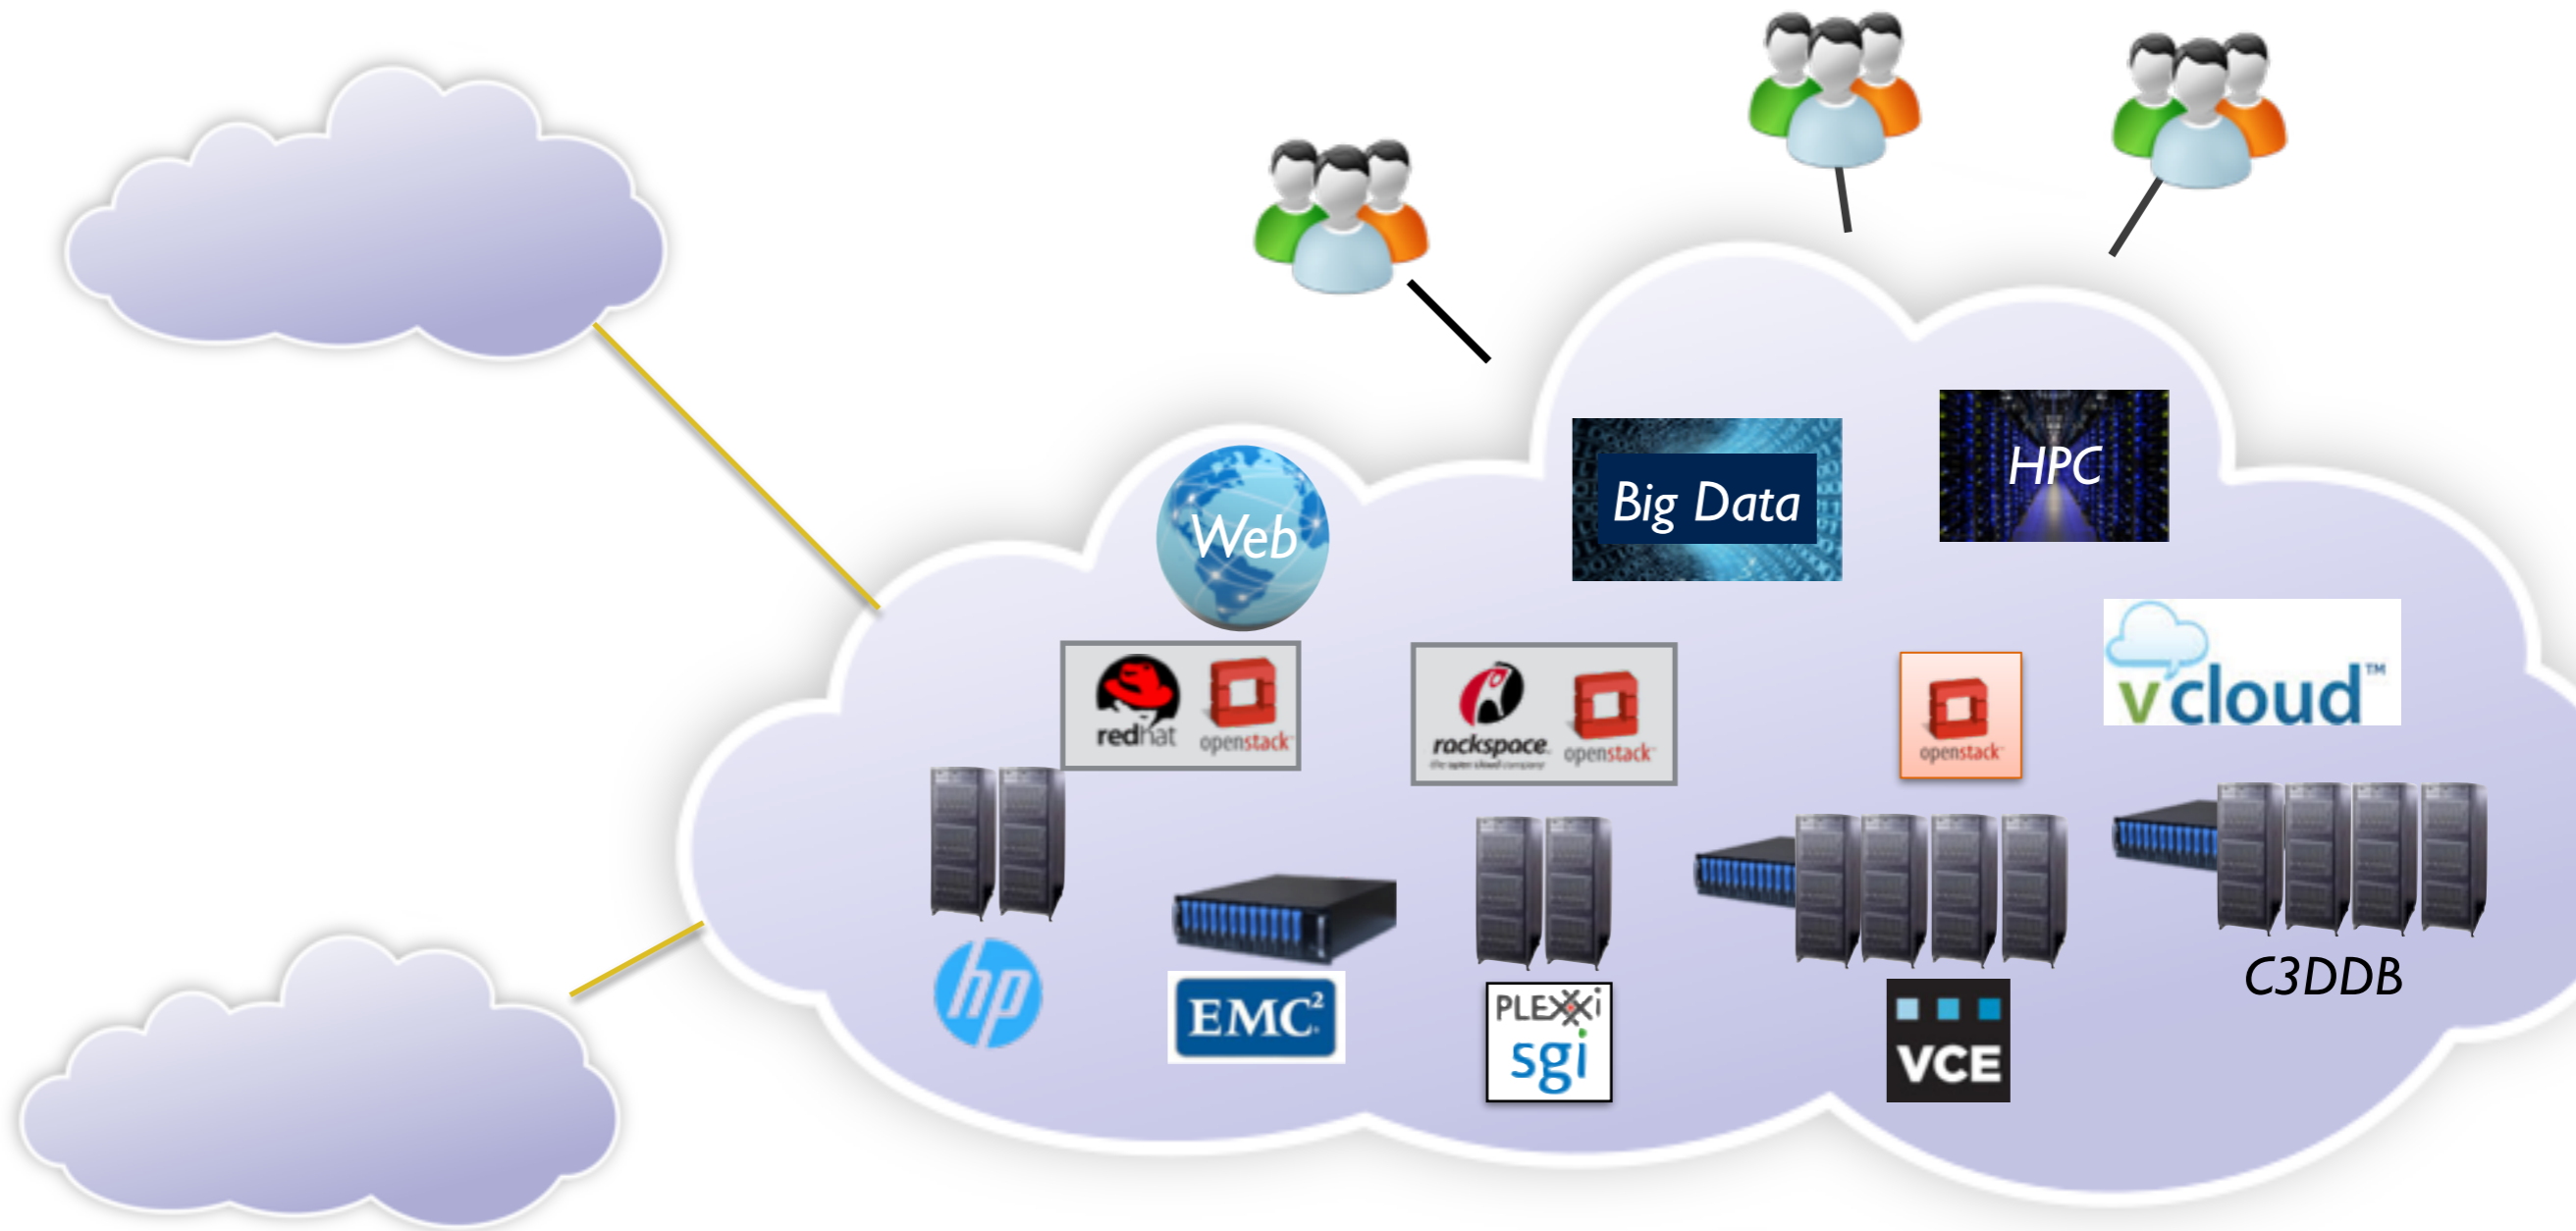


Is a different model possible?


An “Open Cloud eXchange (OCX)”




Is a different model possible? An “Open Cloud eXchange (OCX)”



Is a different model possible? An “Open Cloud eXchange (OCX)”



Is a different model possible? An “Open Cloud eXchange (OCX)”







Why is this important

- Anyone can add a new service and compete in a level playing field
- History tells us the opening up to rich community/marketplace competition results in innovation/efficiency:
 - “The Cathedral and the Bazaar” by Eric Steven Raymond
 - “The Master Switch: The Rise and Fall of Information Empires” by Tim Wu
- This could fundamentally change systems research:
 - access to real data
 - access to real users
 - access to scale
 - seamless boundary for transition to practice

Without that...solving the spherical horse problem...



This isn't crazy... really

- Current clouds are incredibly expensive...
- Much of industry locked out of current clouds
- lots of great open source software
- lots of great niche markets; markets important to us...
- lots of users alienated by business practices of big guys...
- this doesn't need to be AWS scale to be worth it
 - “Past a certain scale; little advantage to economy of scale” — John Goodhue

MGHPCC



15 MW, 90,000 square feet + can grow

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 Performance Computing - Sign up now!



The Massachusetts Open Cloud



We are confident initial service

- Changes to OpenStack to enable resource federation
- HaaS service
- Small number of production providers
- Small number of fixed storage providers
- Simple directory service locate all the providers
- Very simple pricing model
- Status:
 - Initial production 48 servers & 50 TBytes
 - Moving to larger number servers and .5PBytes storage in next few months
 - Simple federation with CSAIL OpenStack cloud

Research challenges

- Marketplace mechanisms (David Parks, Margo Seltzer):
 - How can users/intermediaries describe their needs/constraints/preferences?
 - How do providers describe their services?
 - How do we provide incentives for providers to provide high quality services? To provide rich information?
- Secure HaaS to enable marketplace all the way down
- Multi-provider cloudlet
- Software defined storage

Research challenges

- Marketplace mechanisms (David Parks, Margo Seltzer):
 - How can users/intermediaries describe their needs/constraints/preferences?
 - How do providers describe their services?
 - How do we provide incentives for providers to provide high quality services? To provide rich information?
- Secure HaaS to enable marketplace all the way down
- Multi-provider cloudlet
- Software defined storage (Peter Desnoyers)



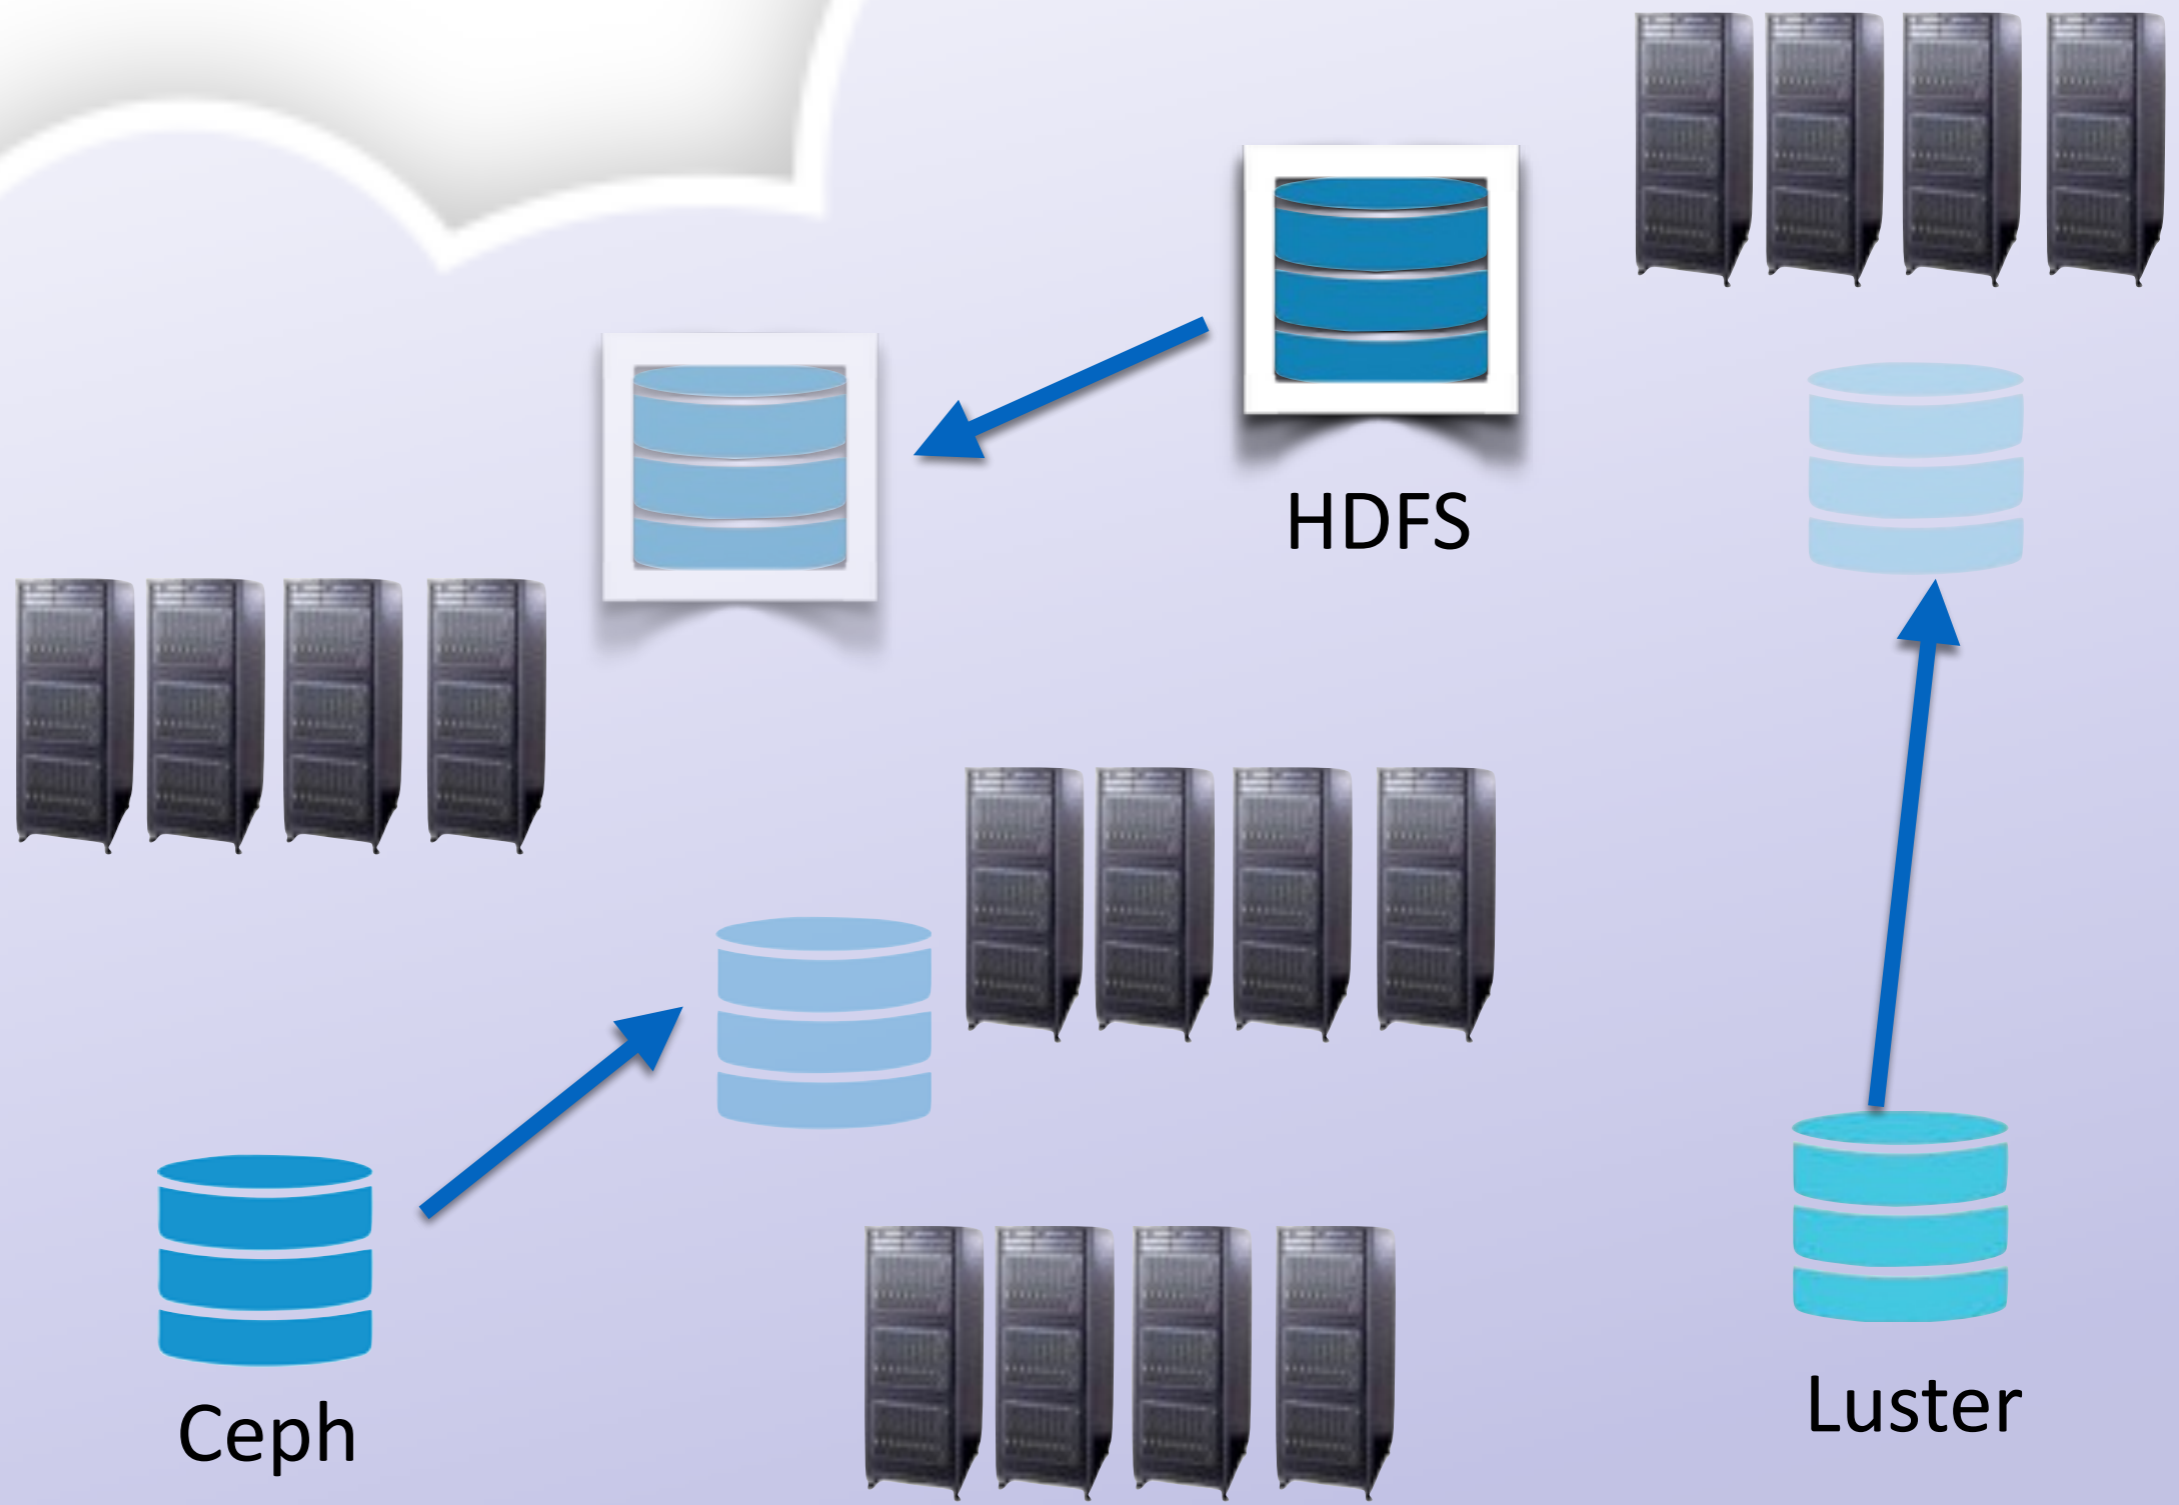
HDFS



Ceph



Luster



Research challenges

- Networking:
 - Multi-provider SDN, SDX on top of SDI in cloud DC, between cloud DC.
 - How give user choice of SDX?
 - Hierarchical SDN, app specific customization on IaaS on top of HaaS
- Can we expose rich information about services while not violating customer privacy
- How can we correlate between the information between the different layers?
- How can we identify source of failures?

Research challenges


- **Networking: (Rodrigo Fonseca)**
 - Multi-provider SDN, SDX on top of SDI in cloud DC, between cloud DC.
 - How give user choice of SDX?
 - Hierarchical SDN, app specific customization on IaaS on top of HaaS
- Can we expose rich information about services while not violating customer privacy
- How can we correlate between the information between the different layers?
- How can we identify source of failures?



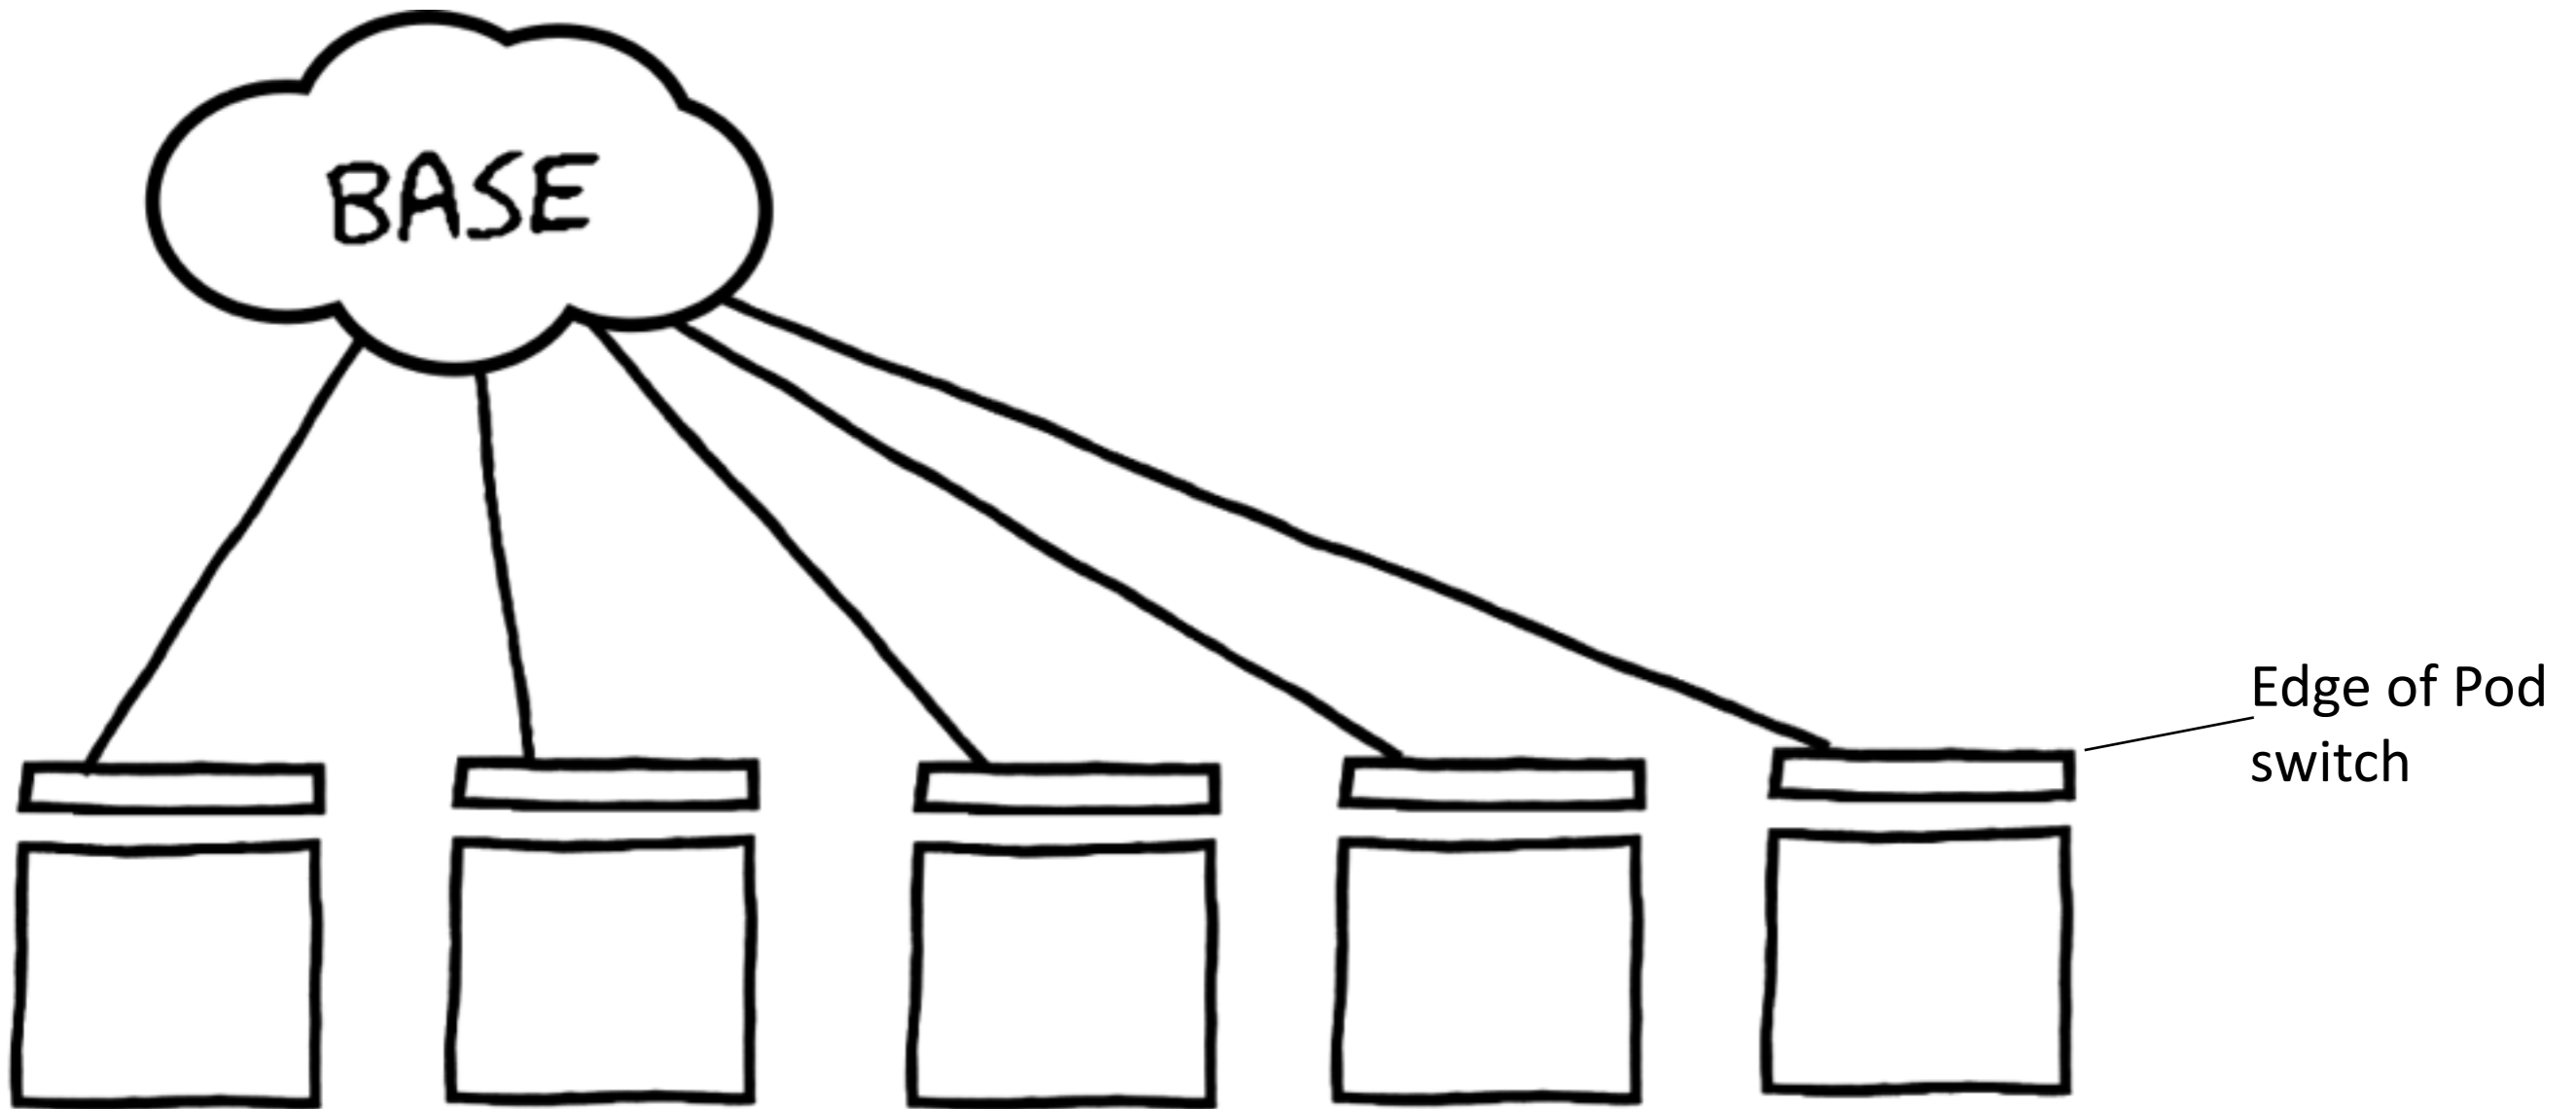
Common view:

Networking is like air conditioning, or power
Part of the infrastructure, provided by the
datacenter

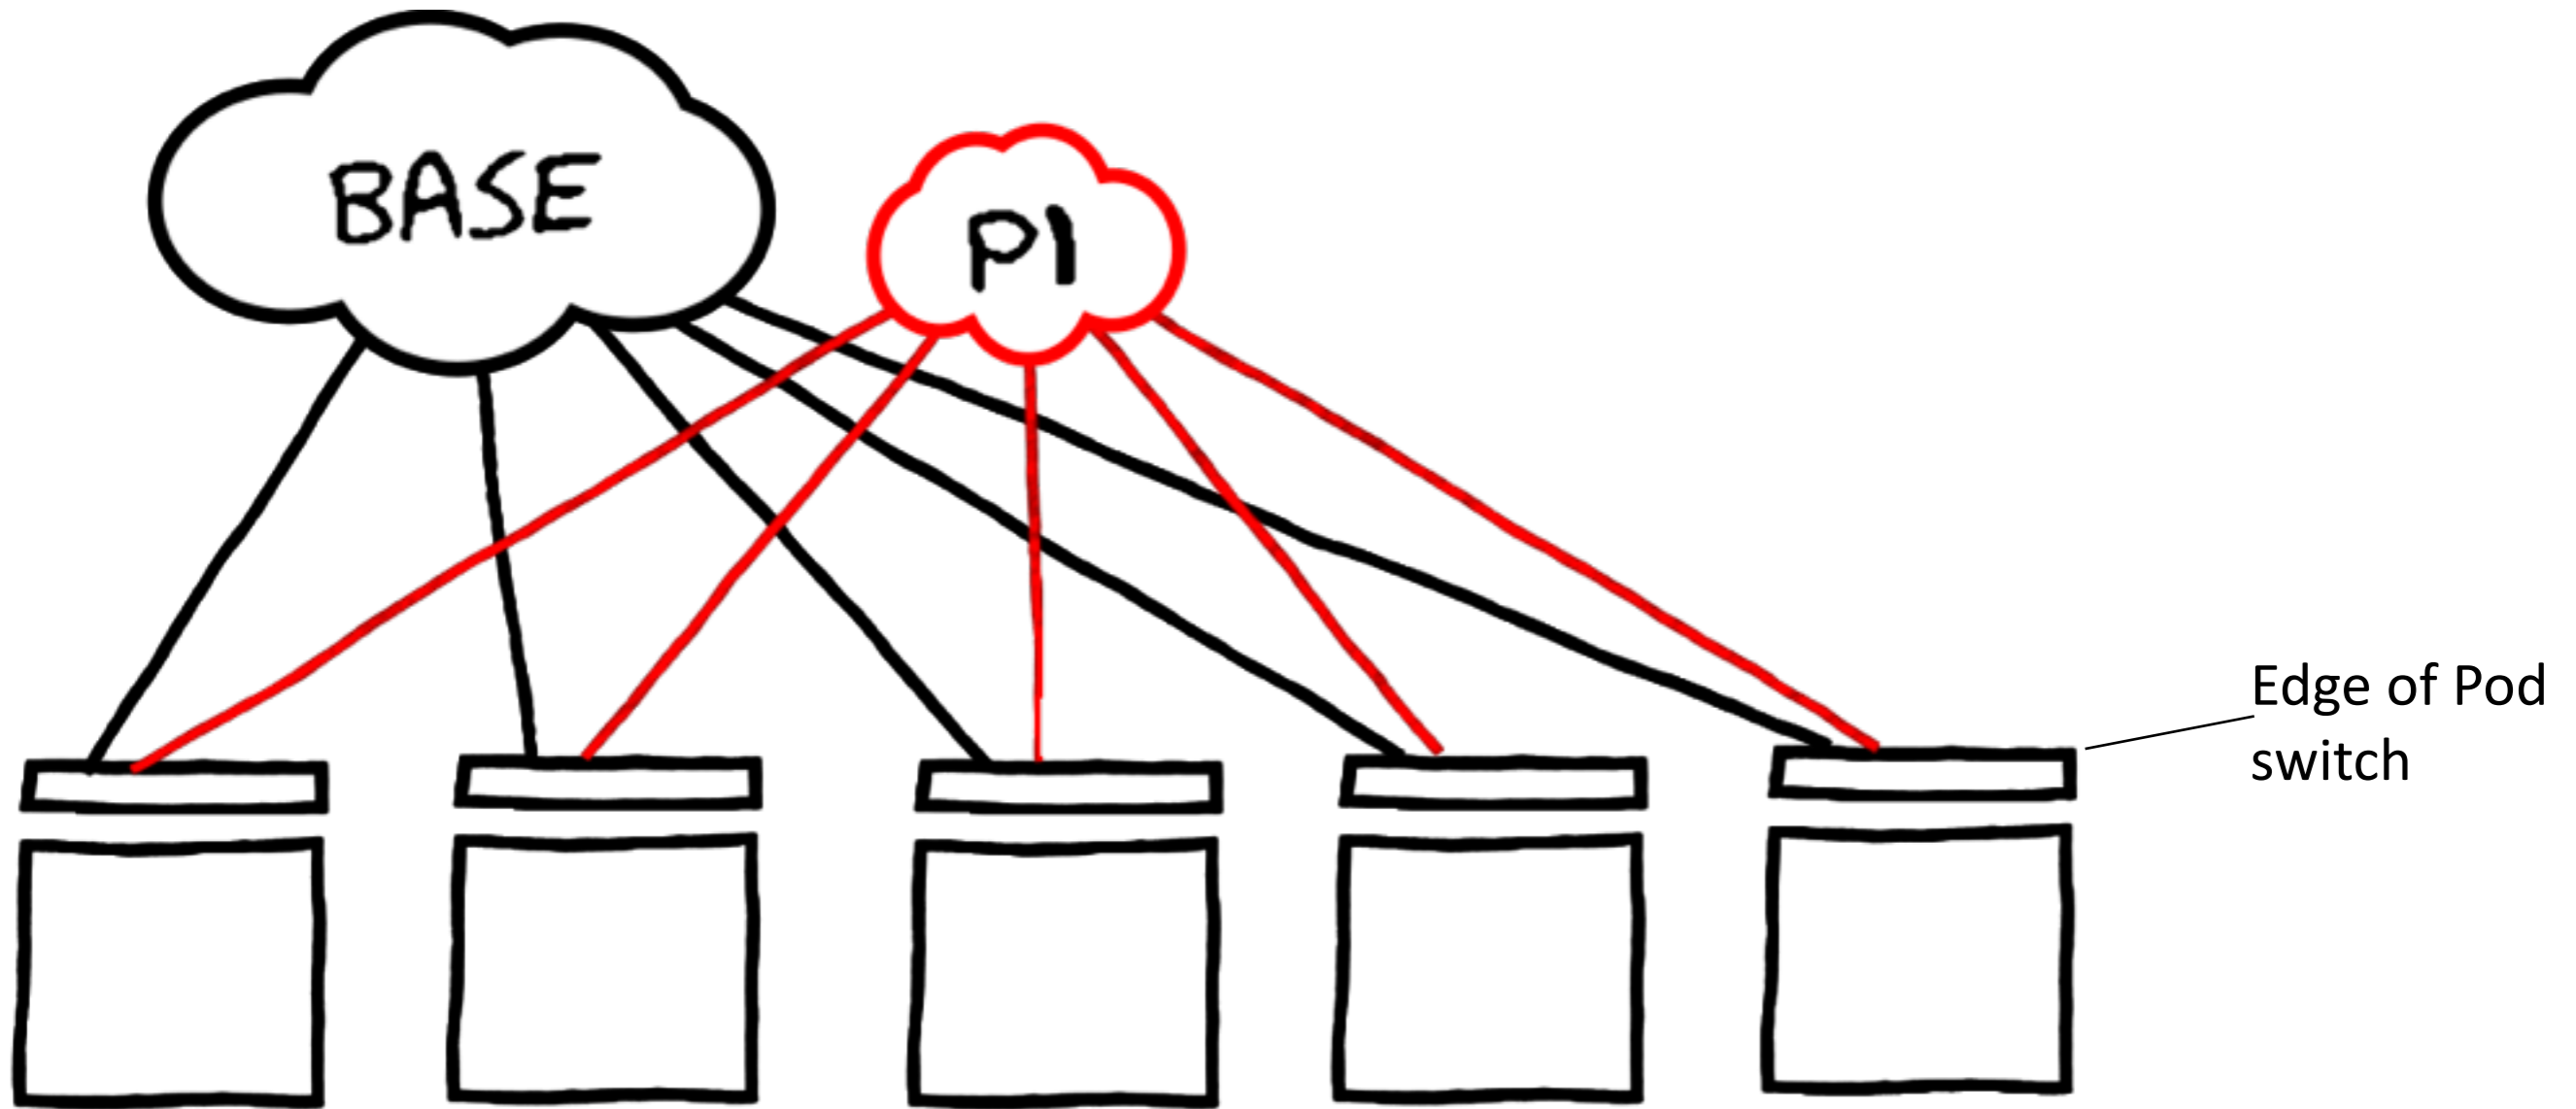
Basic Architecture




Multi-Provider Inter-Pod Network



Multi-Provider Inter-Pod Network

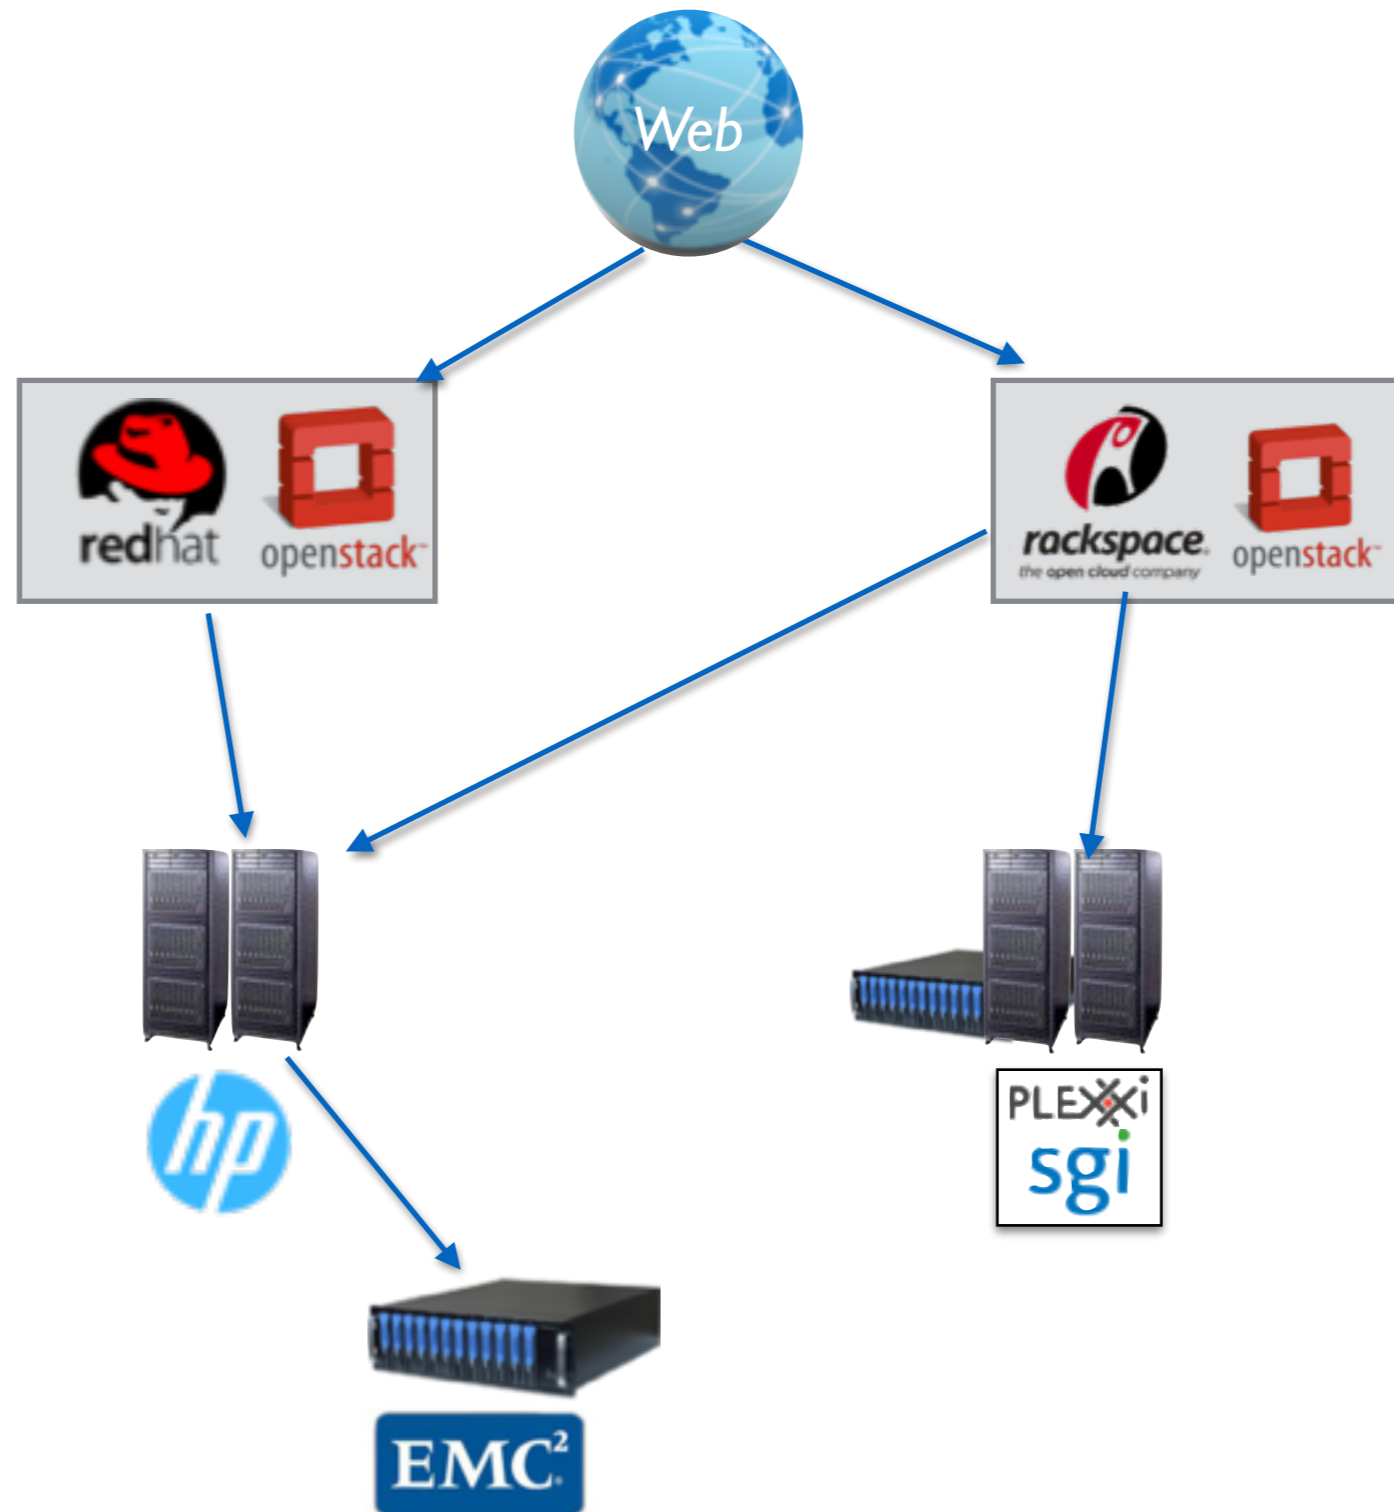







Multi-Provider Inter-Pod Network



Research challenges

- Networking:
 - Multi-provider SDN, SDX on top of SDI in cloud DC, between cloud DC.
 - How give user choice of SDX?
 - Hierarchical SDN, app specific customization on IaaS on top of HaaS
- Can we expose rich information about services while not violating customer privacy
- How can we correlate between the information between the different layers?
- How can we identify source of failures?





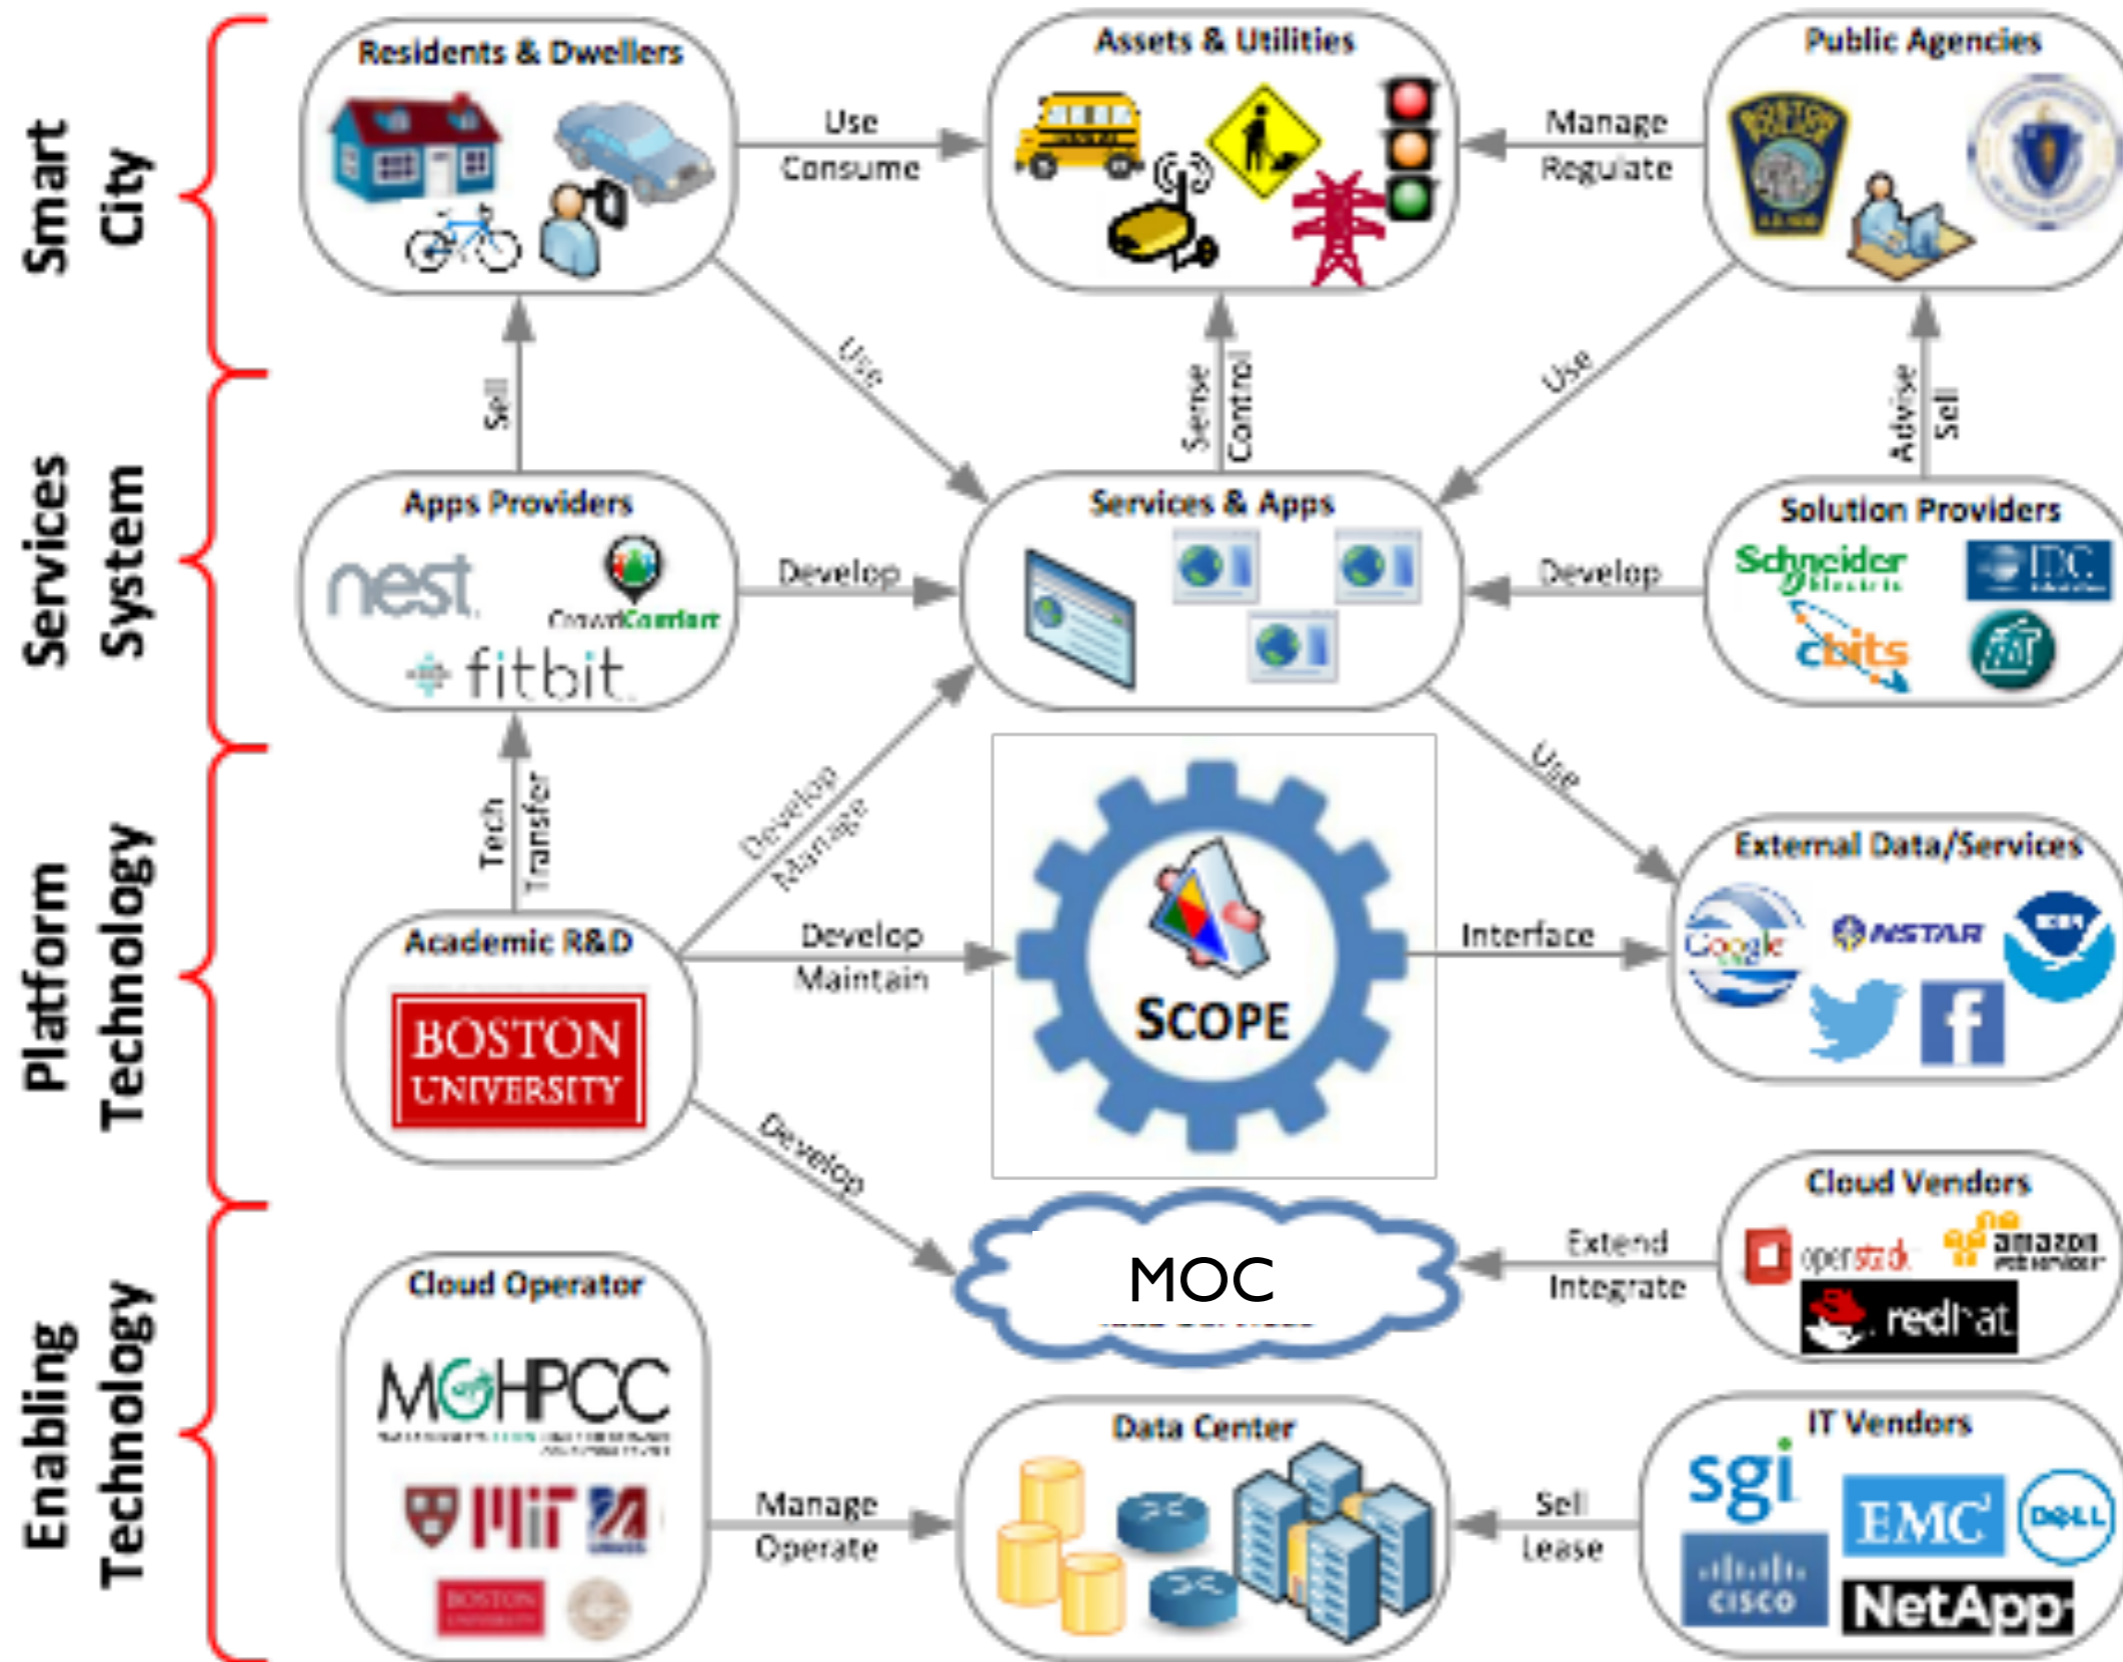
Research enabled

- New hardware infrastructure; e.g. FPGAs, new processors
- Cloud security and composability of security properties; e.g., MACS project
- Smart cities
- Analysis of cloud internal information (logs, metrics) for security, for optimization...
- Highly elastic environments; e.g., 1000 servers for a minute:

Research enabled

- New hardware infrastructure; e.g. FPGAs, new processors
- Cloud security and composability of security properties; e.g., MACS project
- **Smart cities: Azer Bestavros**
- Analysis of cloud internal information (logs, metrics) for security, for optimization...
- Highly elastic environments; e.g., 1000 servers for a minute:

Example: Smart cities




Research enabled


- New hardware infrastructure; e.g. FPGAs, new processors
- Cloud security and composability of security properties; e.g., MACS project
- Smart cities
- Analysis of cloud internal information (logs, metrics) for security, for optimization...: Alina Oprea
- Highly elastic environments; e.g., 1000 servers for a minute:

MOC


VM




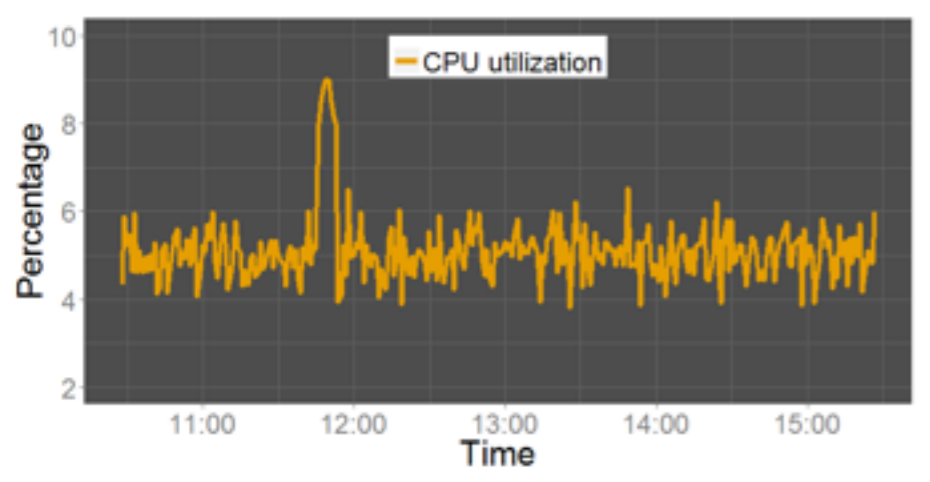
VM



VM



MOC




VM








VM



VM




MOC



Additional user protection

MOC

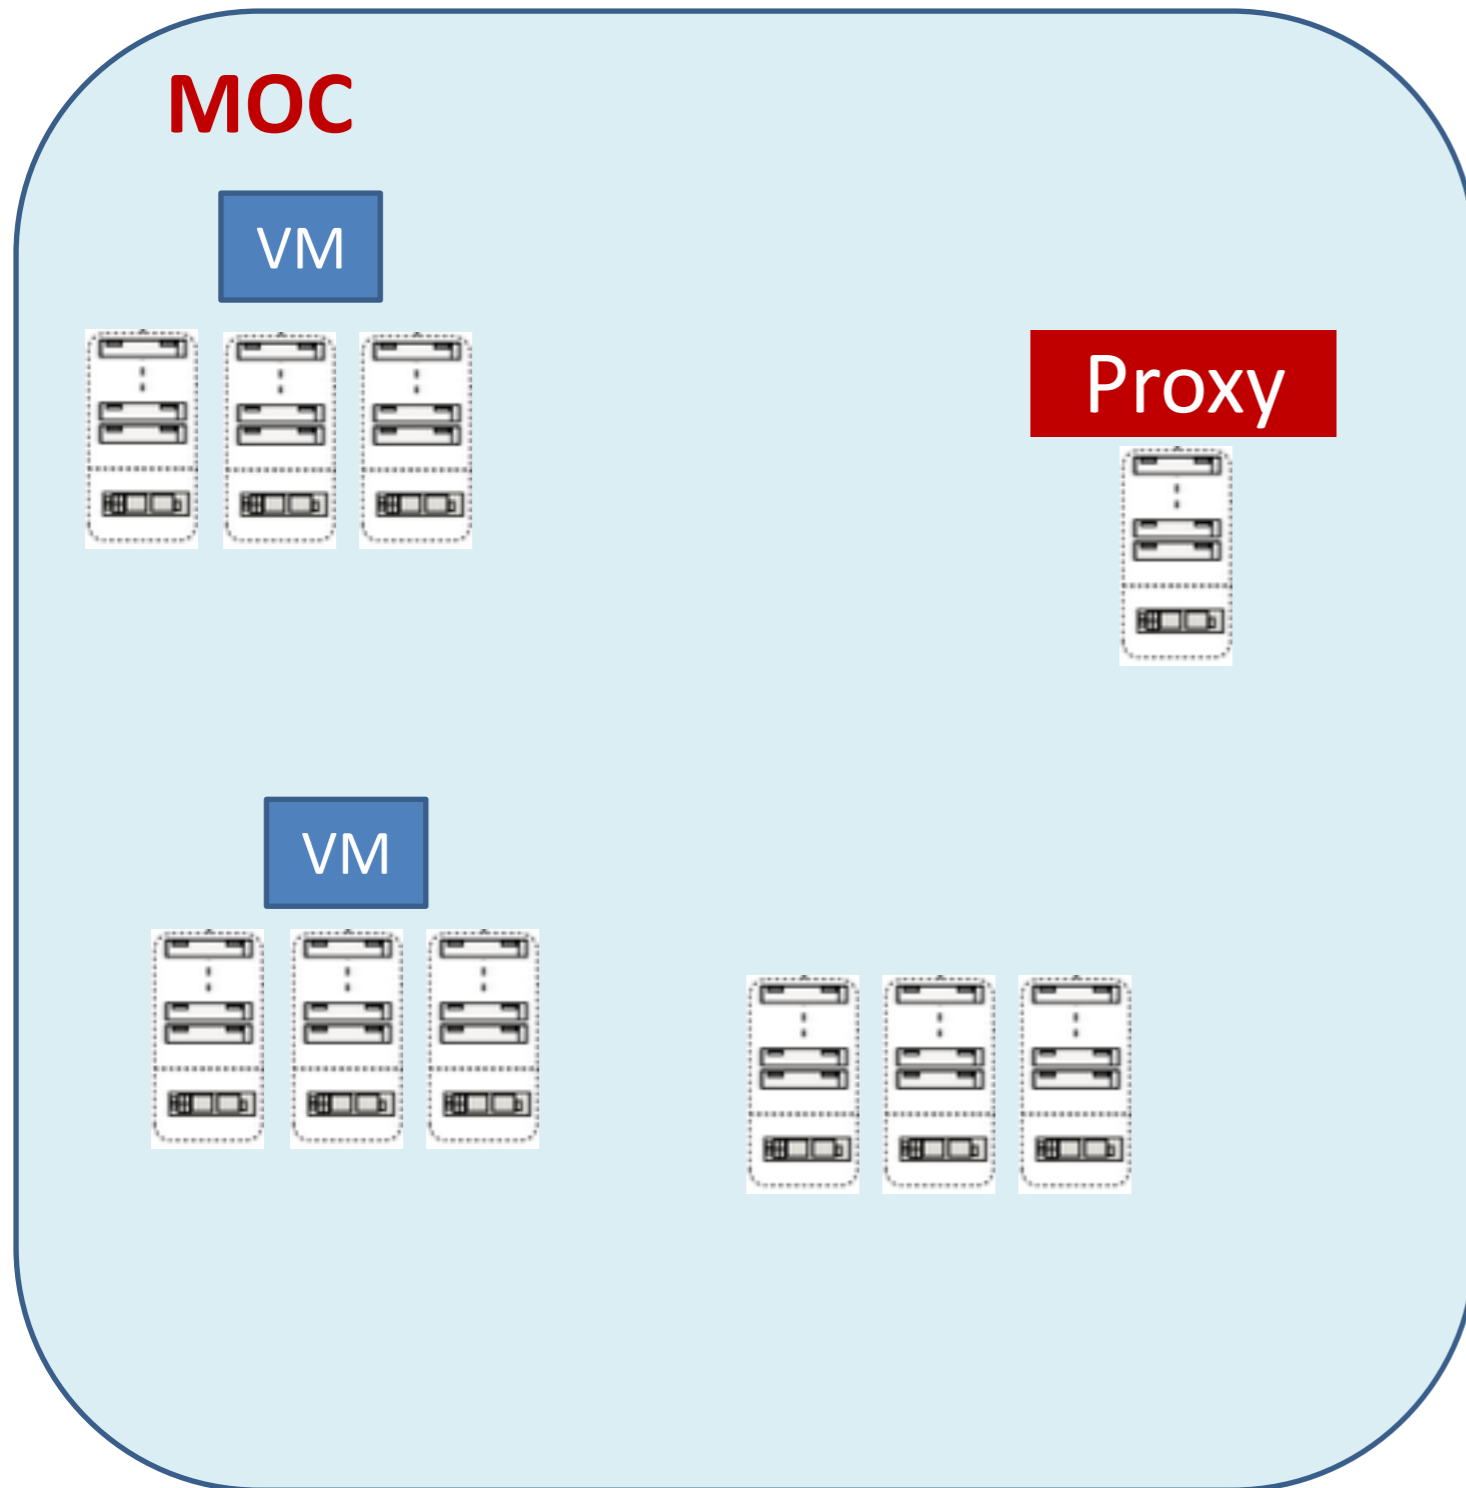
VM



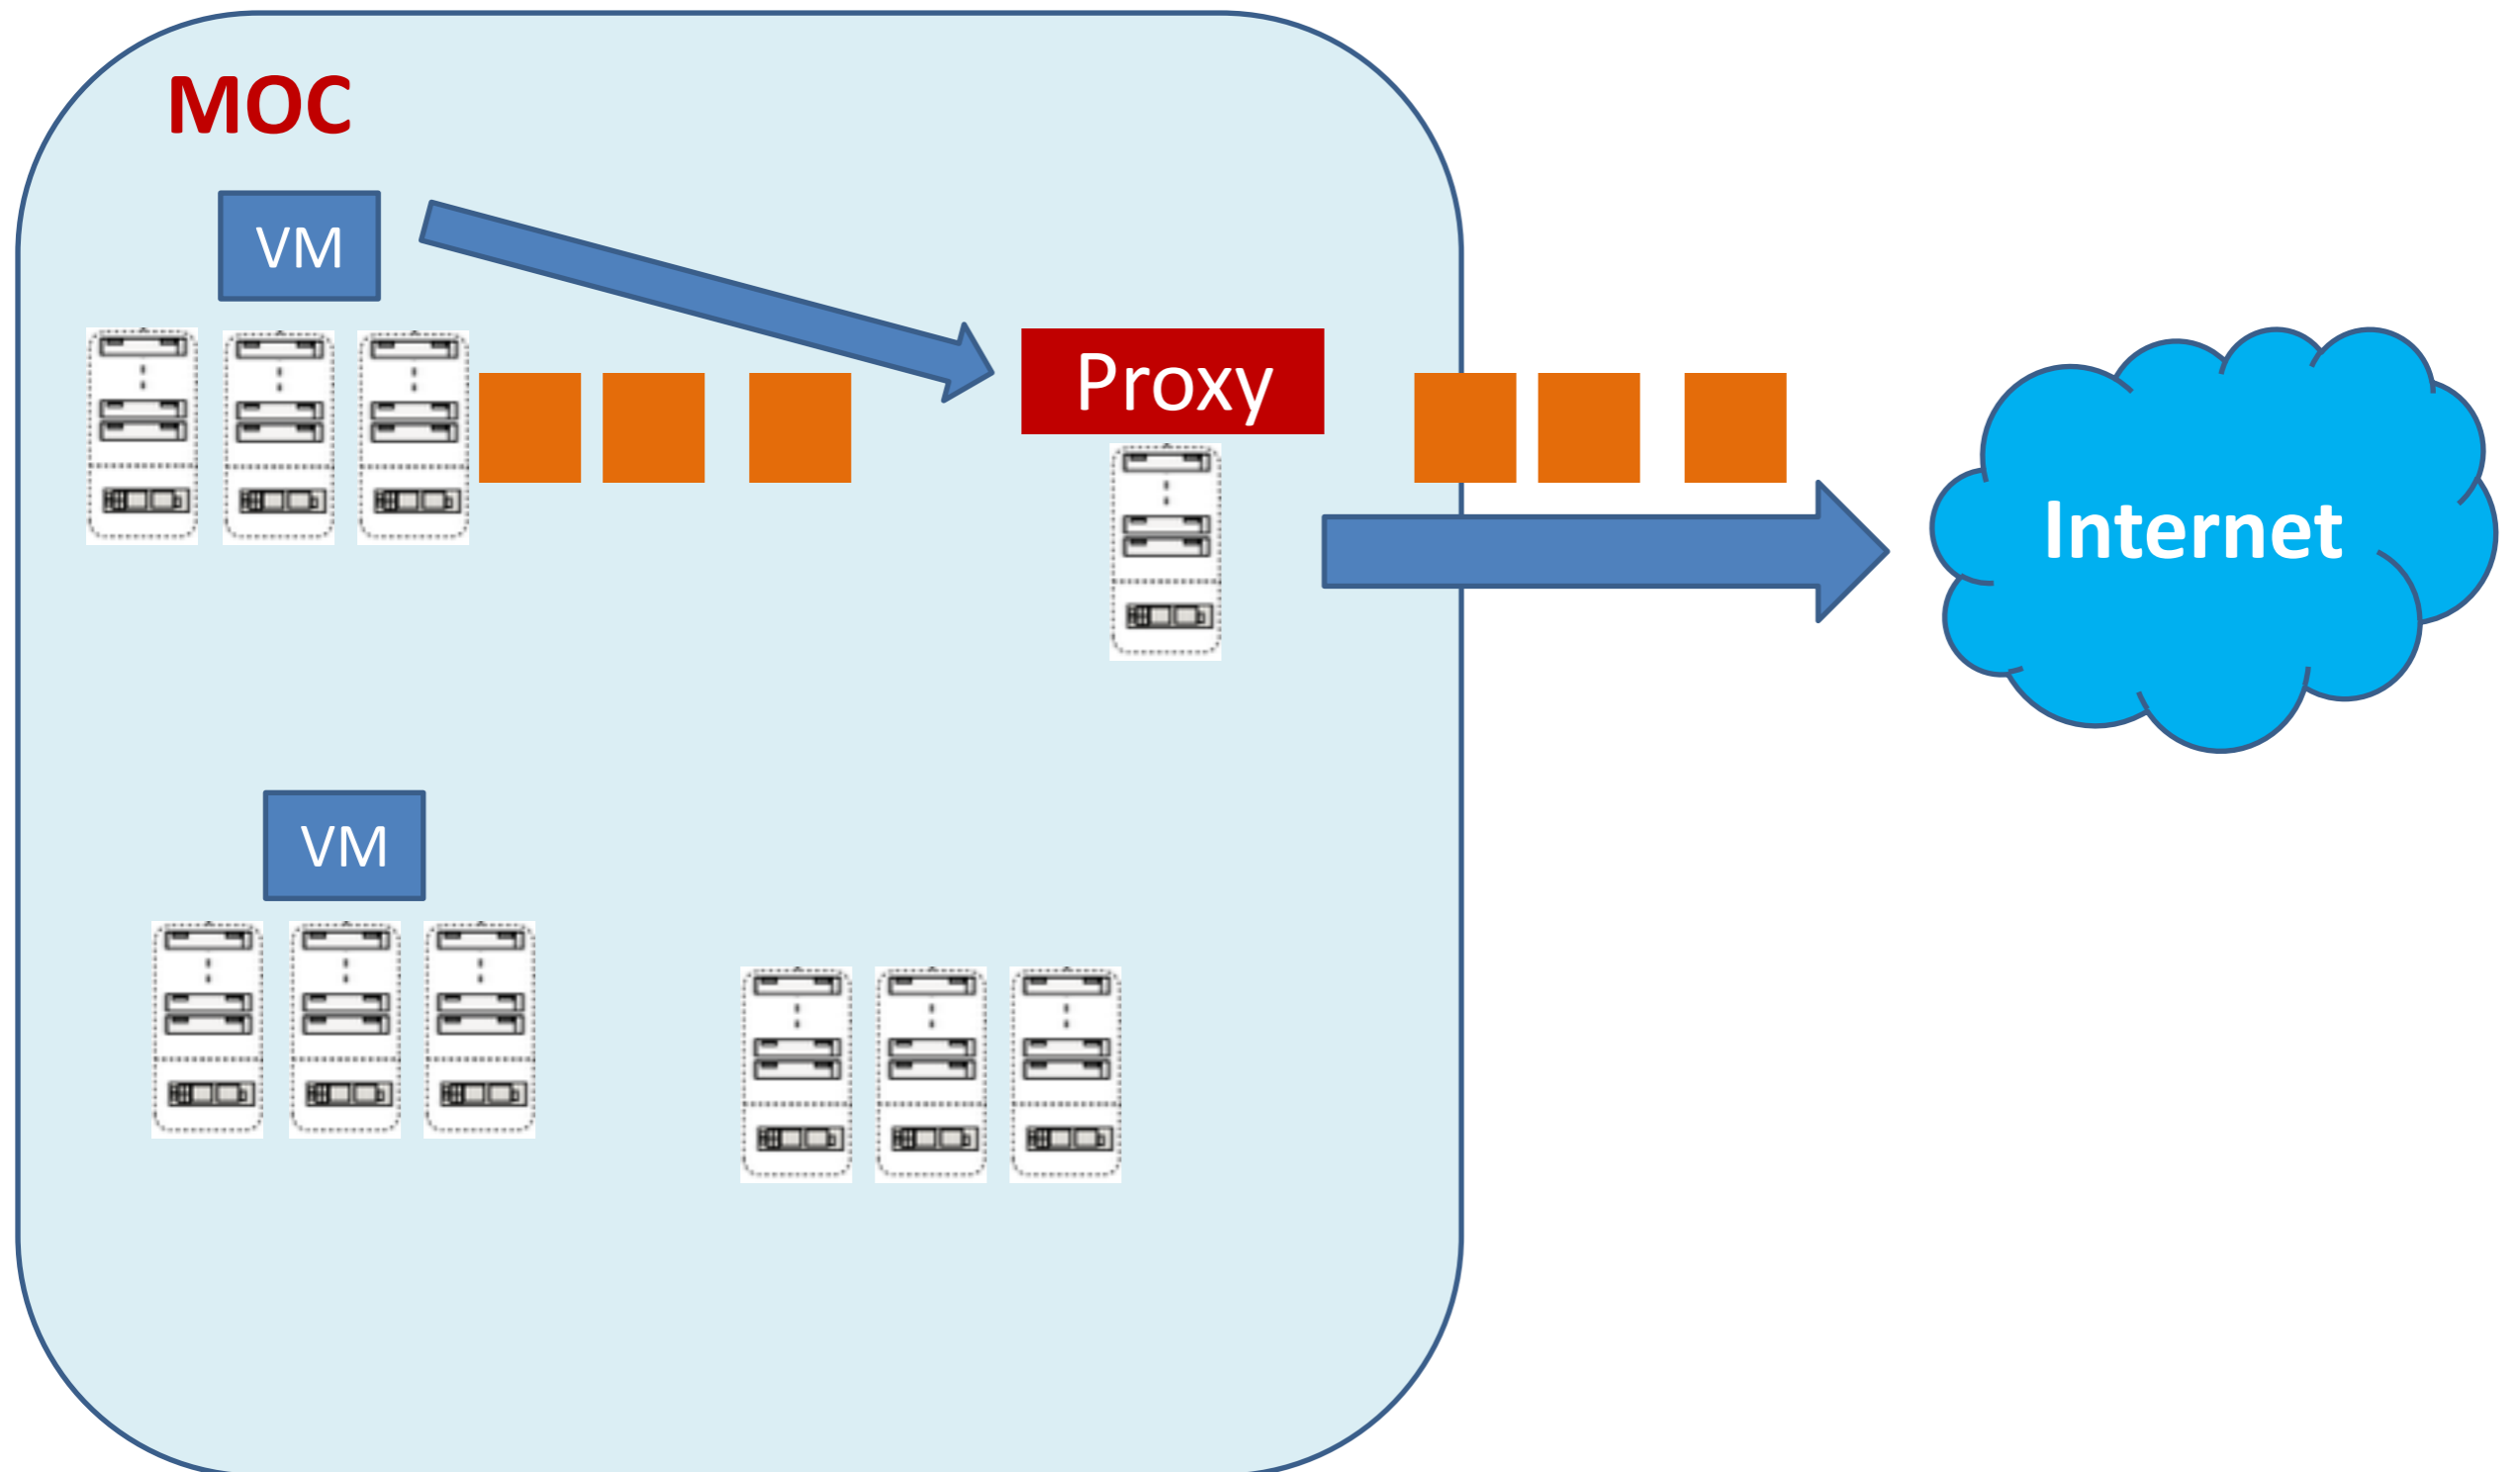
VM



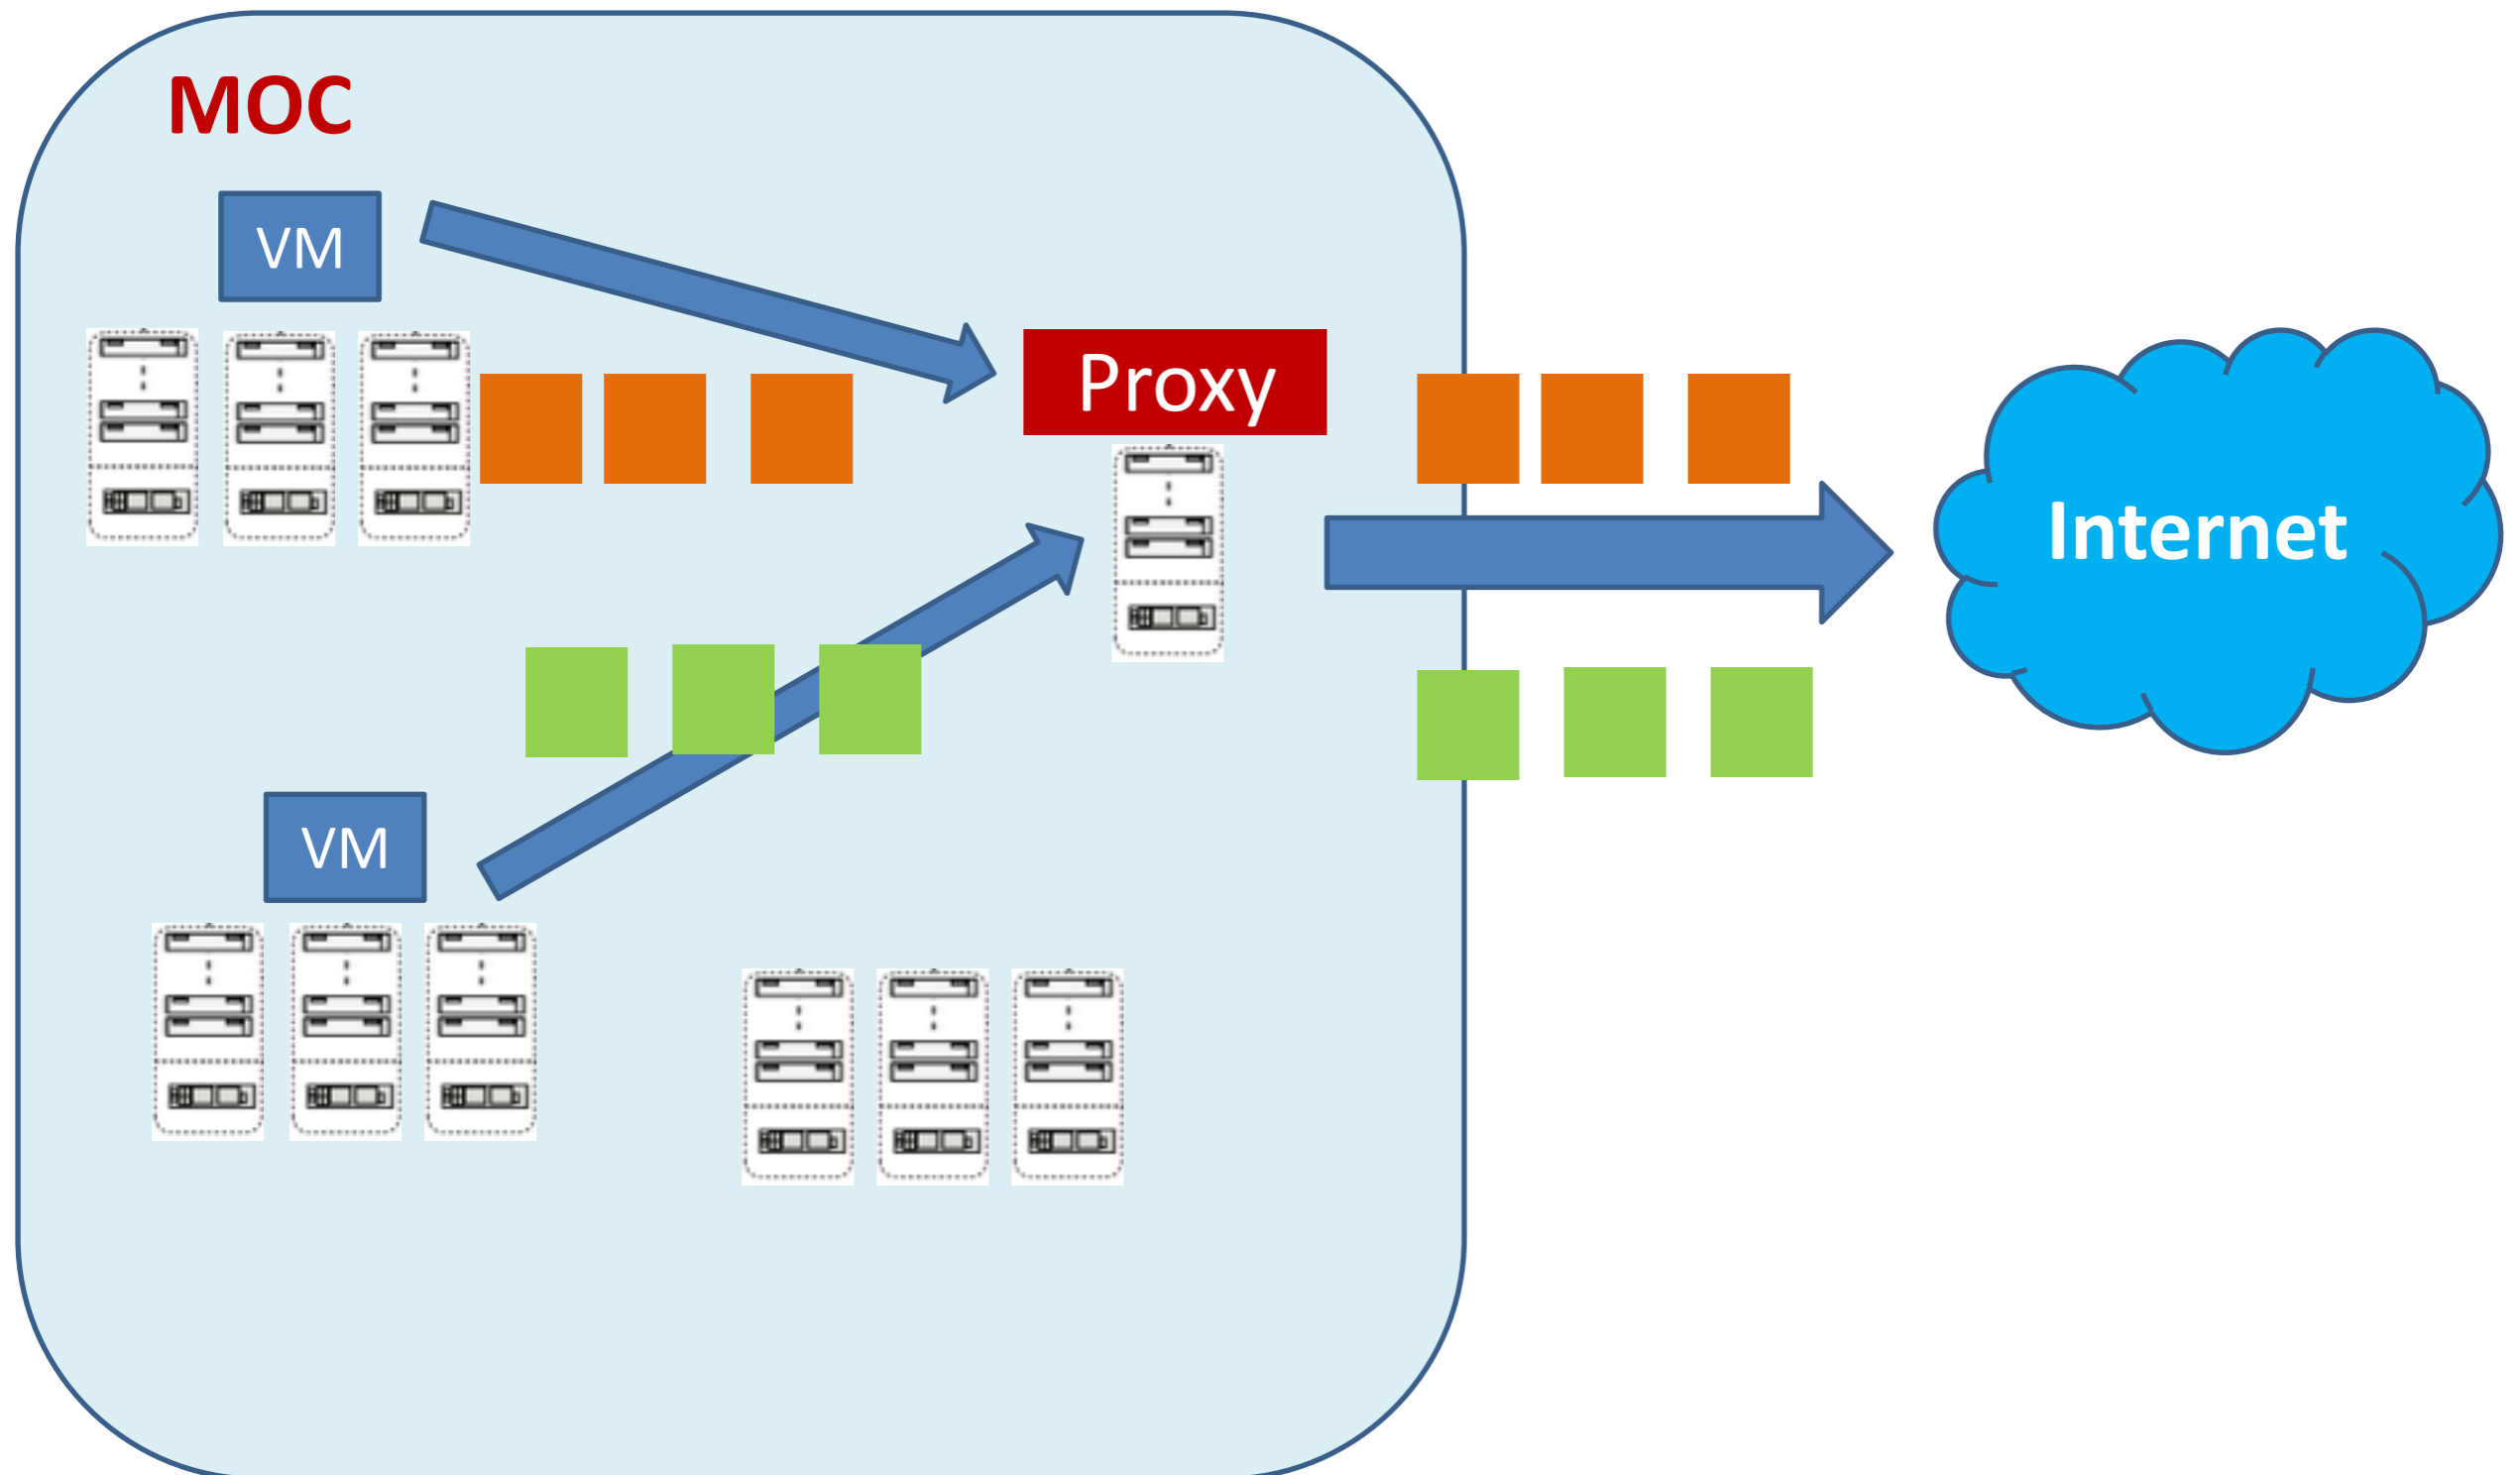
Additional user protection



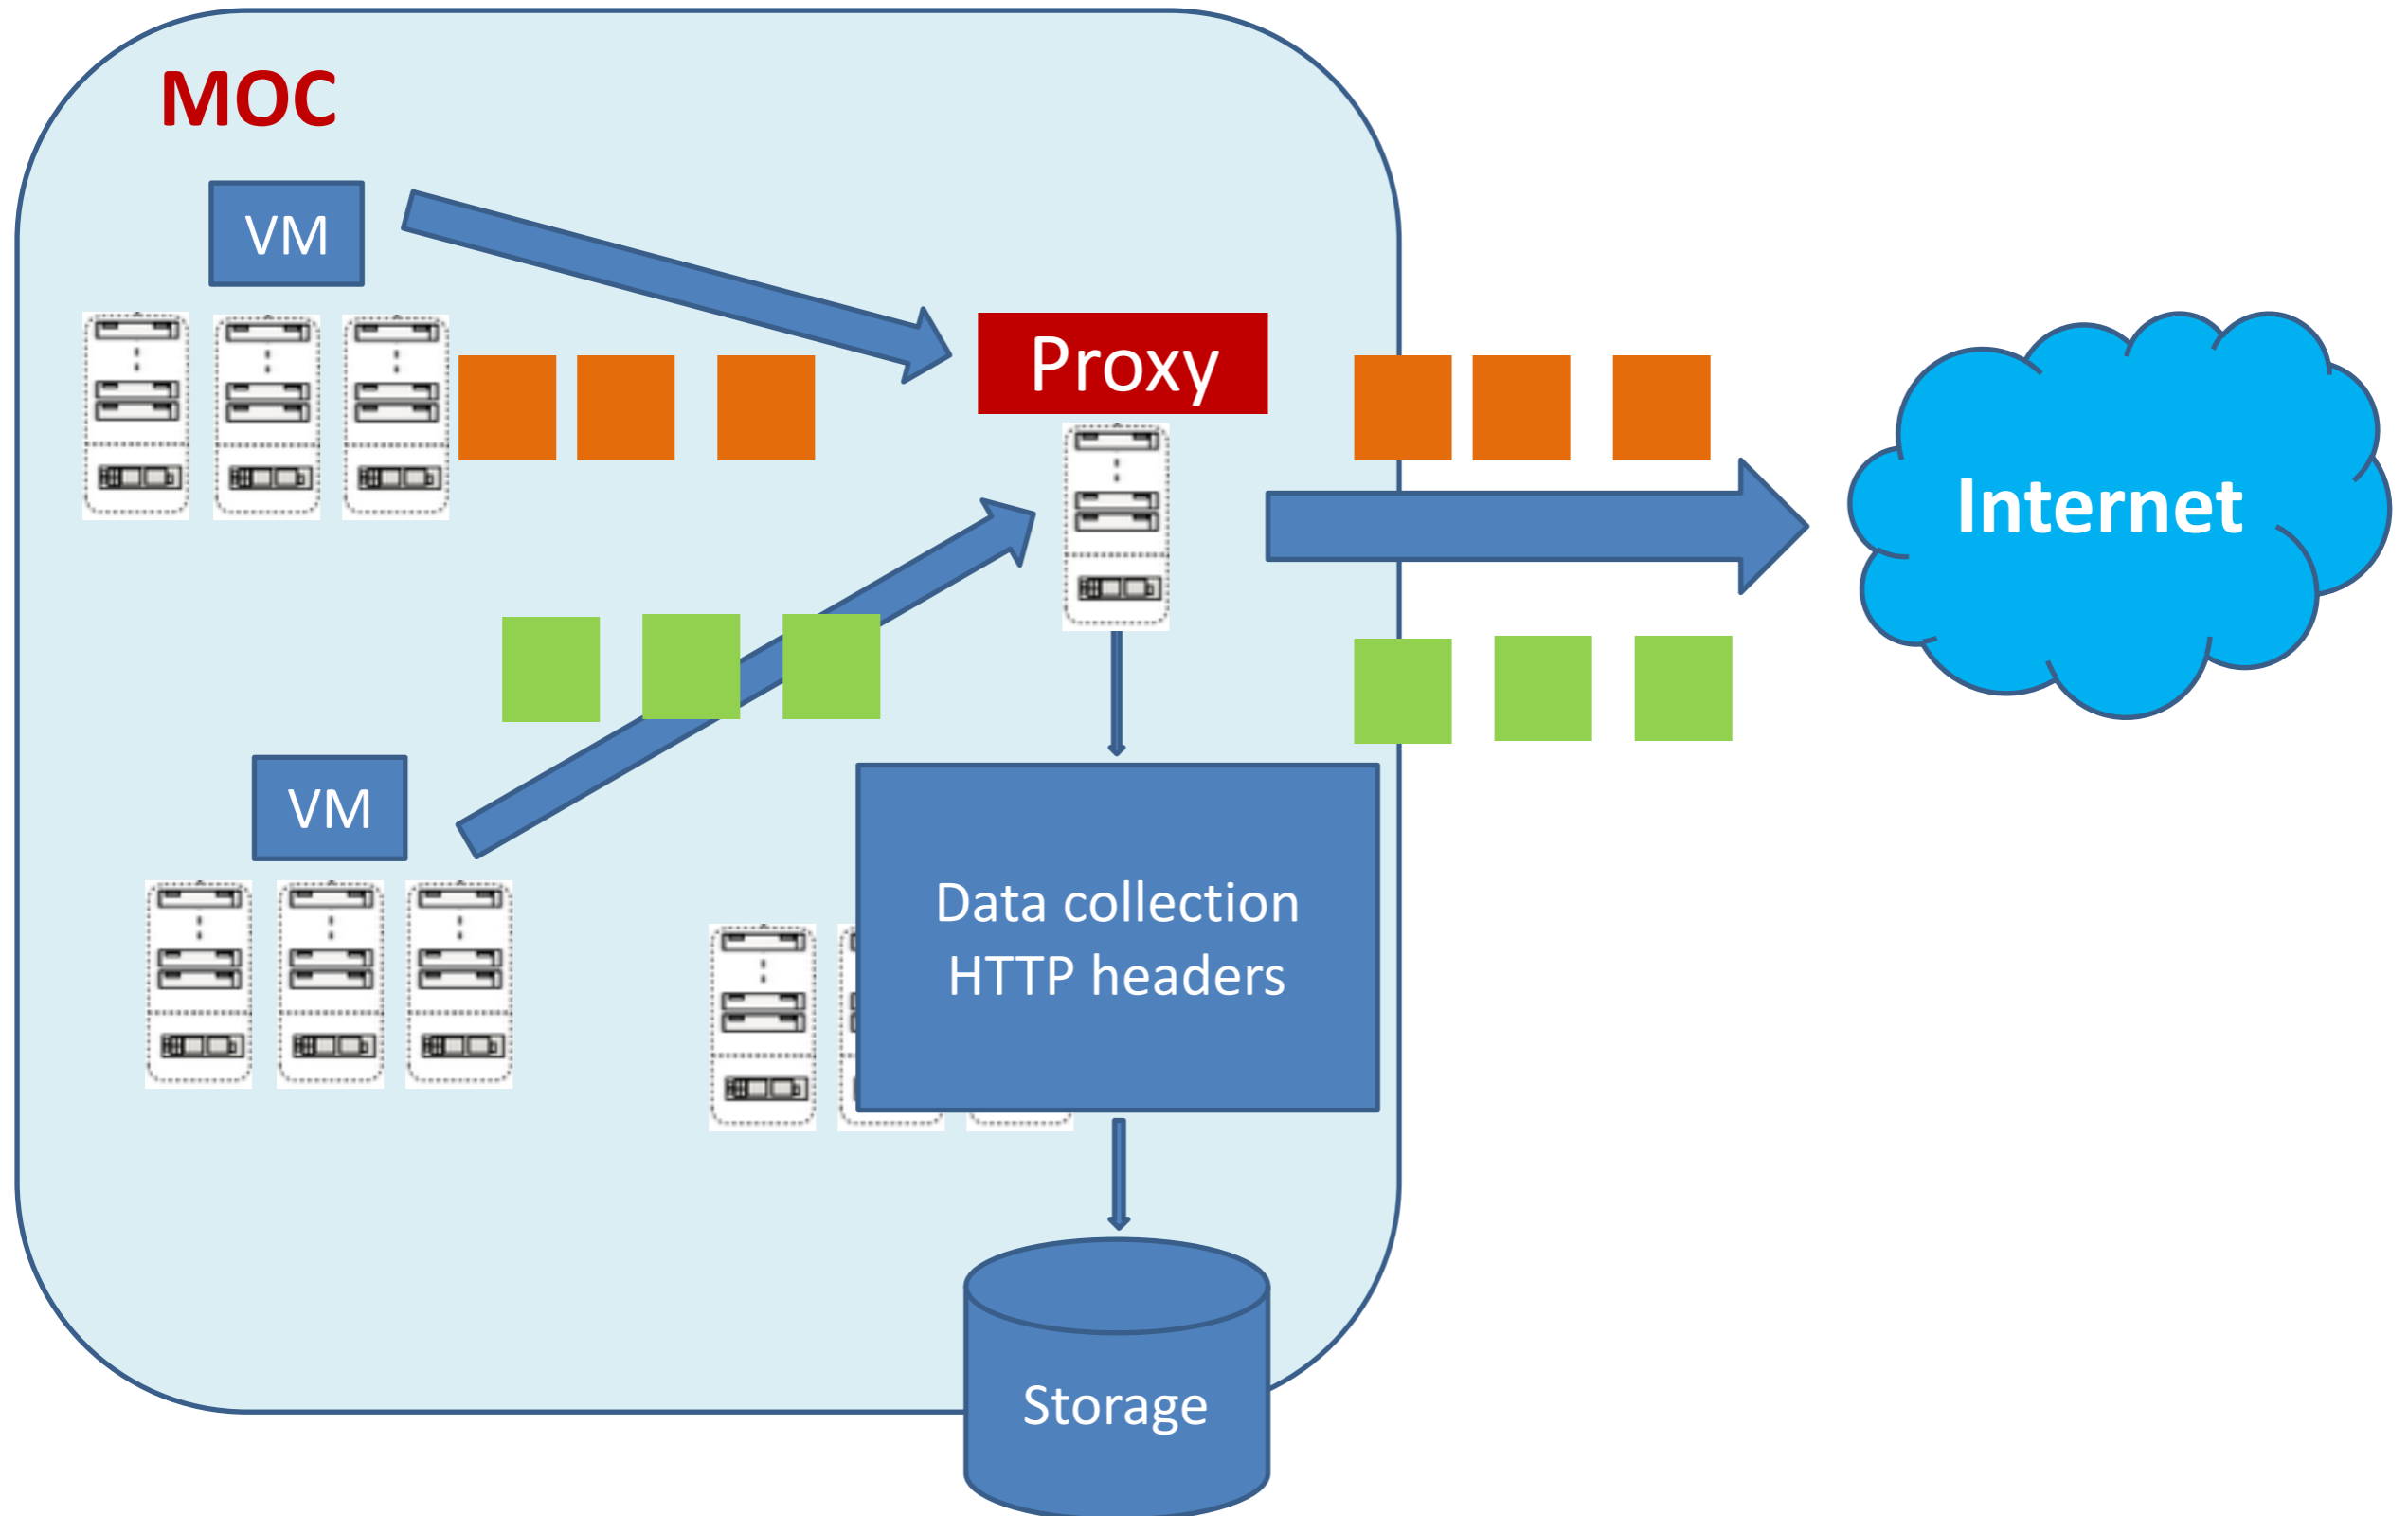
Additional user protection



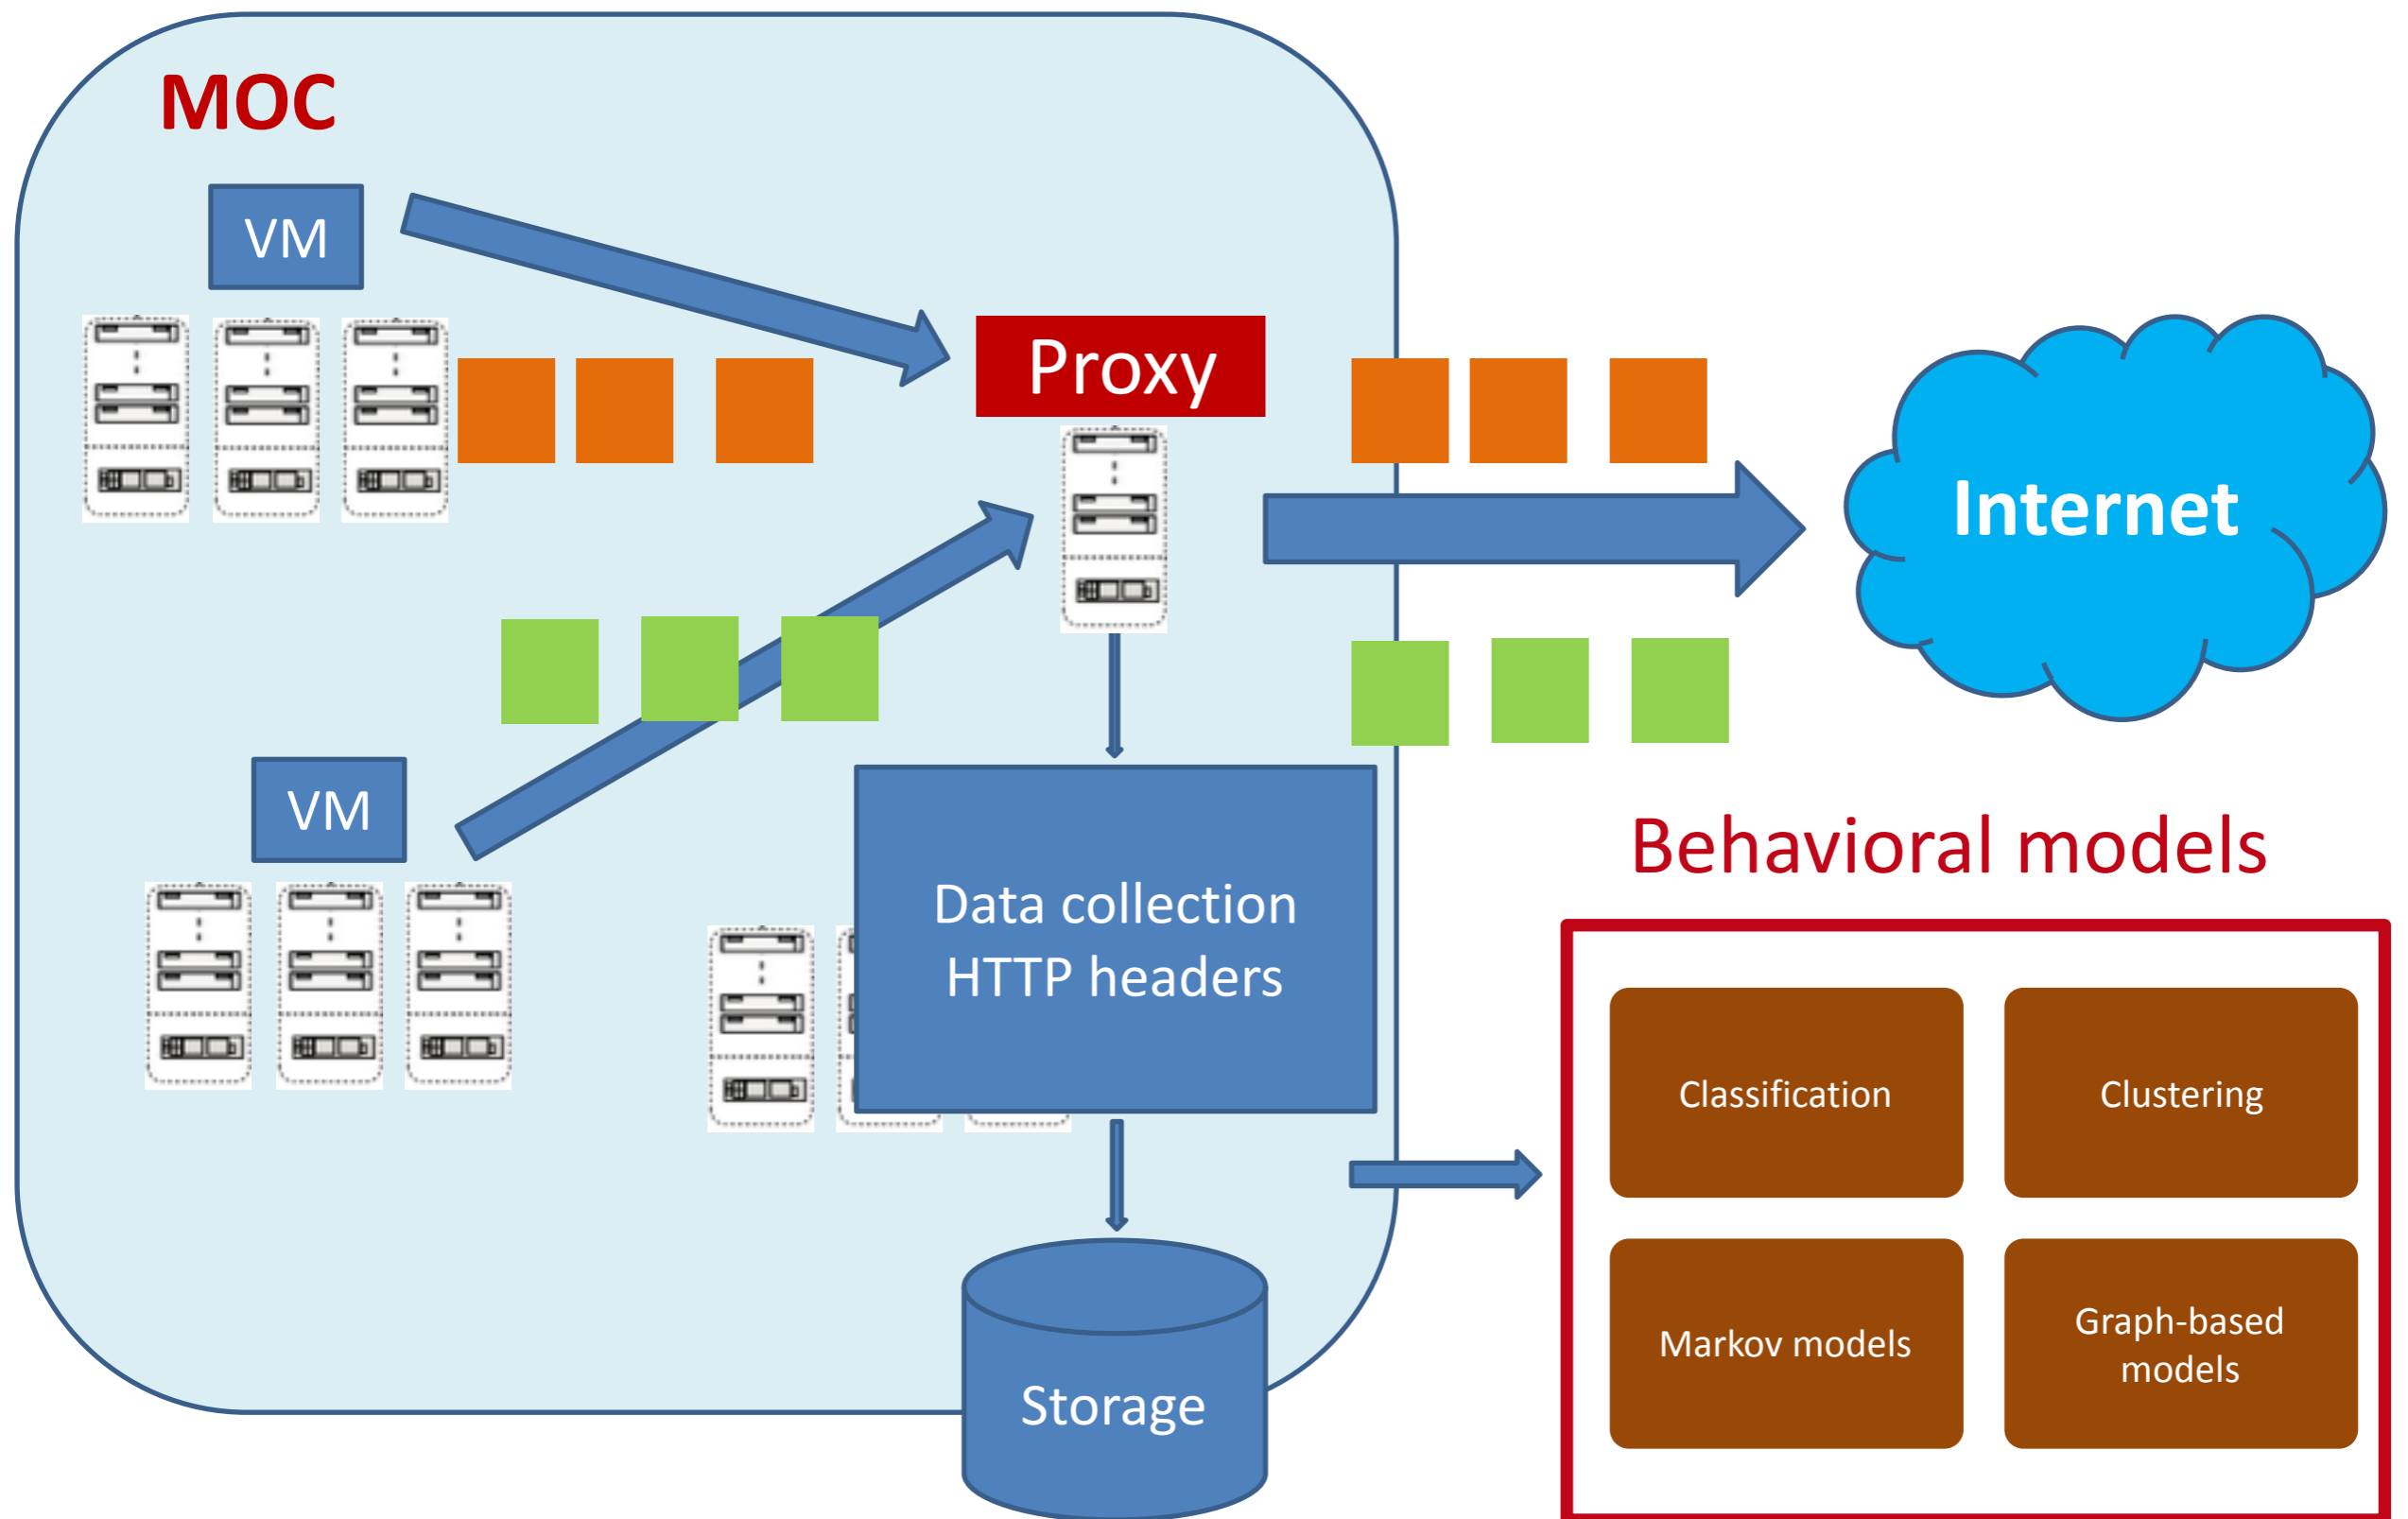
Additional user protection



Additional user protection



Additional user protection




Research enabled

- New hardware infrastructure; e.g. FPGAs, new processors
- Cloud security and composability of security properties; e.g., MACS project
- Smart cities
- Analysis of cloud internal information (logs, metrics) for security, for optimization...
- Highly elastic environments; e.g., 1000 servers for a minute: Jonathan Appavoo

Example 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



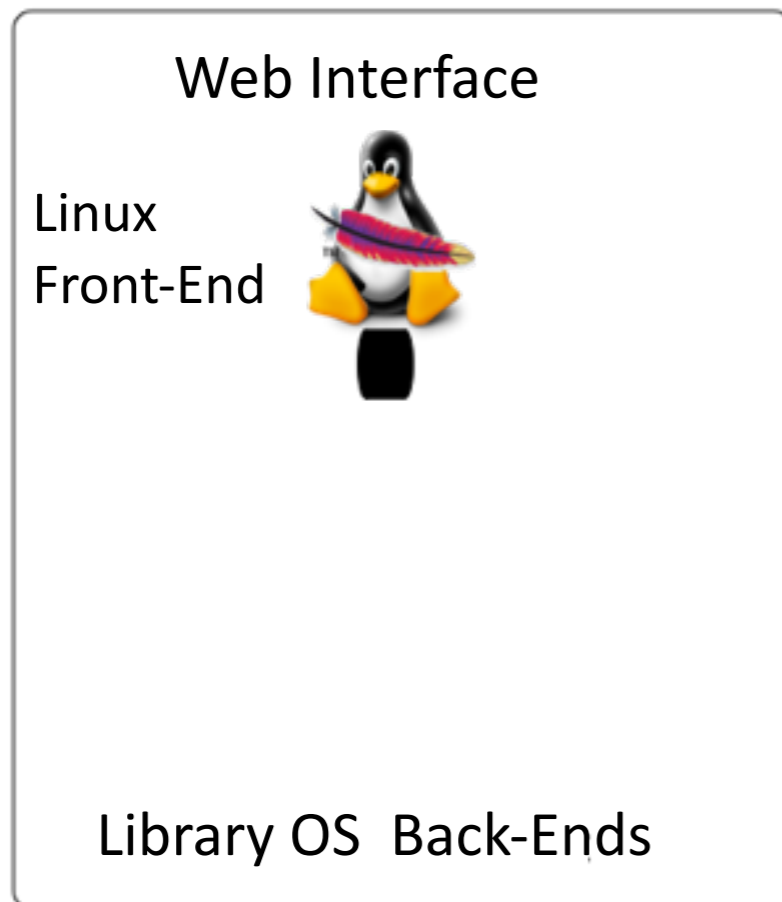
Infrastructure as Elastic Resource Pool

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

Example 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




Infrastructure as Elastic Resource Pool

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




Fetal Image Reconstruction




Fetal Image Reconstruction




Fetal Image Reconstruction




resized,
cropped
96x96
50 slices




Fetal Image Reconstruction





24hrs
→




resized,
cropped
96x96
50 slices




Fetal Image Reconstruction



24hrs
→

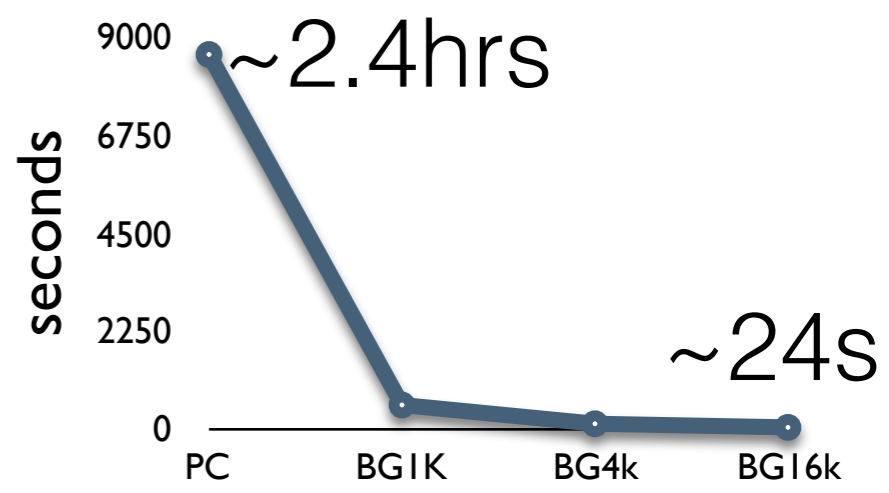



resized,
cropped
96x96
50 slices



Fetal Image Reconstruction

synthetic
1024x1024
200 slices



Concluding remarks

- OCX model could become important component of clouds
- Major research challenge
- New research opportunities in this world
- OCX model important to enable research co-exists with commercial efforts with:
 - real data, real users, real scale