

Contextualization: Providing One-Click Virtual Clusters

Kate Keahey

Tim Freeman

Argonne National Laboratory

University of Chicago

{keahey, tfreeman}@mcs.anl.gov

Nimbus: <http://workspace.globus.org>

Cloud Computing

Infrastructure-as-a-Service (IaaS)



Science Cloud



How do I log in?
How does my colleague log in?
How do I point my VM at the server I need?
Not this server, the other server...
How can I make all these things
happen automatically?

Azure Services Platform


Nimbus: <http://workspace.globus.org>

Parameterizable Appliance

- Appliance = environment for application
- Parameterizable = can be customized

deployment time

boot time




To learn more:

- *Paper: “A Scalable Approach To Deploying And Managing Appliances”, TeraGrid 2007*
- *Open Virtualization Format (OVF) Specification*

Nimbus: <http://workspace.globus.org>

What if not all information is available at deployment?

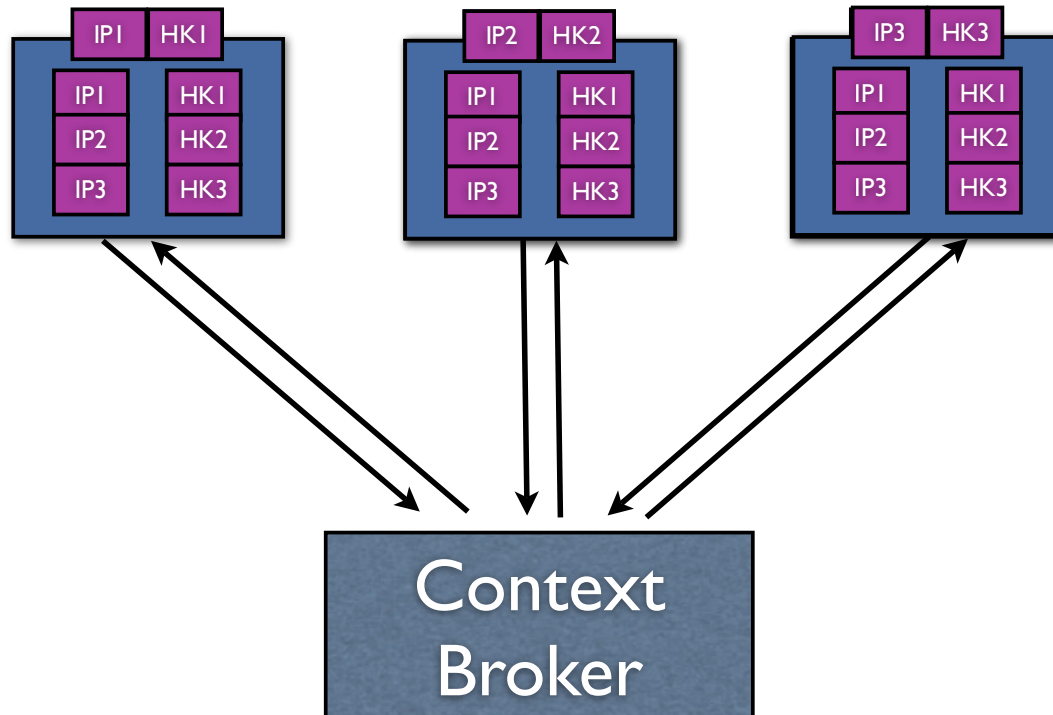
Tightly-Coupled Cluster



Reciprocal exchange of information

Nimbus: <http://workspace.globus.org>

Solution: Brokering Context Information



Nimbus: <http://workspace.globus.org>

More generally:

- Who orchestrates this exchange?
- Will this entity work with my service/appliance?
- Will my appliance work with your service?
- How do we exchange this information securely?
- Can we orchestrate this exchange over multiple providers?

Nimbus: <http://workspace.globus.org>

Roles and Responsibilities


Appliance Providers

marketplaces
commercial providers
communities

Deployment Orchestrator

orchestrate the deployment of
environments across possibly
many resource providers:
Context Broker

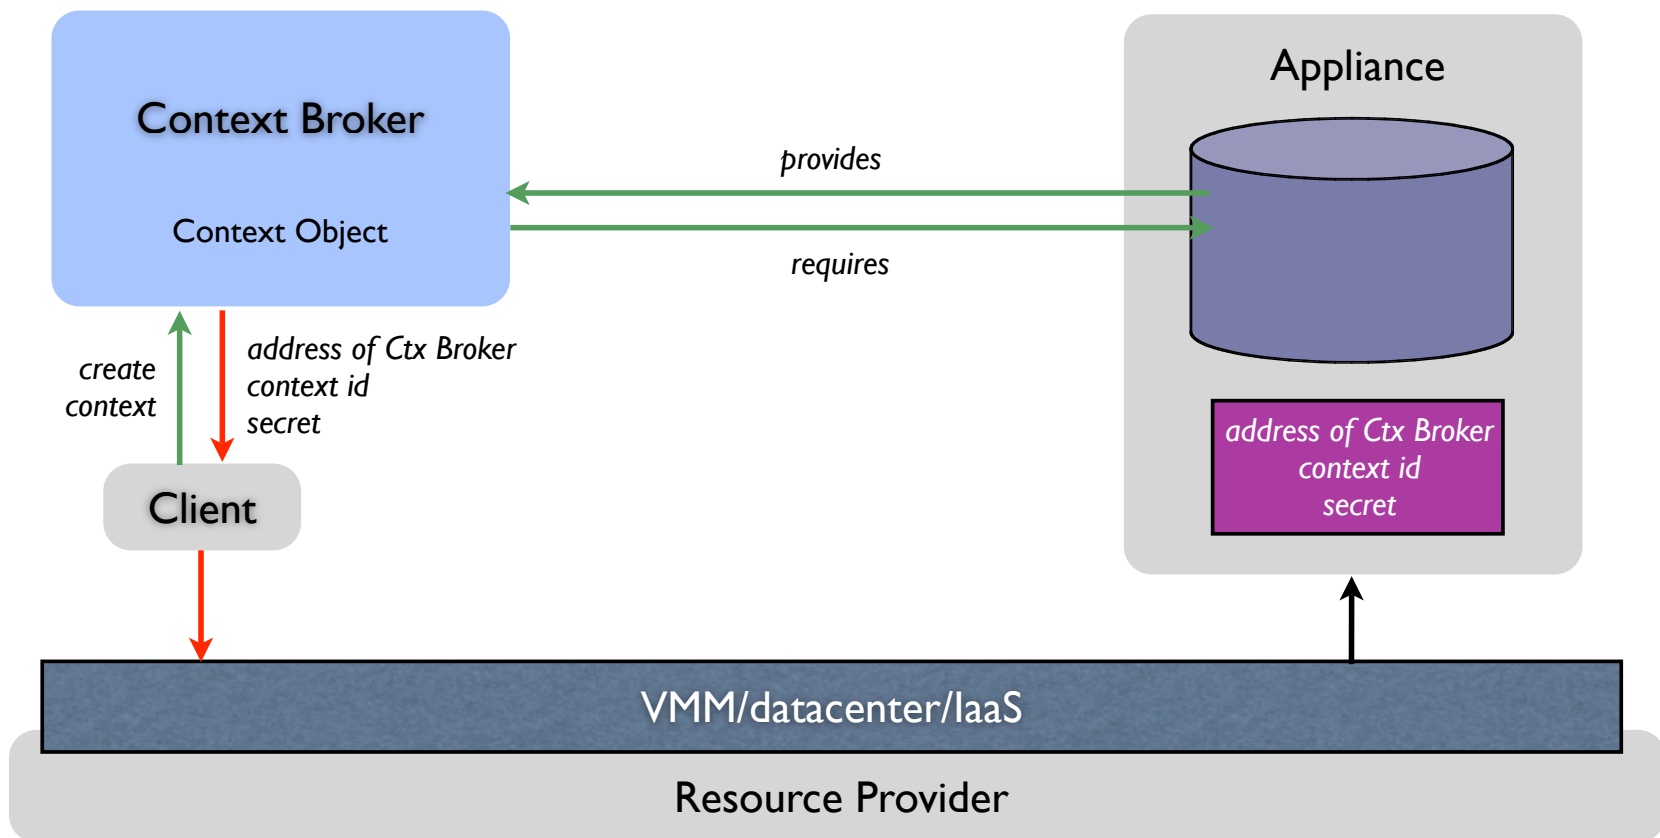
VMM/datacenter/laaS



Resource Providers

Nimbus: <http://workspace.globus.org>

The Architecture



Nimbus: <http://workspace.globus.org>

Interactions

- **Appliance Provider**
 - Ctx template including generic information
 - Means of integrating information
- **Appliance Deployer/Resource provider**
 - Obtain and provide generic context information securely
- **Deployment Orchestrator**
 - Interpret and broker non-generic information
Nimbus: <http://workspace.globus.org>

Implementation (Context Broker)

- **WSRF service**
 - Create: security context and id
 - AddWorkspace: register an appliance
 - AddInformation (poll for information release)
- **Contextualization template**
 - provides and requires
 - roles

Nimbus: <http://workspace.globus.org>

Implementation (Generic Context)

- EC2: use instance metadata (user-data field -- up to 16K of unstructured information)
- Nimbus (WSRF): provide information by appliance patching (OVF-like)
- In both cases:
 - provide information to the service secured using HTTPS channel
 - rely on secure networking within the implementation

Nimbus: <http://workspace.globus.org>

Example (NFS Server)

NFS Server
10.0.0.1

Provides:
id: 10.0.0.1
label: nfsserver
Requires:
role: nfscient
10.0.0.2
10.0.0.3

NFS Client
10.0.0.2

Provides:
id: 10.0.0.2
label: nfscient
Requires:
role: nfsserver
10.0.0.1

NFS Client
10.0.0.3

Provides:
id: 10.0.0.3
label: nfscient
Requires:
role: nfsserver
10.0.0.1

What is this good for?

- Production STAR clusters of up to 100 nodes on Science Clouds and EC2
- Deployment and integration into community infrastructure of Alice nodes
- Evaluating Montage Workflows, Hoffa et al
- CloudBLAST, Matsunaga et al.
- Ready-to-use generic OSG cluster

Nimbus: <http://workspace.globus.org>

Appliance Providers

- Right now VMs are contextualized manually
 - Black Art
 - But: contextualize once -- deploy many times
- Working with appliance providers
 - rPath: review and implementation
 - Bcfg2: contextualizing generic appliance

Nimbus: <http://workspace.globus.org>

Related Work

- Open Virtualization Format (OVF) Specification
- Configuration on the fly
 - VMPlant and “Cluster on the Fly”
- Sapuntzakis et al., “Virtual Appliance in the Collective”
- Configuration Tools
 - Bcfg2, rBuilder, Quattor, Puppet, etc.
Nimbus: <http://workspace.globus.org>

Conclusions

- Increasing importance of appliance providers
 - We need higher-level languages to “code” appliances
- Standards in many dimensions
 - “rough consensus and working code”
- Increased interest in cloud interoperability

Nimbus: <http://workspace.globus.org>

Acknowledgements

- Funding
 - NSF SDCI “Missing Links”
 - NSF CSR “Virtual Playgrounds”
- The Nimbus Team
 - Come see us at <http://workspace.globus.org/>

Nimbus: <http://workspace.globus.org>