

# Context-aware cloud computing for HEP applications

Randall Sobie

Institute of Particle Physics (Canada)

University of Victoria

# Context-aware systems in mobile communication



How can we use the information (context) of a user to their benefit?

Location  
Direction  
Time  
Health  
Social

# Context-aware systems in mobile communication



The information can meet on-demand request or be proactively sent to the user

Directions  
Food and lodging  
Financial  
Commercial  
Health

## Entity

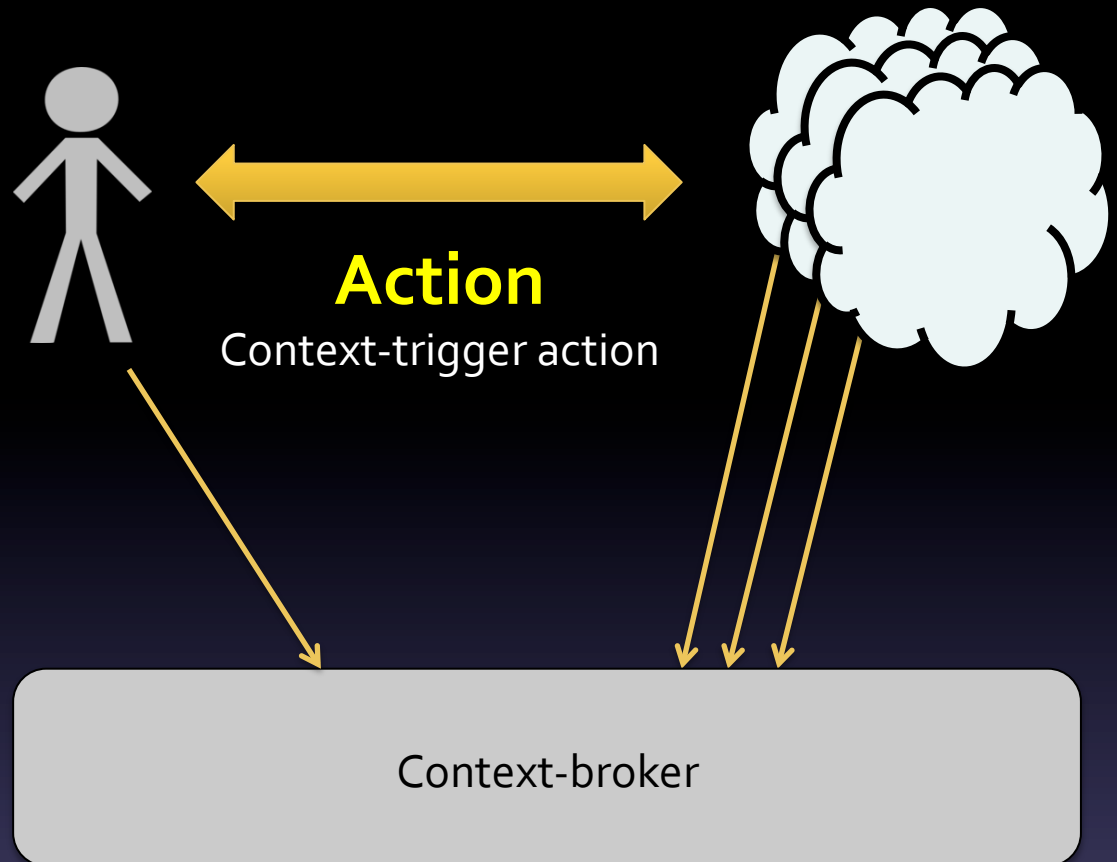
A person, place, or object that is relevant to their interaction

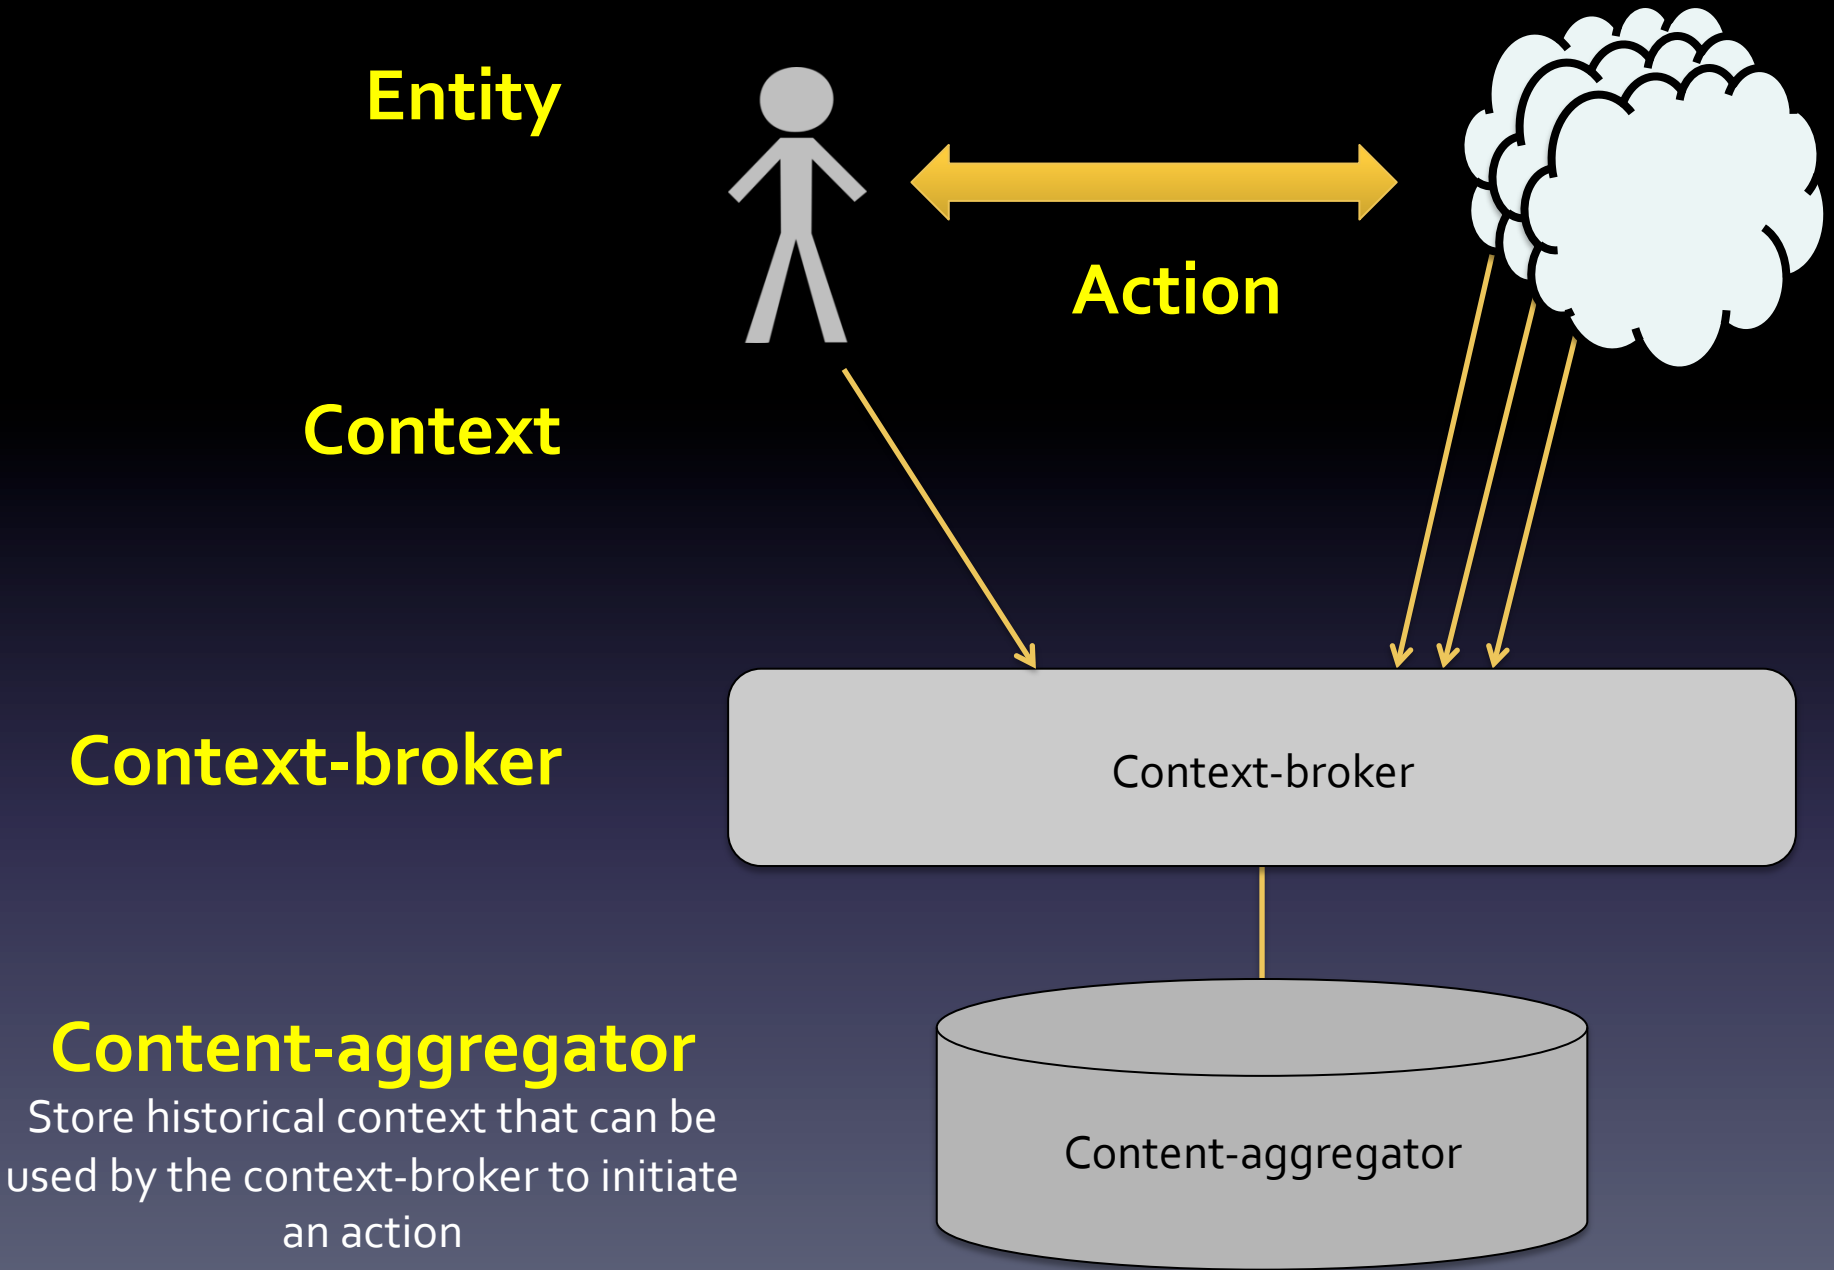
## Context

Information used to characterize the situation of an entity

## Context-broker

Combines real-time context



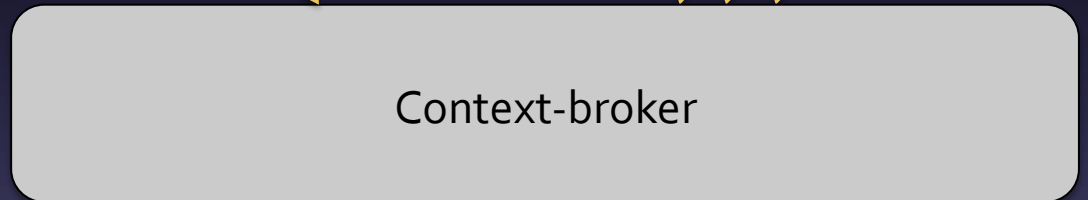


**Entity**  
ISGC Physicist



**Context**  
Hungry  
Dinner time  
Central Taipei


**Context-broker**



**Entity**  
ISGC Physicist



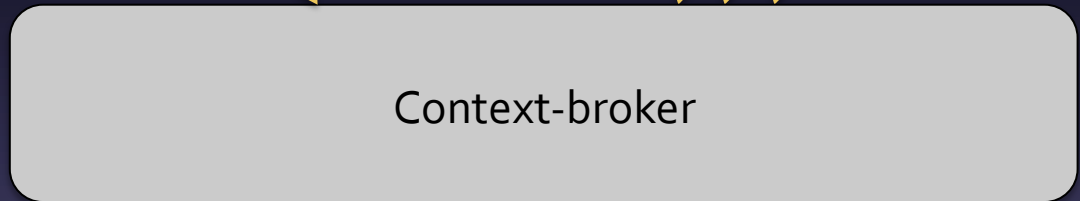
**Action**  
List of fast food



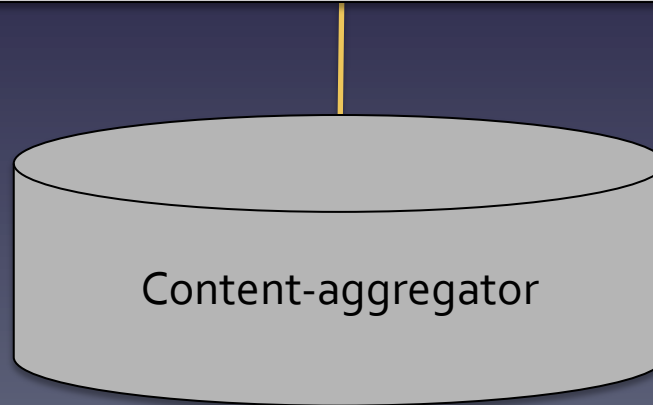
**Context**  
Hungry  
Dinner time  
Central Taipei



**Context-broker**



**Content-aggregator**  
Graduate student



# Context-aware cloud computing system

## Entities

Workload systems  
Clouds  
VMs  
VM-provisioner  
Batch job scheduler



**Can a context-aware design make it easier, more efficient to run HEP applications?**

Utilize dynamic and opportunistic resources  
Automatically self-configure clouds and virtual machines  
Locate optimal software caches and data repositories  
Fault-detection and error correction  
Intelligent design and self-learning



# Cloud computing in HEP

Cloud computing in HEP is typically providing 5-20% of the processing of current projects

## Dedicated

Virtual cluster



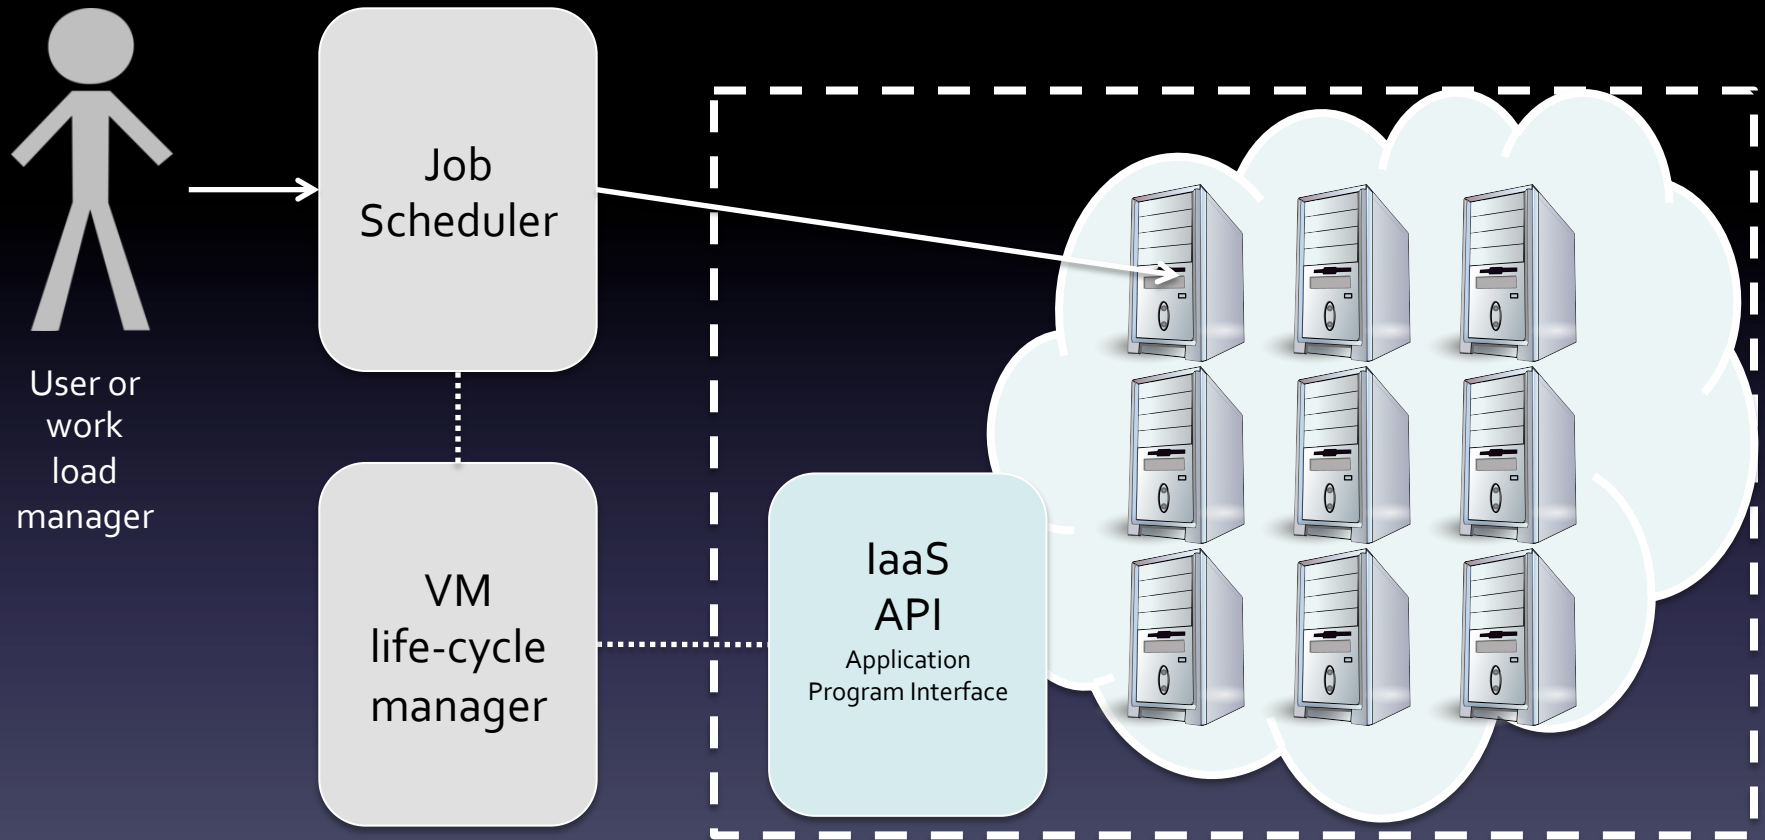
## Opportunistic



**“Dedicated” clouds**  
(Owned by HEP)

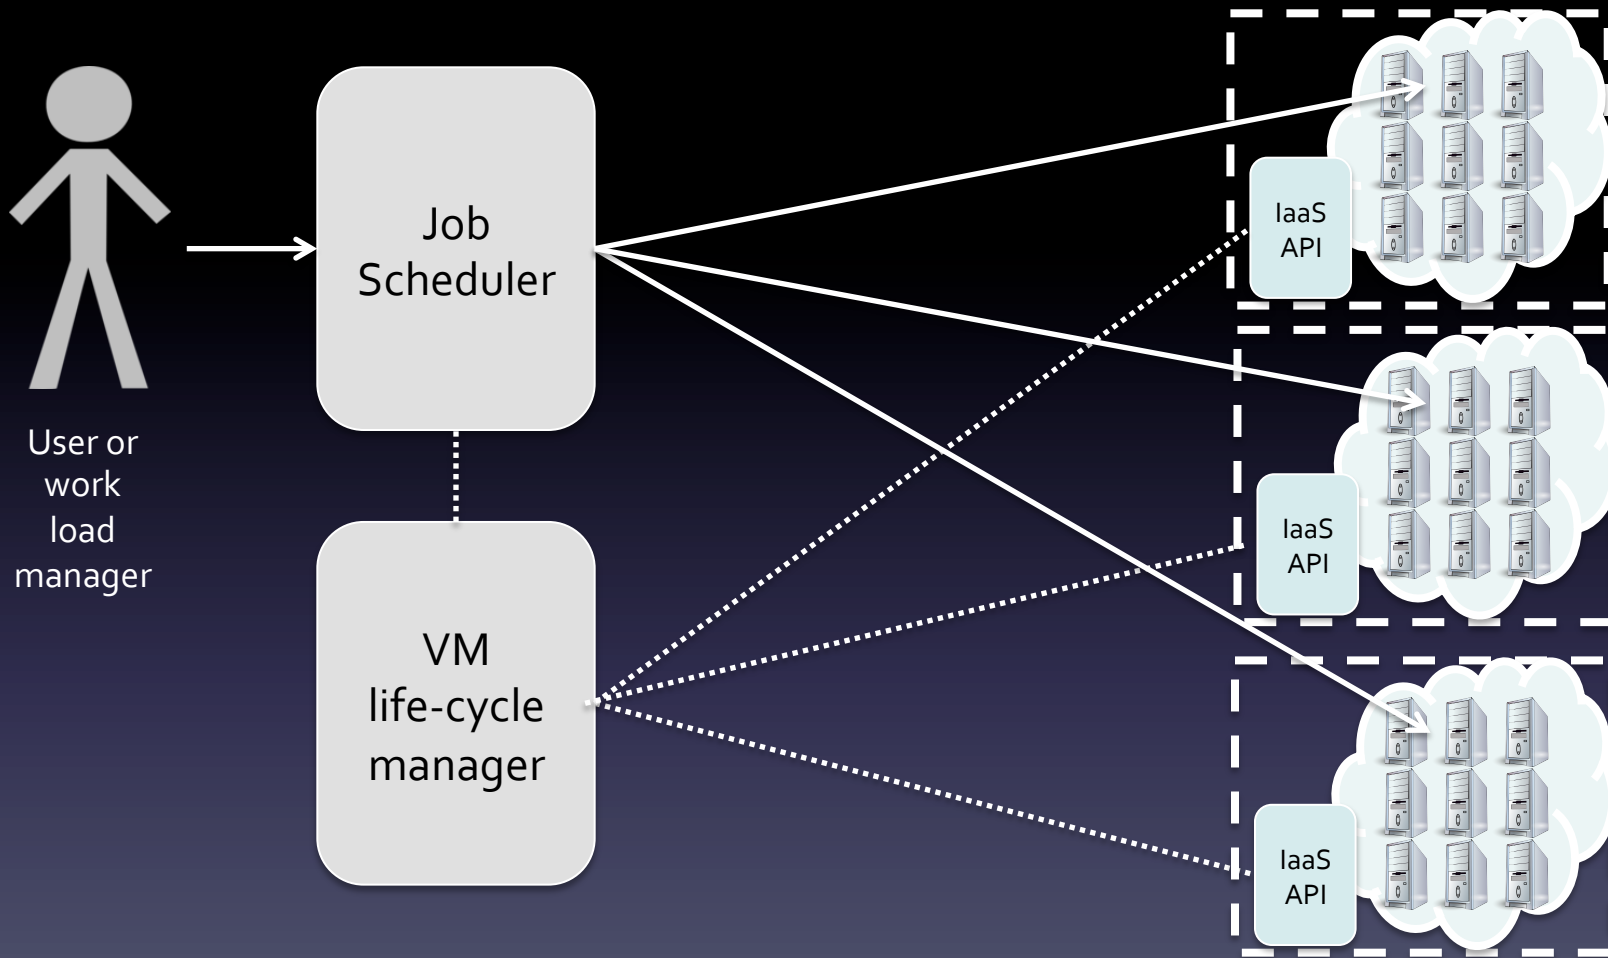
**“Opportunistic” clouds**  
(private and commercial)

# HEP use of clouds for batch workloads




A "VM life-cycle manager" manages VMs based on the job queue  
Cloud has no application-specific requirements

# Multiple remote clouds

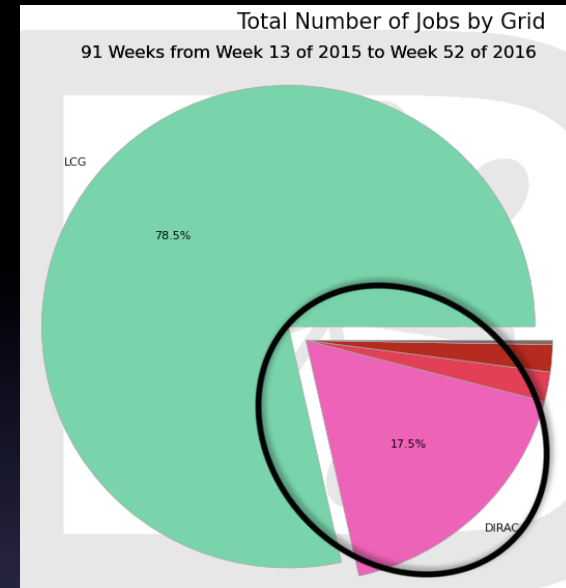


*Workload manager has no knowledge of the resources*

# Distributed batch cloud system



Dedicated and opportunistic resources  
(ATLAS and BelleII)



17% of BelleII computing in  
2015 used clouds

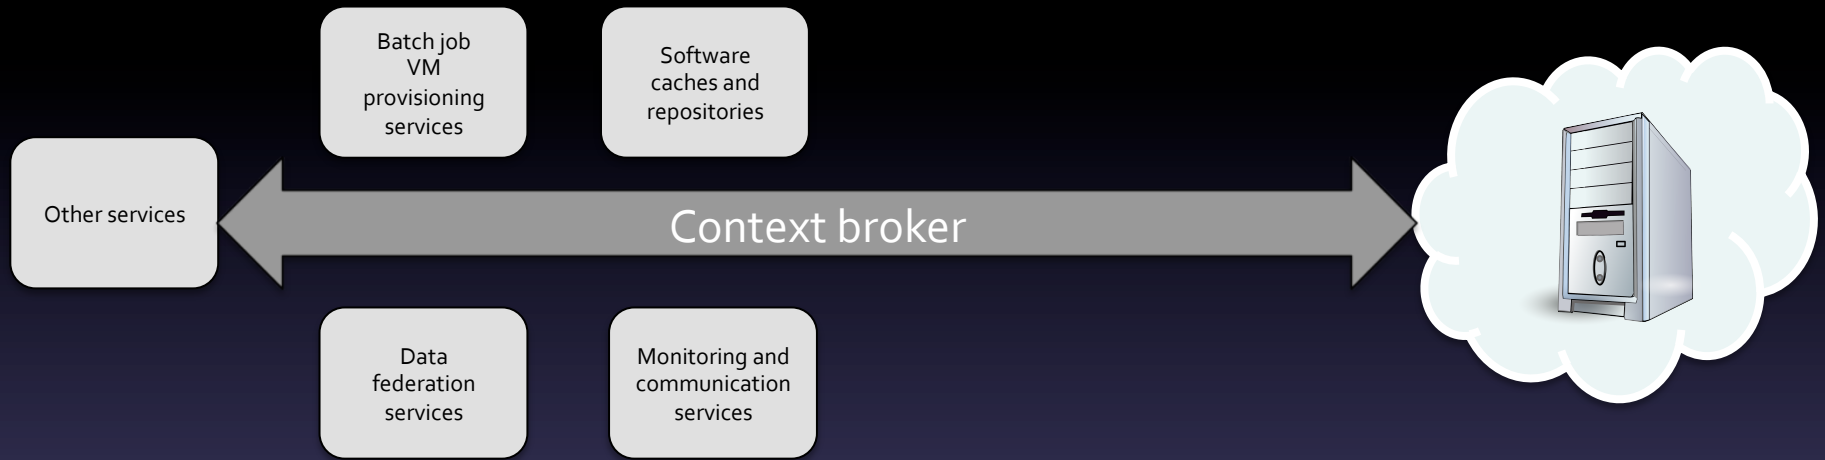
***Production use of clouds for 3-4 years with gradual increase in utilization***

*Technology is still young and rapidly evolving*

*Integrating new technologies into a production system is challenging*

# Evolution to a context-aware architecture

**Goal is a system that is dynamic, automated, intelligent, error-correcting and able to exploit any cloud willing to provide resources**



**Context information system(s) keeping track of real-time information**

State of clouds, services, squid-caches, data federations

**Use the context information to configure, operate and monitor**

Identify and monitor clouds – match jobs to resources

VMs self configure and monitor, locate the software and data

Recognize changes and errors – take corrective actions

# Context-collection

We already collect data and logs on services, clouds, user job, VMs

Real-time (message passing) communication between some services

ATLAS Victoria

cc-east (100) cc-west (75) dair-ab (4) dair-qc (6)


Cloud	CloudScheduler VMs				Condor Slots								Idle VMs
	Starting	Running	Retiring	Error	1	2	3	4	5	6	7	8	
cc-east atlas-worker	0	0	9	0	6	3	5	4	5	5	5	5	0
cc-east atlas-mcore-worker	0	23	68	0	91								0
cc-west atlas-worker	0	0	4	0	2	2	3	1	0	3	0	1	0
cc-west atlas-mcore-worker	0	71	0	0	71								0
dair-ab atlas-mcore-worker	0	3	1	0	4								0
dair-qc atlas-mcore-worker	0	6	0	0	6								0
Jade jade-worker	0	0	0	2	0	1	1	1	1				0

Jobs	Total	Condor Job Status			
		Idle	Running	Completed	Held
All	287	43	242	2	0
1 Core	85	18	66	1	0
8 Core	198	25	172	1	0
Alberta	4	0	4	0	0

Snapshot of current system

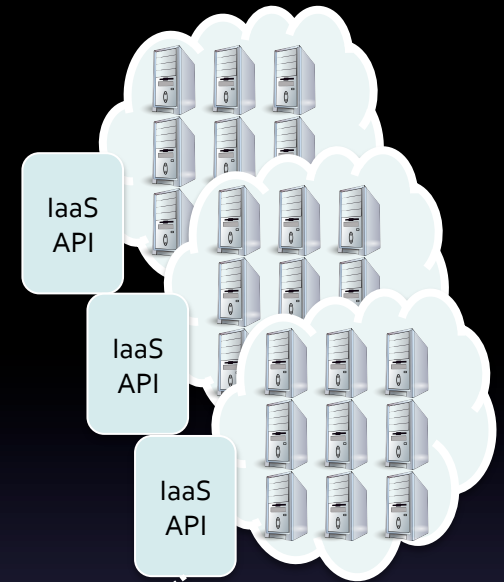
For example, monitoring page for one of the ATLAS cloud systems in Canada

#active clouds  
#jobs (1c/8C)  
#VMs  
#job slots



# VM provisioning

Entity



Context

User job requirements

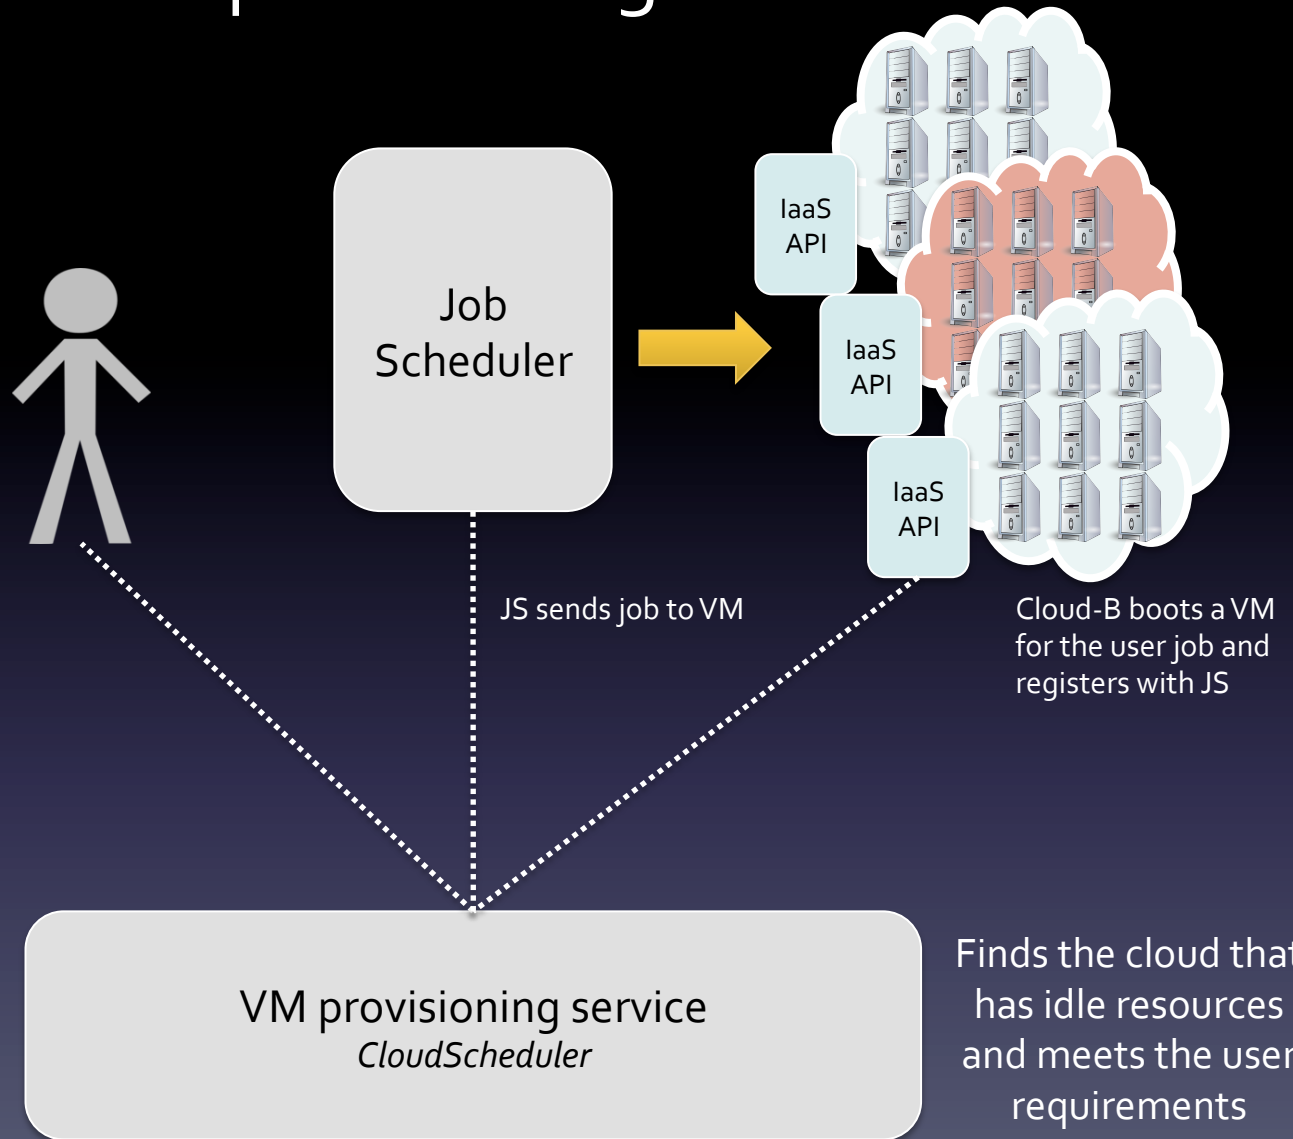
Job list

Resource availability

Context-broker



# VM provisioning



Entity


Action

Context-broker



# Software delivery service

Entity




Context

List of squids  
Location  
Load

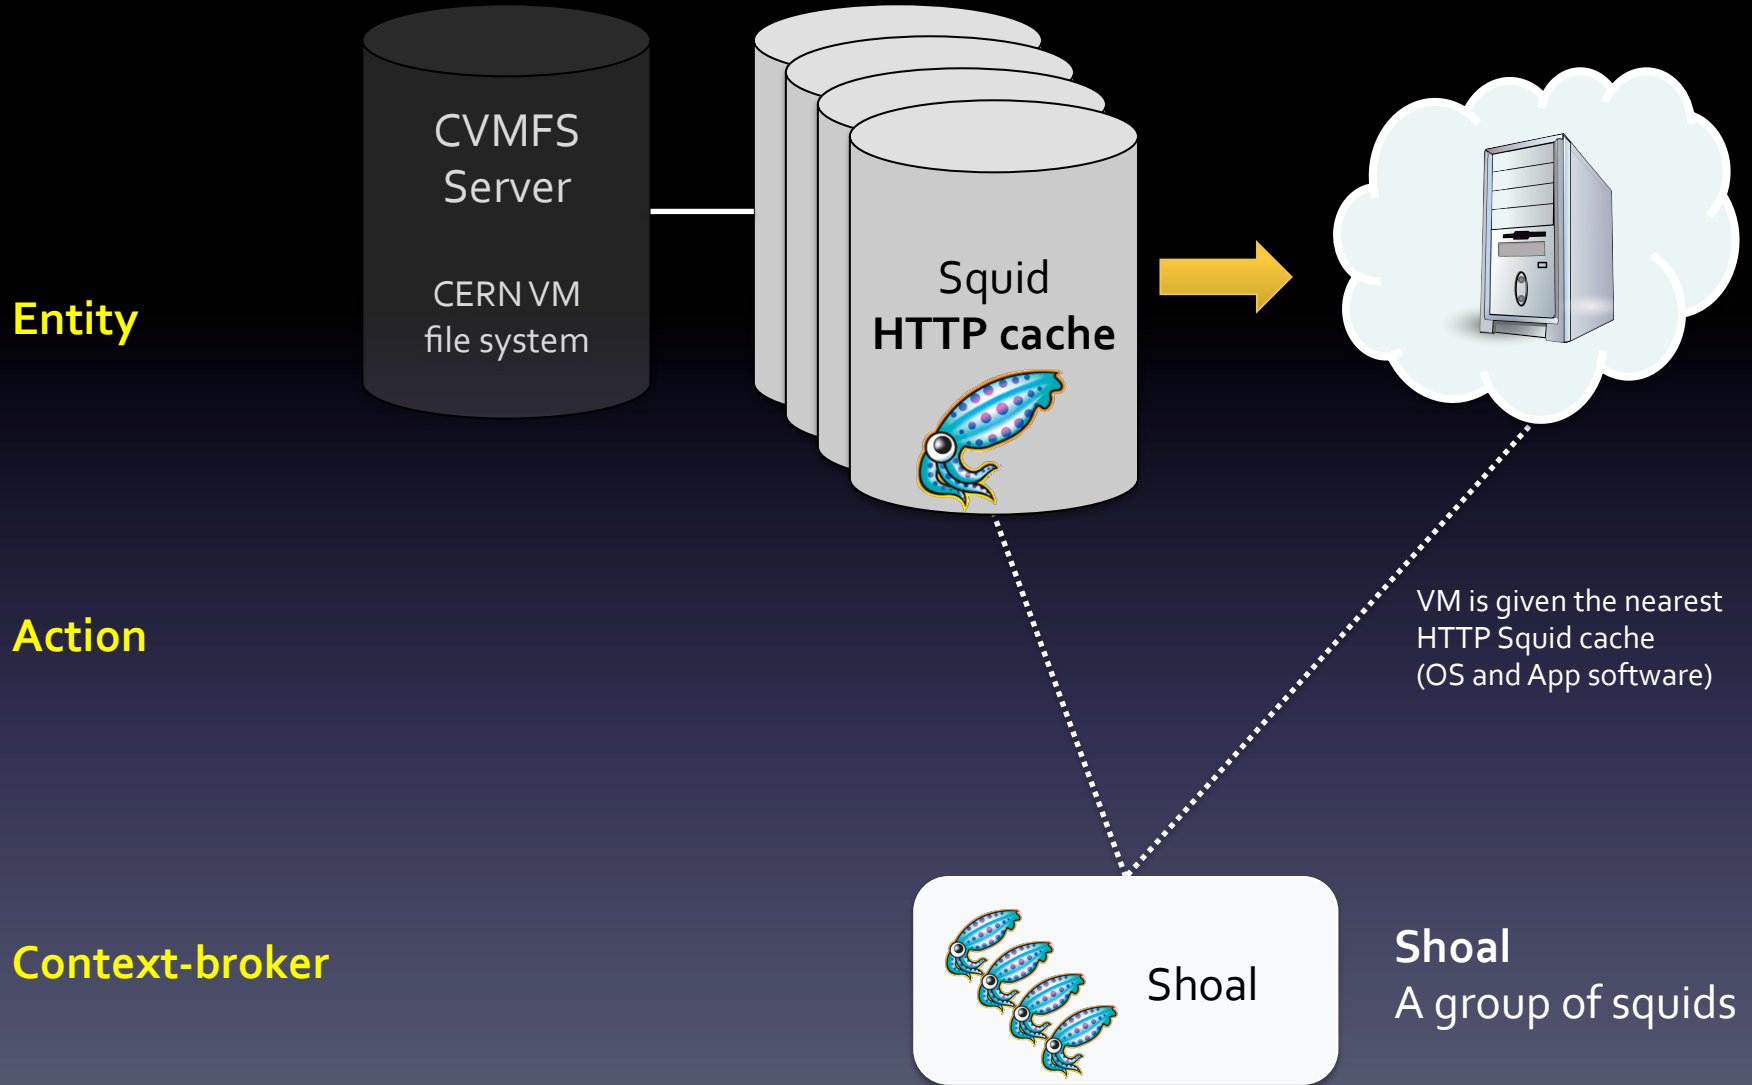
Virtual  
machine  
location

Context-broker



Shoal  
A group of squids

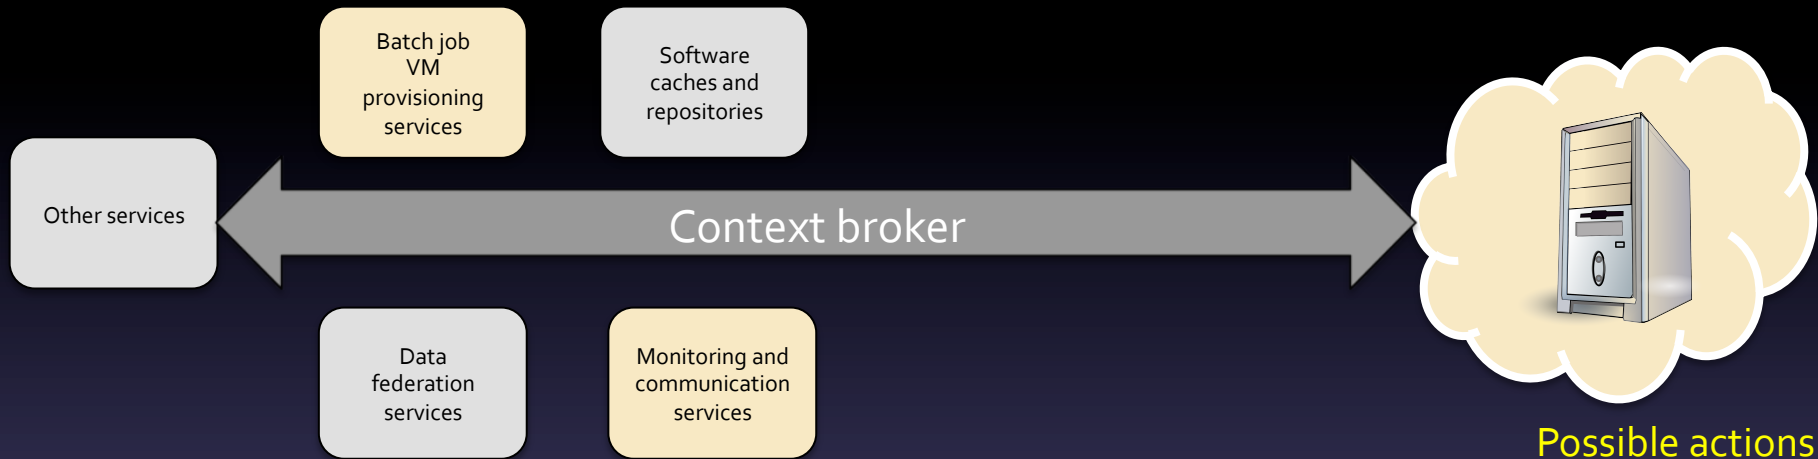
# Software delivery service



# Real-time Cloud and VM status

We would like real-time monitoring information of the remote cloud and individual VM

## Automated fault-recovery response




### Possible actions:

Cloud is automatically disabled  
VM de-registered from HTCondor

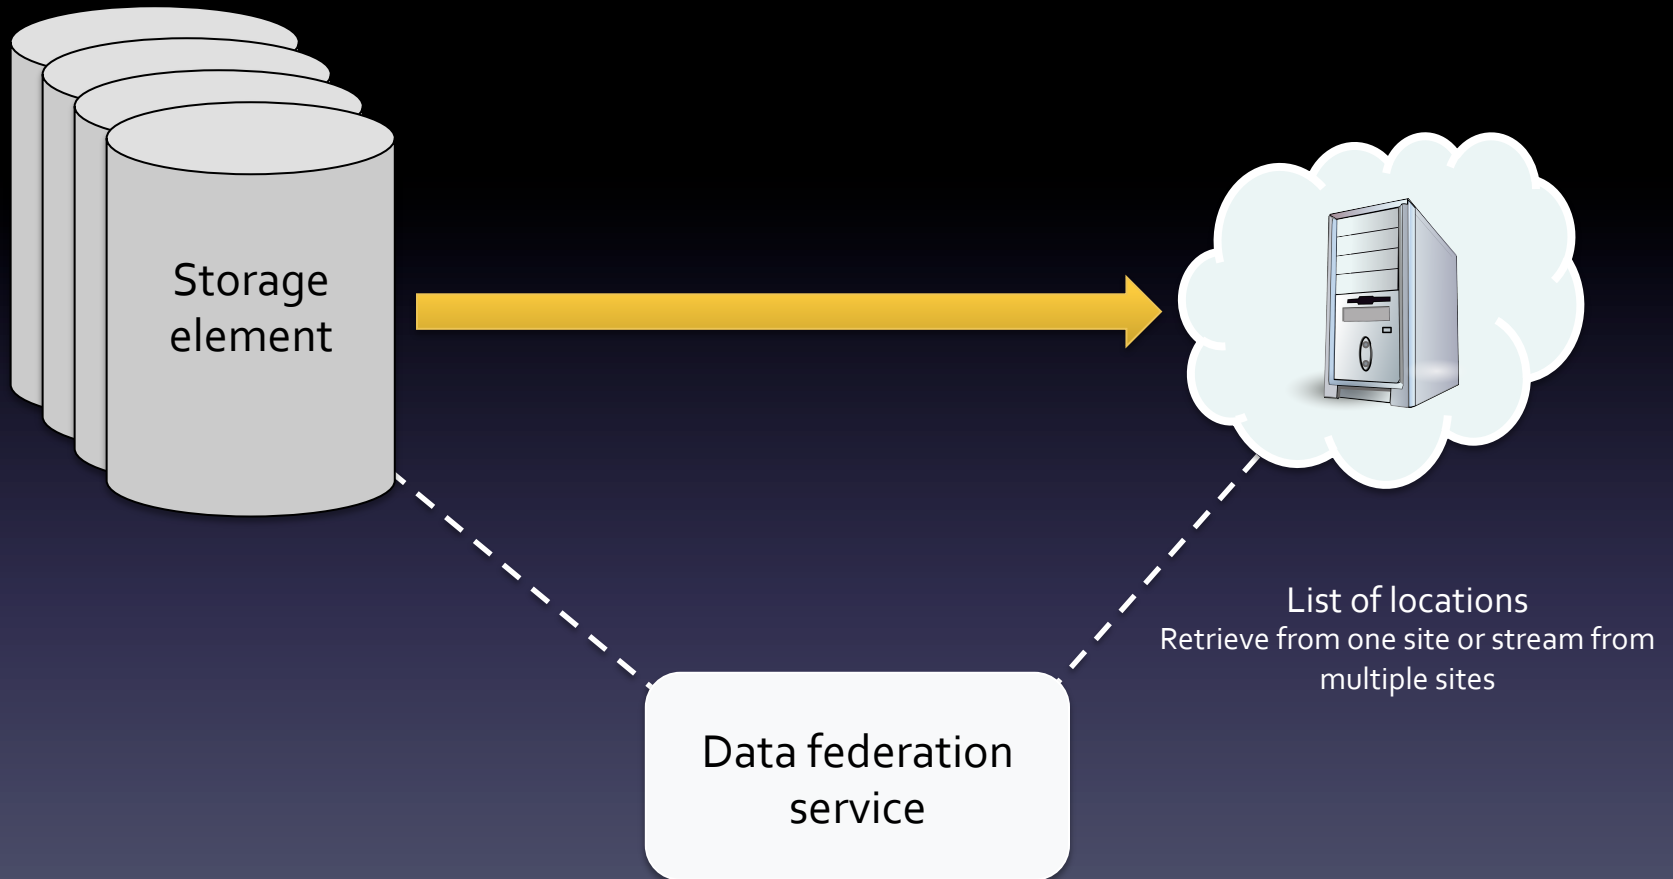
### Common issues:

- Periodically a cloud goes offline
- Application job is stalled
- VM is stalled – not responsive to the VM-provision service
- VM is lost – not registered with the VM-provision service
- Local network issues in the cloud
- VM evictions due to price fluctuations in commercial costs

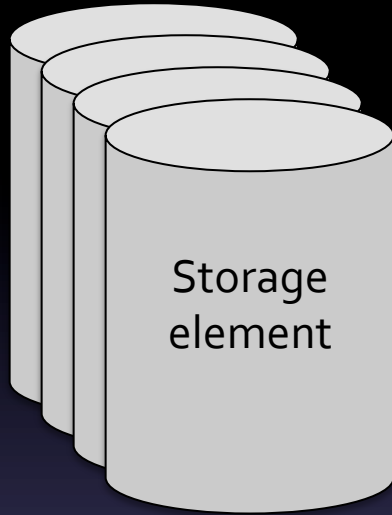
# Data federation service



# Data federation service



# Data federation service




```
~/myfed [AOD]mc11_7TeV.105011.J2_pythia_jetjet.merge.AOD.e815_s1273_s1273_12780_04541566_00 - Mozilla Firefox
~/myfed [AOD]mc11_7T...
@ug.hep.cw.ca/~/myfed [AOD]mc11_7TeV.105011.J2_pythia_jetjet.merge.AOD.e815_s1273_s1273_12780_045415...
Disable Cookies CSS Forms Images Information Miscellaneous Outline Resize Tools
/myfed
/AOD/mc11_7TeV.105011.J2_pythia_jetjet.merge.AOD.e815_s1273_s1273_12780_04541566_00
Mode UID GID Size Modified Name
-rw-rw-rw- 0 0 1.16 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000001.pool.root.1]
-rw-rw-rw- 0 0 1.20 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000002.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000003.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000004.pool.root.1]
-rw-rw-rw- 0 0 1.26 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000005.pool.root.1]
-rw-rw-rw- 0 0 1014.9M Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000006.pool.root.1]
-rw-rw-rw- 0 0 1011.6M Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000007.pool.root.1]
-rw-rw-rw- 0 0 1001.6M Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000008.pool.root.1]
-rw-rw-rw- 0 0 1000.7M Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000009.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000010.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000011.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000012.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000013.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000014.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000015.pool.root.1]
-rw-rw-rw- 0 0 1.86 Thu, 01 Jan 1970 00:00:00 GMT [AOD_541566_000016.pool.root.1]
```

Data federation service

HEP-specific: FAX  
HTTP/WebDav: Dynafed

Multiple copies distributed around the world

# Other ideas – future plans




Typically cloud resources in OpenStack are assigned to projects using “tenants”

Efforts to share resources between tenants in an automated manner to fully utilize the resources (e.g. sharing of ATLAS and BelleII resources)

Connect cloud resources to the LHCONE network?  
(e.g. Software defined networks – SDN)

Commercial cloud links to research networks



# Summary



- Use of cloud computing in HEP is growing
  - Using dedicated and opportunistic resources to provide an increasing fraction of the needs of HEP experiments
- Context-aware cloud computing will enable us to scale our use of distributed resources that is required by HEP projects
  - Utilize technologies developed inside and outside HEP to simplify management, expand functionality and ease of use
- Valuable to other research communities