ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Dynamic Resource Provisioning for HPC in the Cloud

Gabriel Mateescu
mateescu@acm.org

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Overview

- The case for dynamic provisioning
- xCAT architecture
- MOAB and xCAT
- Torque and xCAT
- Elastic resource management
- Conclusions

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Removing resource silos

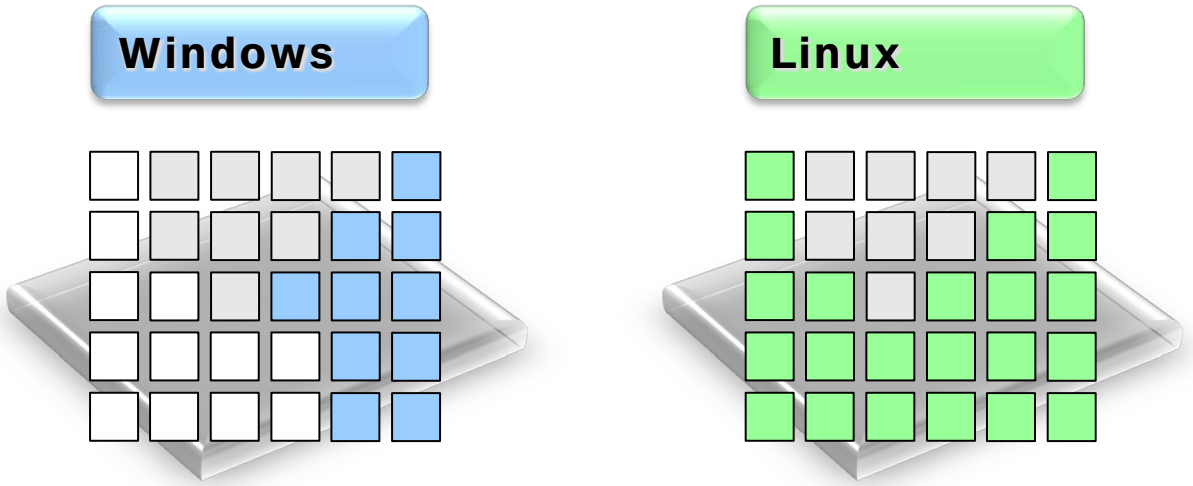

- Static resource provisioning leads to inefficient resource utilization and limits scalability

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Dynamic Provisioning

- Make execution environment changeable using on-demand provisioning
  - Traditional: stateful or stateless
  - Virtual Machine provisioning
- Dynamic provisioning enables elasticity of resources
  - Remove resource silos
  - Reallocate resources based on workload and priorities

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Provisioning Paradigm Shift

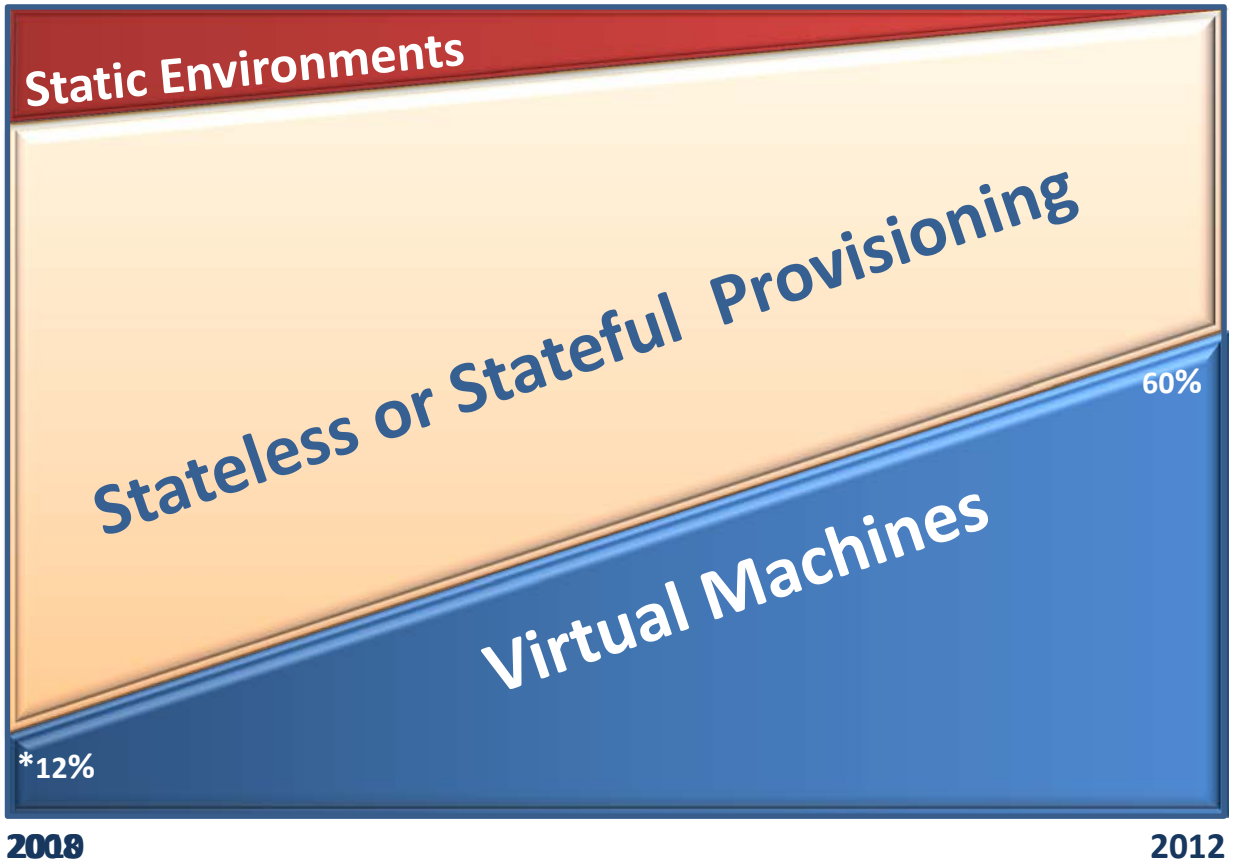

xCAT

eXtreme Cluster/Cloud Administration Toolkit

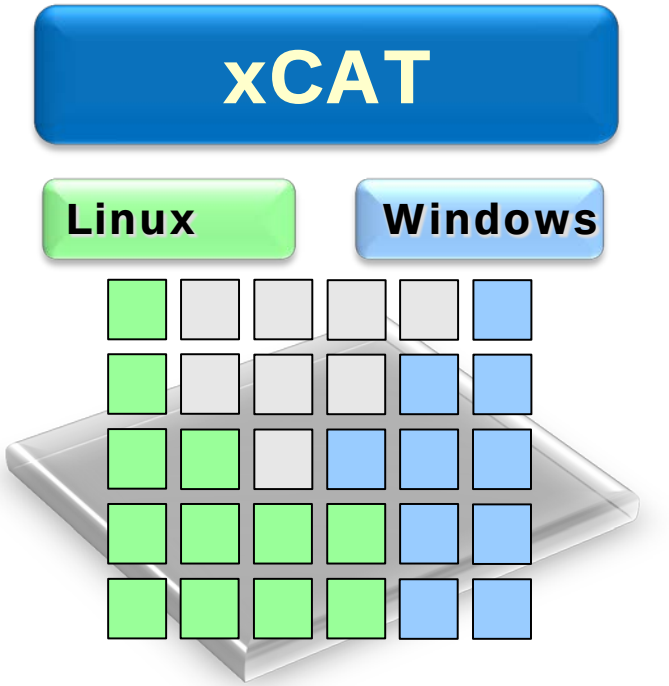

- Dynamic resource provisioning improves resource utilization and manageability

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

xCAT

- Software provisioning
  - Traditional: stateful or stateless
  - Virtual Machine provisioning
- Hardware management
  - monitoring, alerts
- Implementation
  - Configuration database
  - Distributed service

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

xCAT architecture

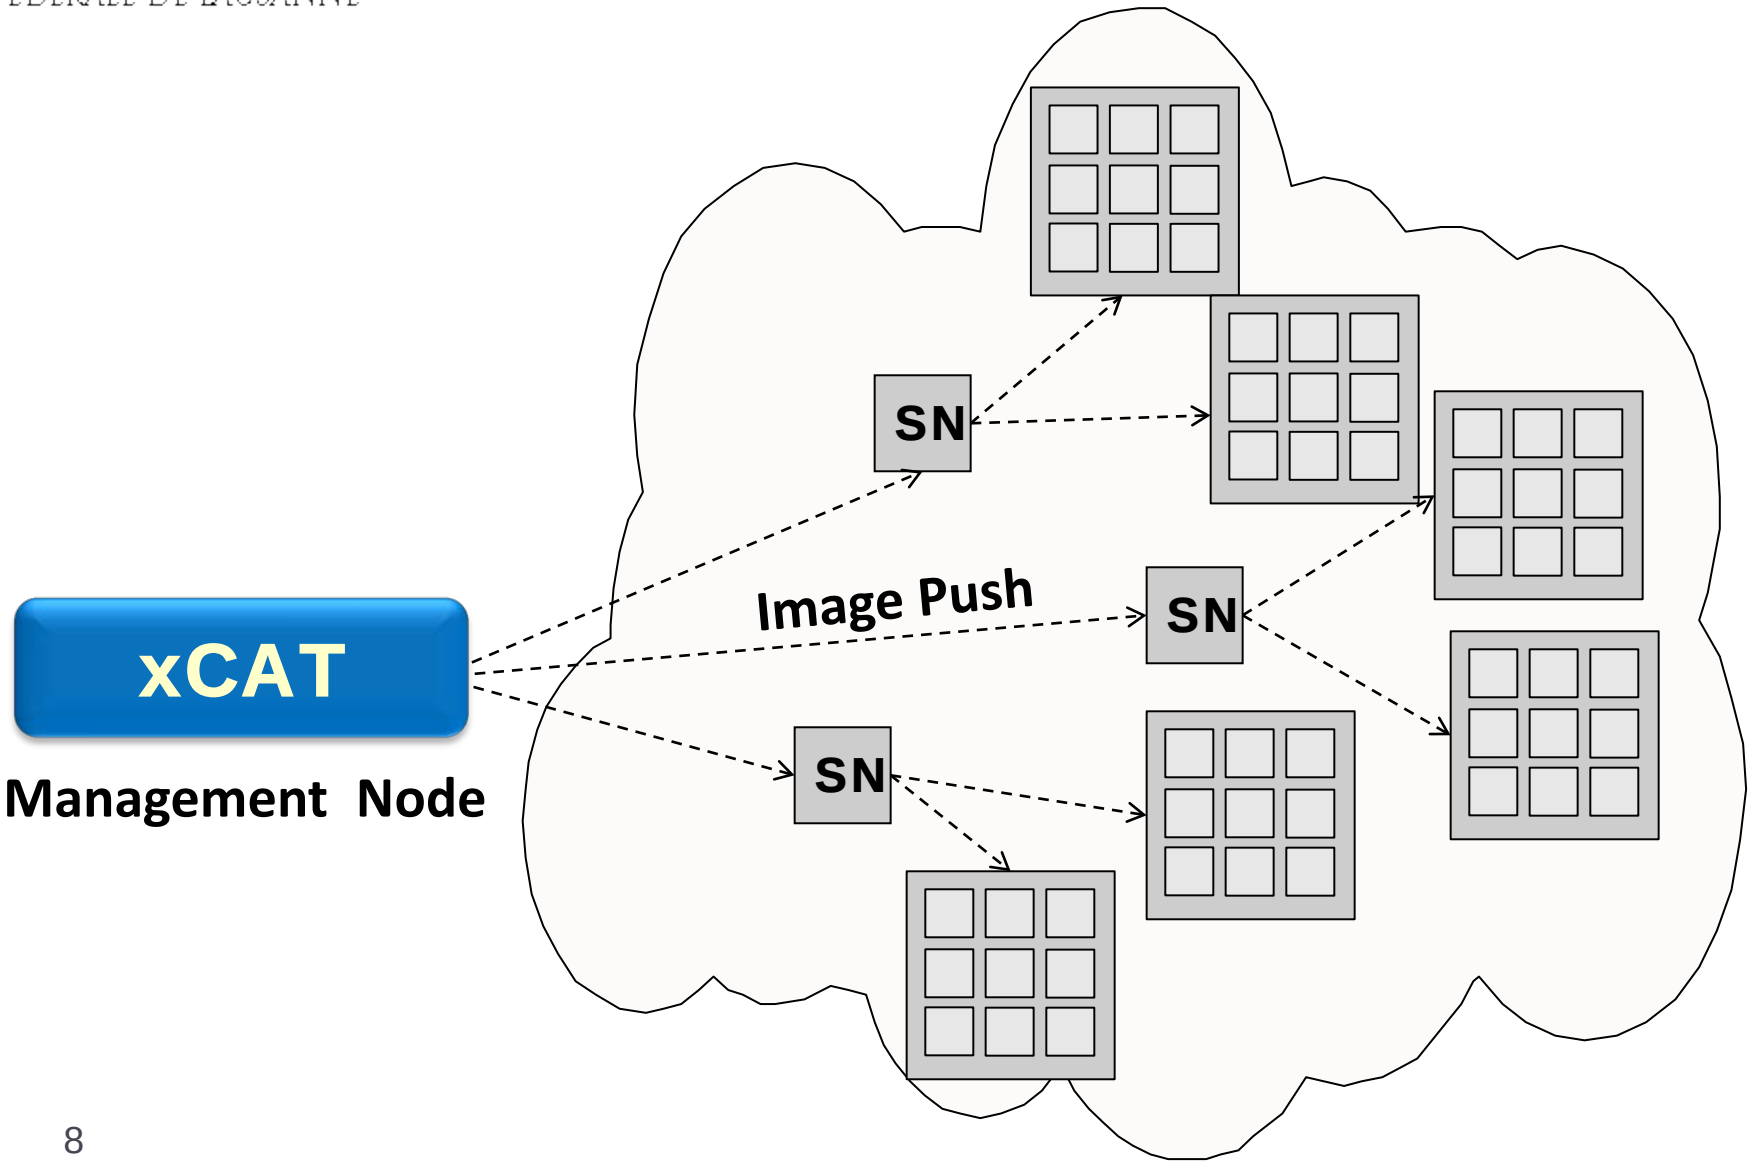

Provisioning methods

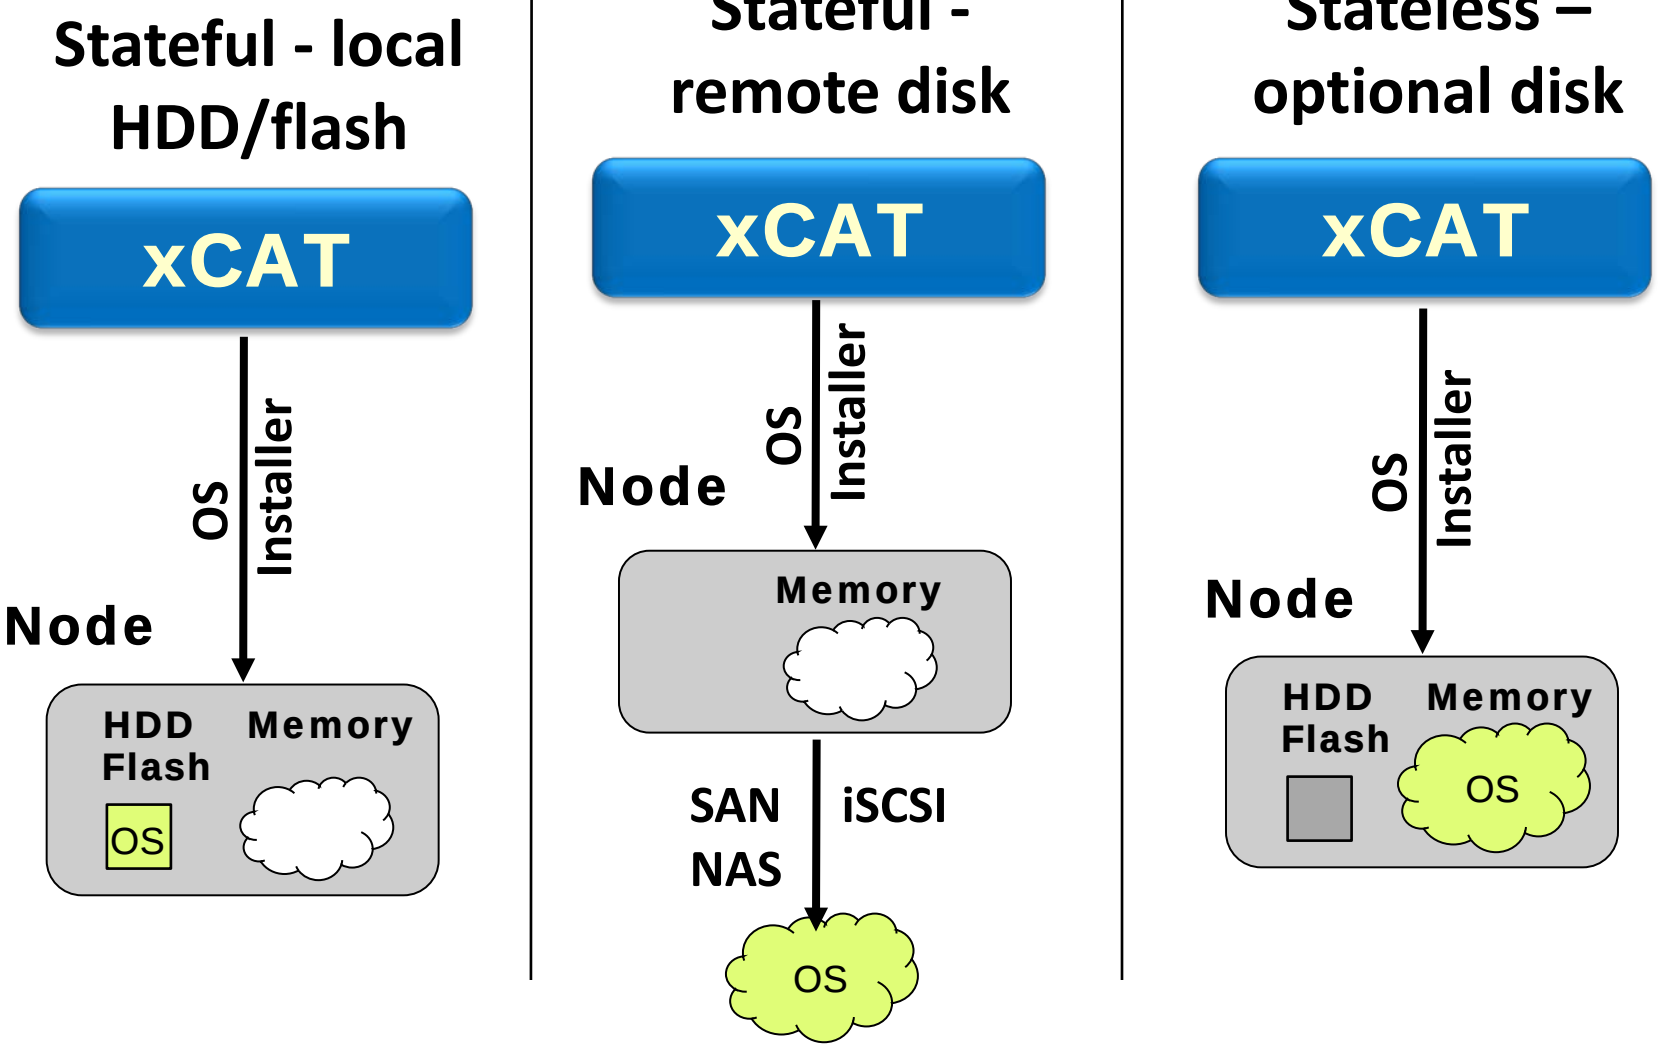

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

MOAB and xCAT

- Software provisioning
  - Traditional: stateful or stateless
  - Virtual Machine provisioning
- Hardware management
  - monitoring, alerts
- Implementation
  - Configuration database
  - Distributed service

MOAB and Dynamic resources

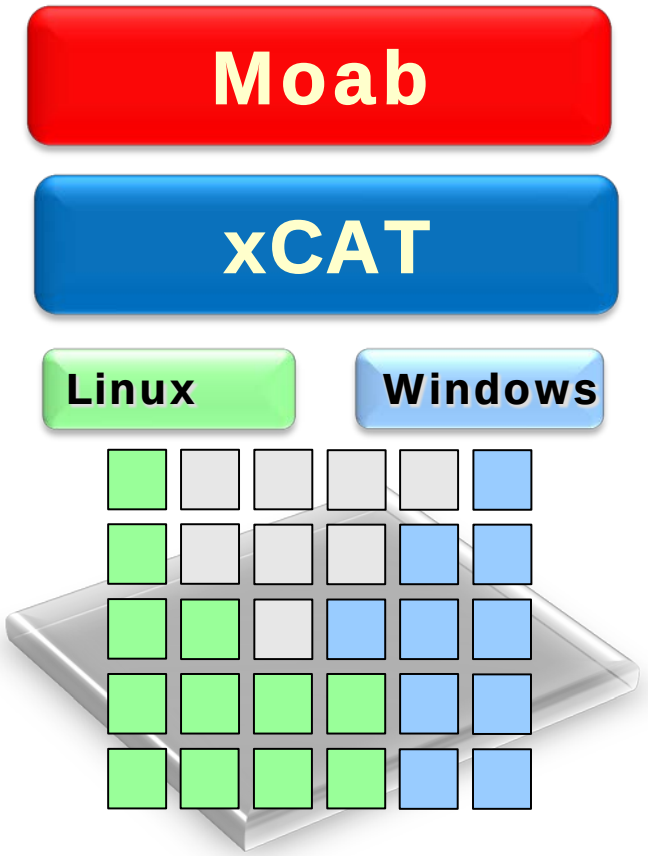

- Dynamic resource provisioning improves resource utilization and manageability

MOAB/xCAT architecture

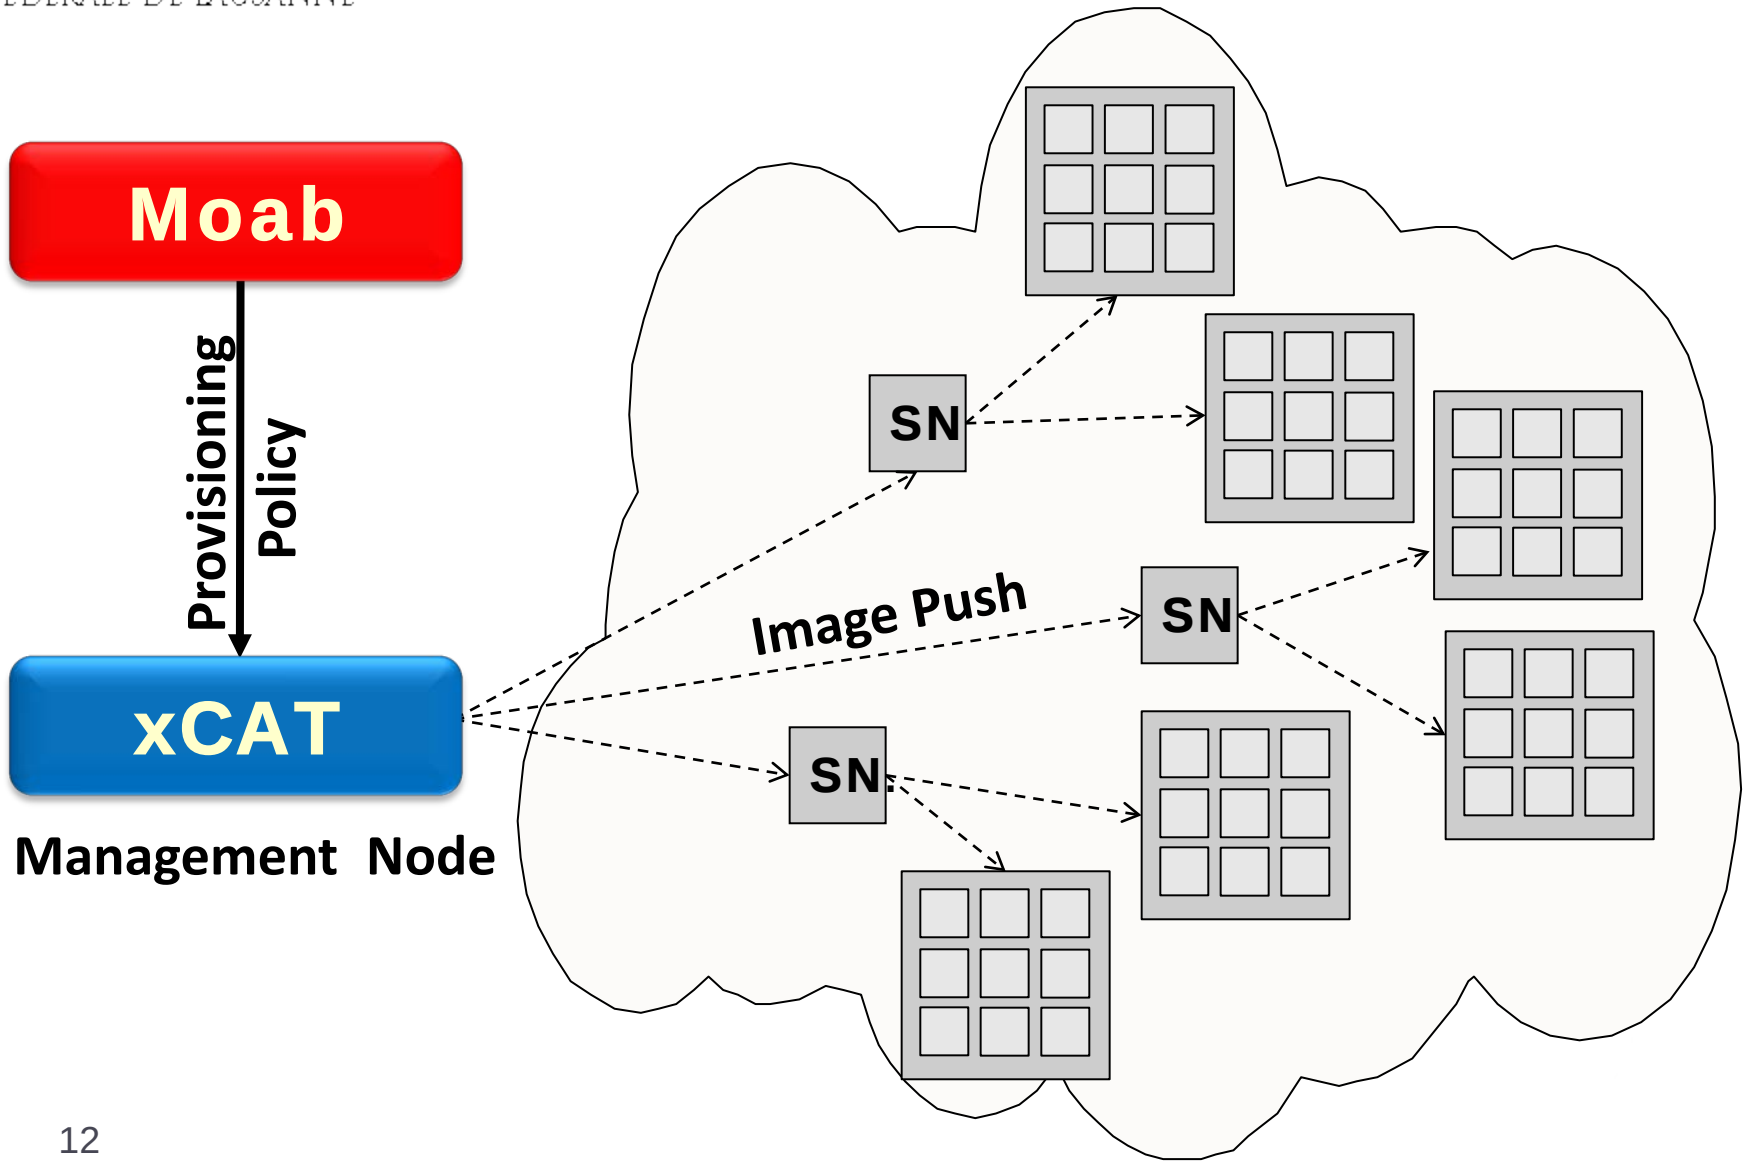

MOAB/xCAT functionality

Intelligent Management

- **Provisioning**
- **Policy driven resource allocation**
- **Power: power off idle nodes**
- **QoS**
- **Fault recovery**

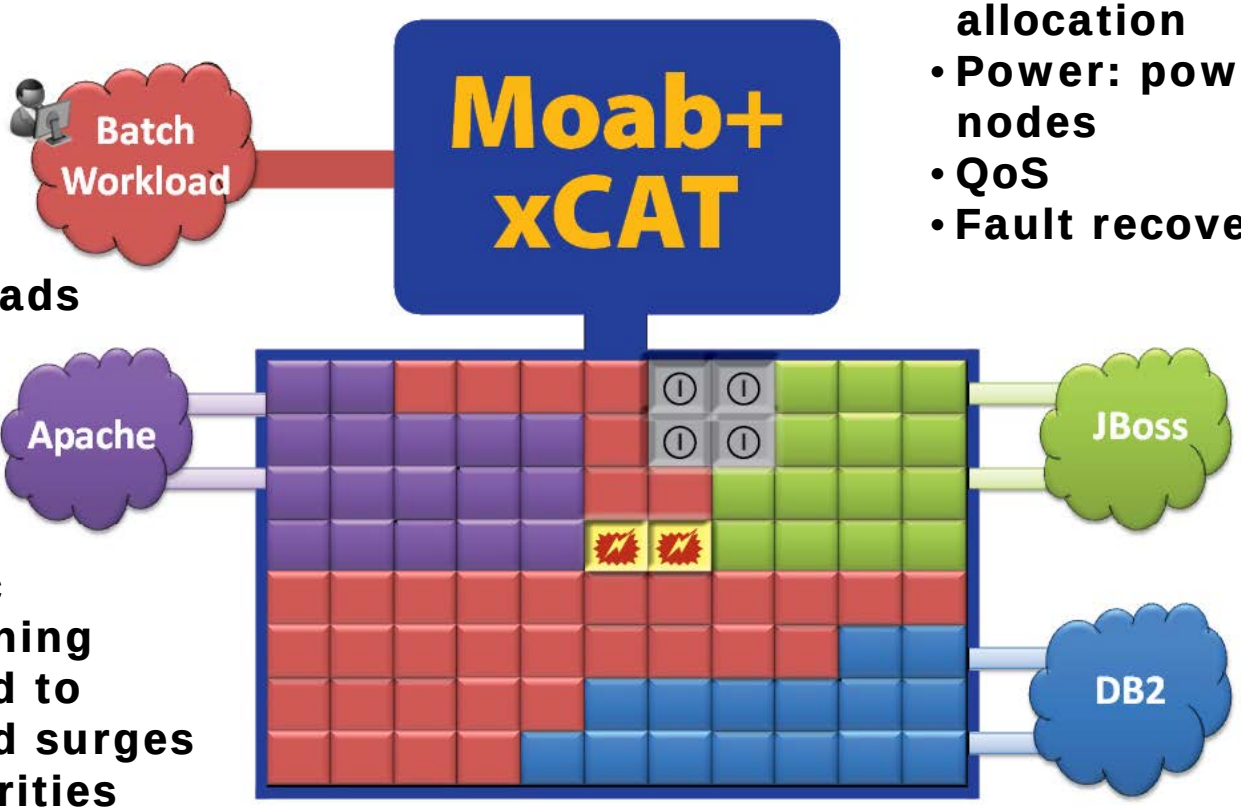

Mixed Workloads

Dynamic Provisioning
•Respond to workload surges and priorities

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Provisioning Sequence

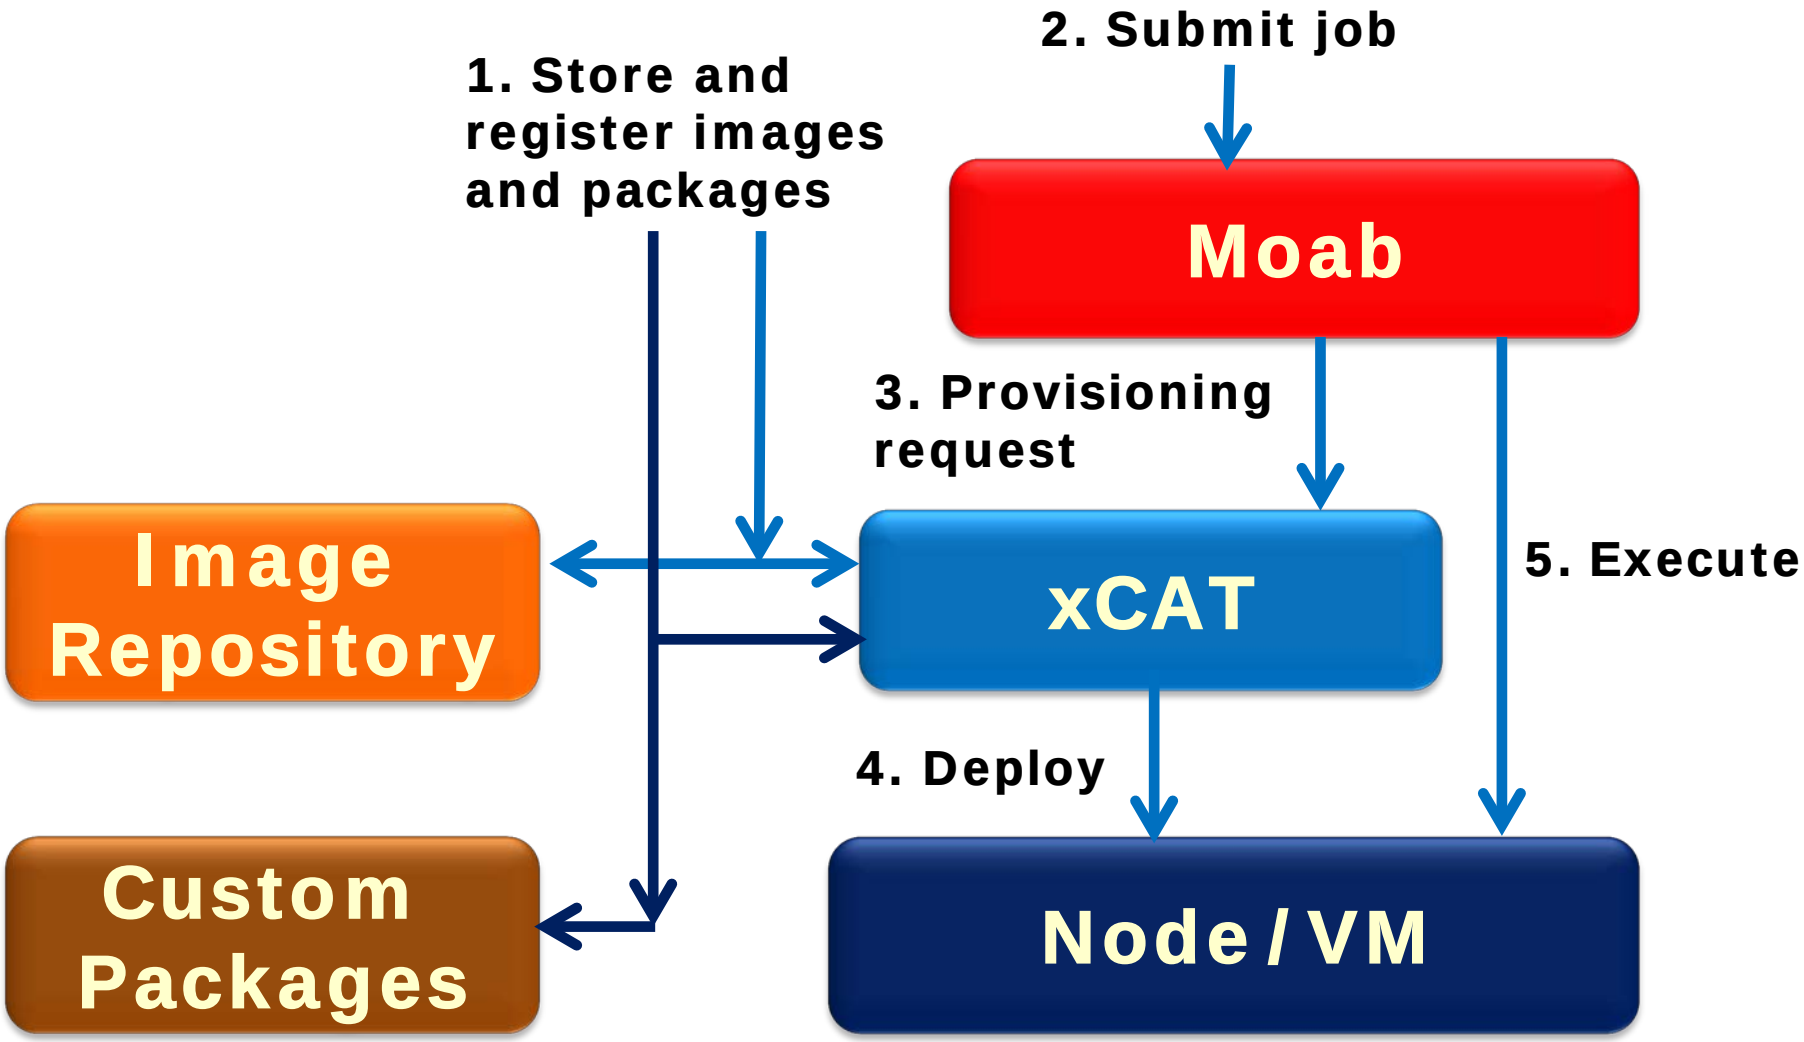

Managing multiple clusters

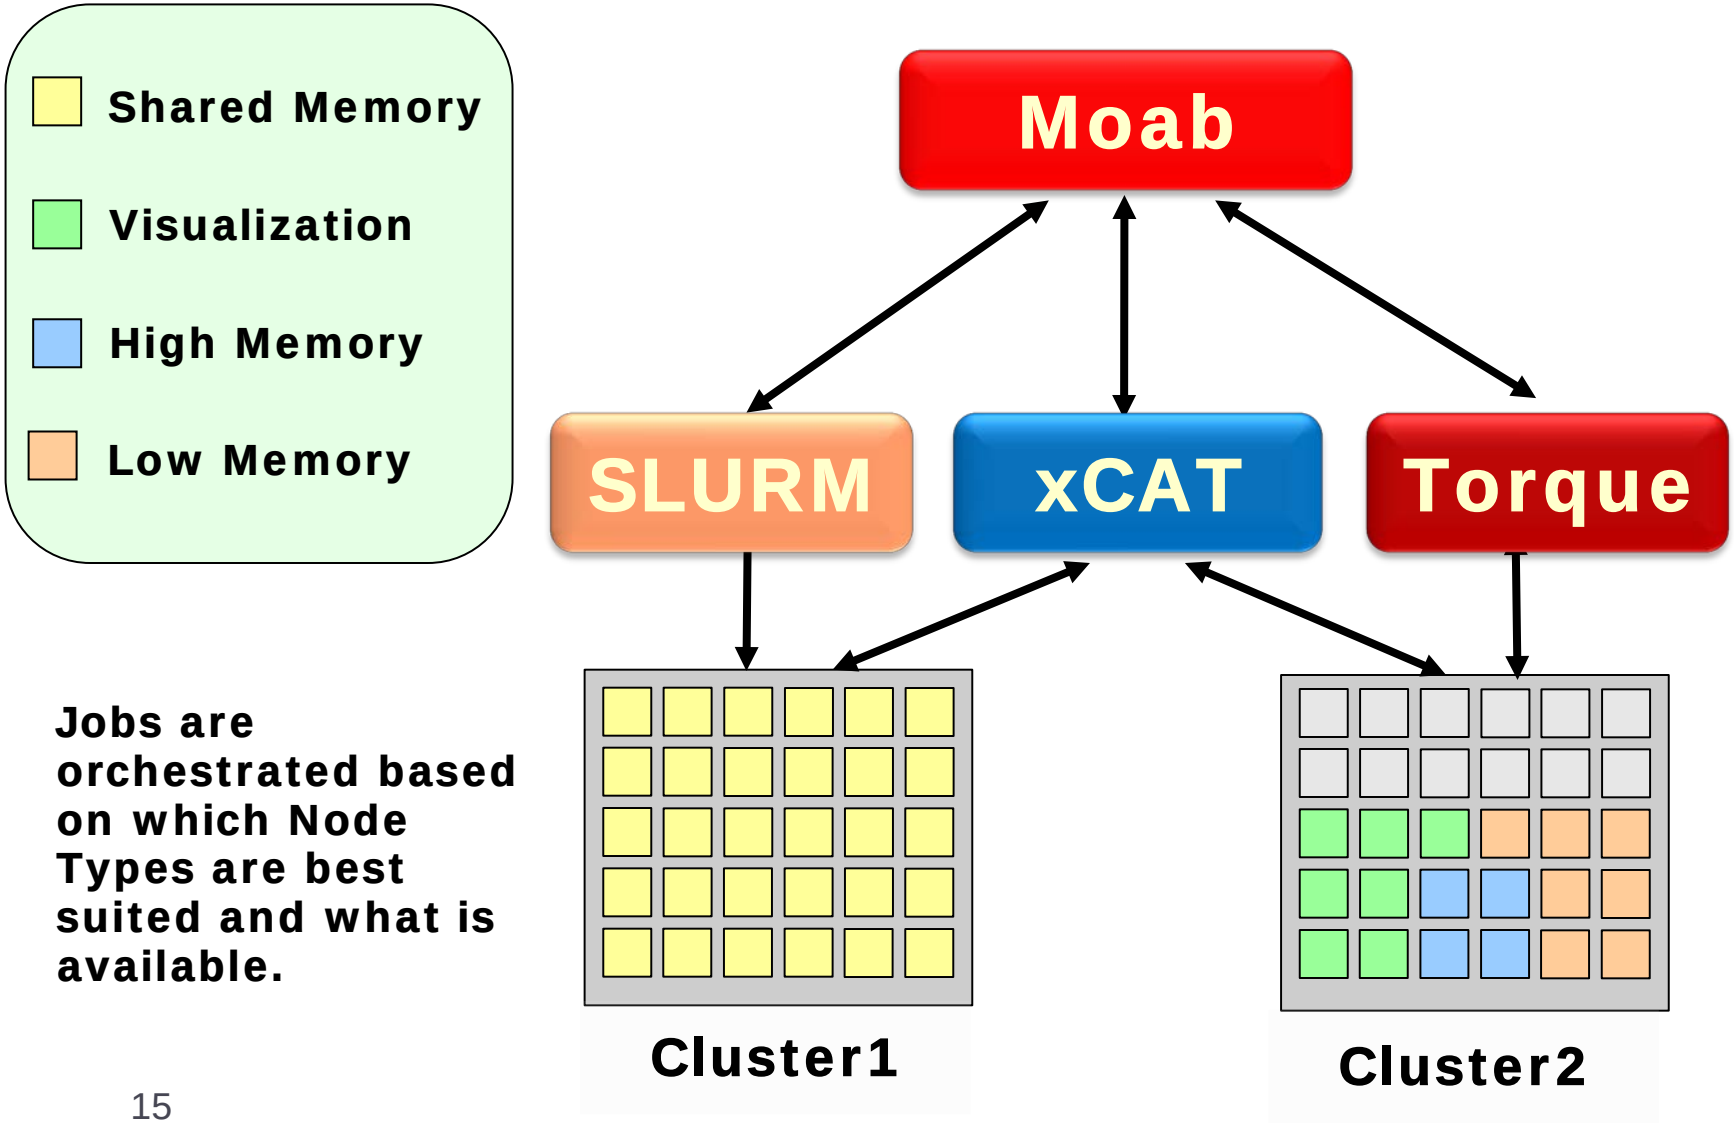

Jobs are orchestrated based on which Node Types are best suited and what is available.

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Torque and xCAT

- The xCAT web site refers to my IBM article
  - xCAT site page
  - IBM developerWorks article
- The article describes diskless provisioning
  - PXE (pre-boot execution environment) network booting
  - xCAT using a management network and the Intelligent Platforms Management Interface (IPMI) to control the nodes

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Elastic Applications

- Elastic applications adapt to variable work amount and resource capacity
  - Iterations that require more work use more processes
- Two types of elastic applications
  - Malleable: changes in job size are initiated by the system
  - Evolving: changes in job size are initiated by the application
- Workflows are evolving applications

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Elastic Applications (2)

- Malleable applications
  - the cloud infrastructure dynamically changes the resources allocated to a job
  - enables auto-scaling based on performance monitoring
  - eliminates preemption of lower priority jobs
- Evolving applications
  - request different amounts of resources over their lifetime
  - the cloud grants the requests

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Conclusions

- An automated dynamic provisioning system includes:
  - provisioning service
  - image and package repositories
  - workload and resource management service
- An HPC cloud supports
  - rigid (traditional) HPC applications
  - elastic applications such as an workflow execution engine