Cloud Deployment at KEK

Wataru Takase

Computing Research Center, KEK, Japan

Background: KEK Linux Cluster

- 10000 CPU cores
- Scientific Linux 6
- IBM Spectrum LSF

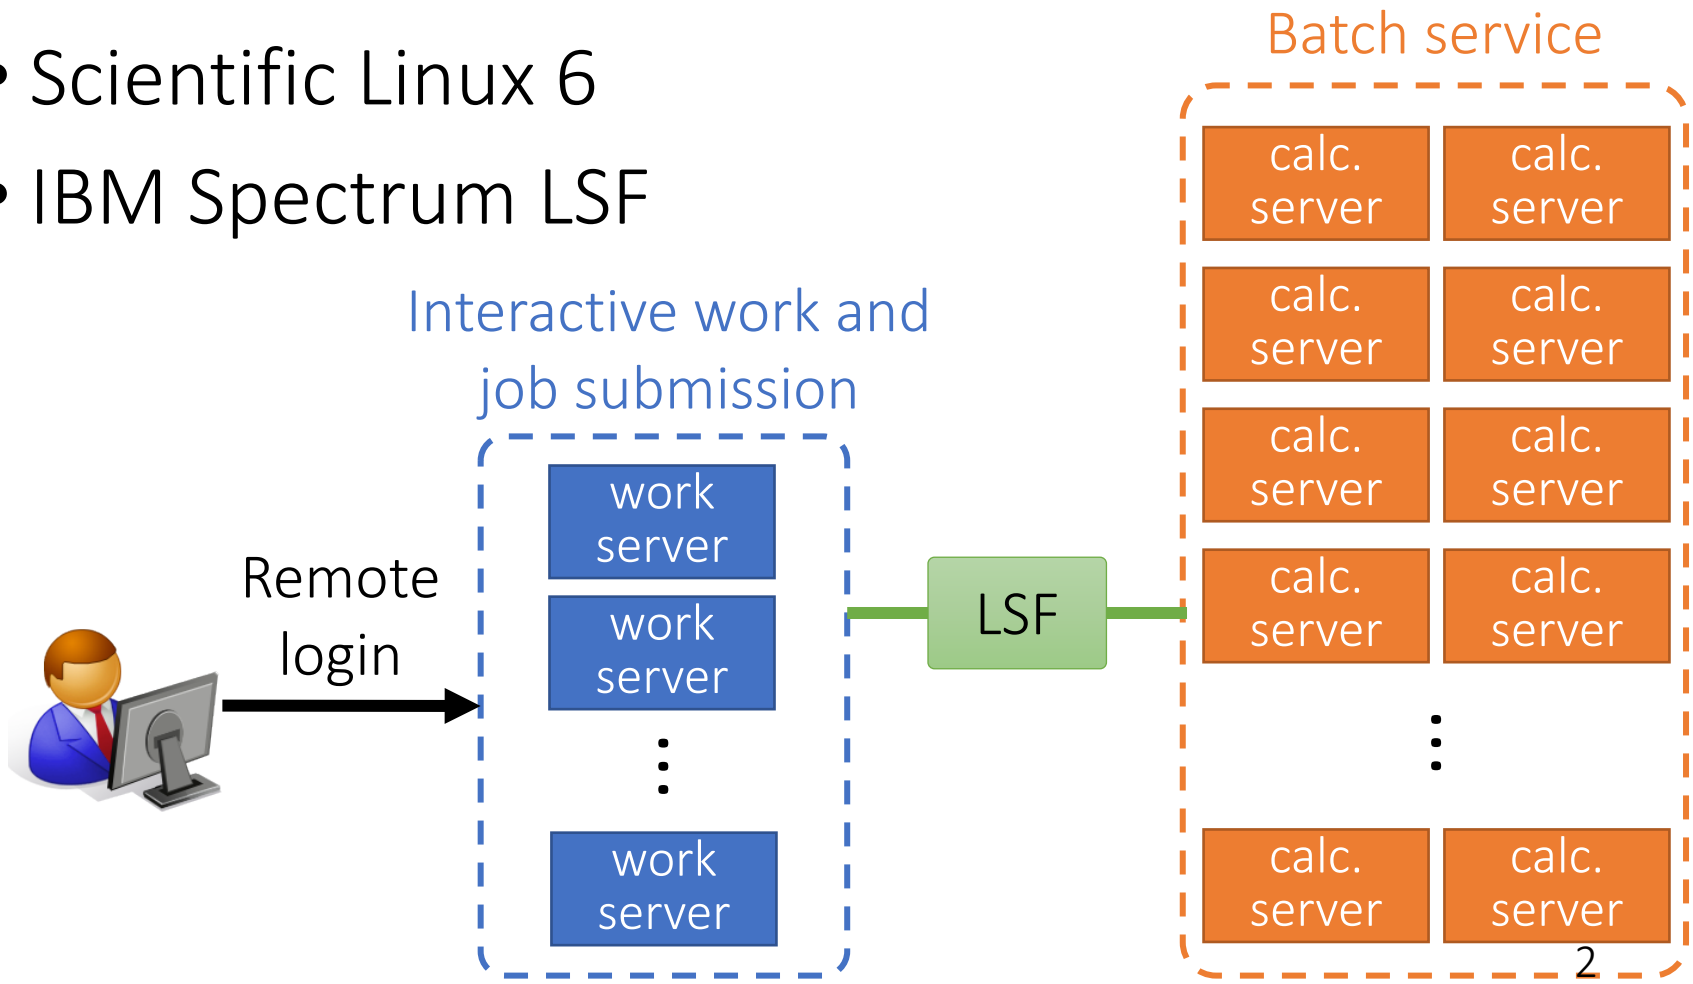

Background: Need Workload Management for Different Groups

- Requirements on specific system
  - Develop an application on the other OS.
  - Test for newer OS/Libraries.
  - Stick to old OS.
- Efficient management of limited resources

Take advantage of Cloud computing

IBM Cloud Manager with OpenStack (CMO) [1]

- CMO
  - IBM Cloud software based on OpenStack
- Additional features:
  - Simplified service portal
  - IBM Platform Resource Scheduler
  - Simplified Cloud deployment by Chef

[1] https://www.ibm.com/support/knowledgecenter/en/SST55W

CMO: Simplified Service Portal

- Quick summary of all of relevant projects
- Activity logs
- Launching instance
- Request for launching instance (optional)

Simplified service portal for users

OpenStack dashboard for cloud admins

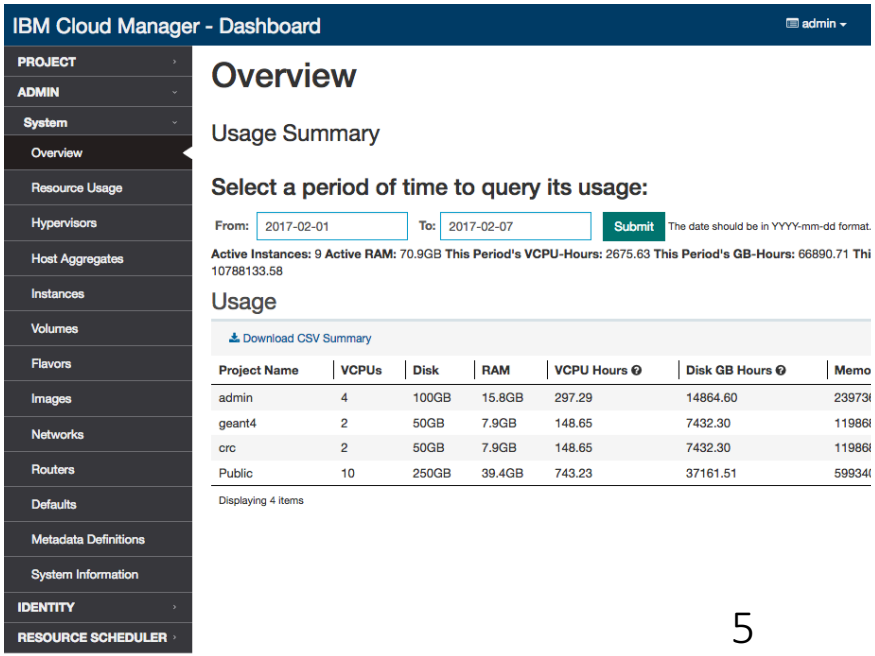

CMO: IBM Platform Resource Scheduler [2]

- Extends Nova compute scheduler.
- Policy based VM deployment / reallocation.
  - Packing policy
    - Pack VMs to minimum number of compute nodes as much as possible.
  - CPU load balance policy
    - Balance CPU load among compute nodes.

[2] https://www.ibm.com/support/knowledgecenter/SS8MU9/

CMO: Simplified Cloud deployment by Chef

- Automate deployment by predefined Cloud topologies.
  - e.g. Minimal, Controller + *N* compute nodes

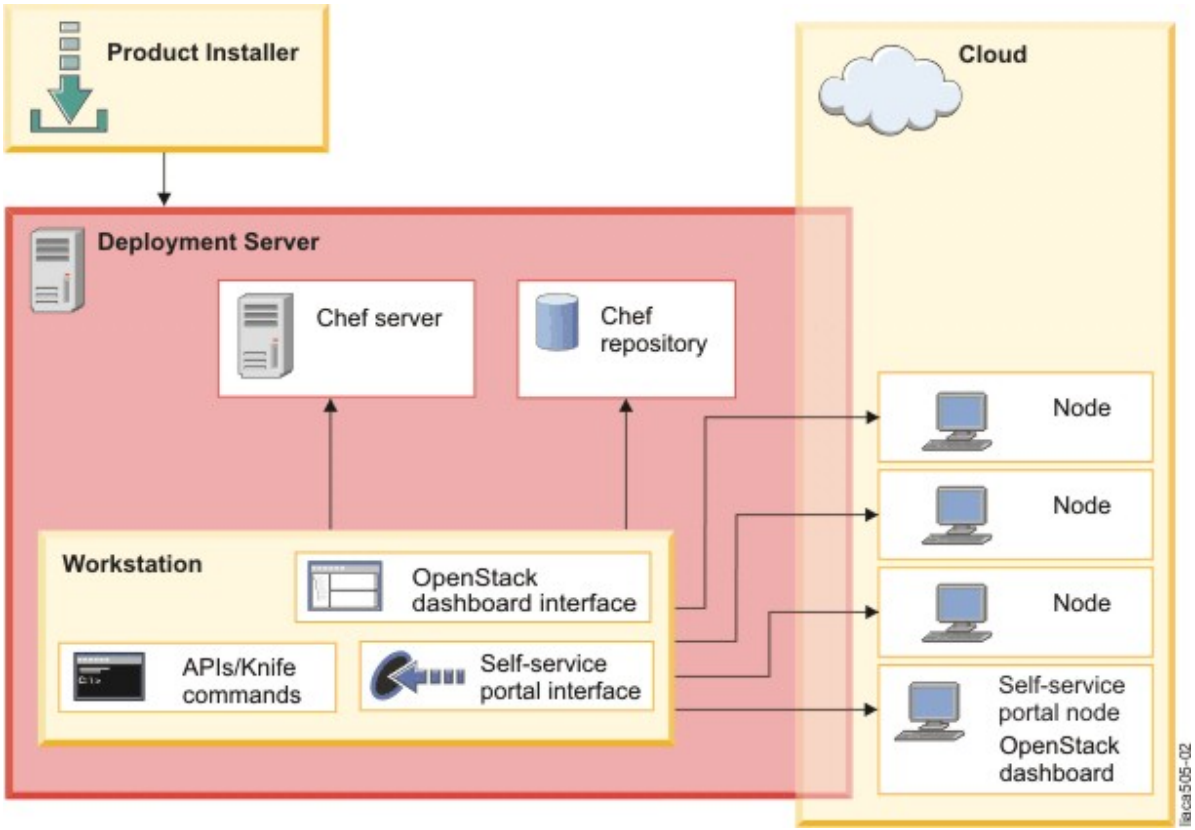

Image source: https://www.ibm.com/support/knowledgecenter/en/SST55W_4.3.0/liaca/liaca_techoverview.html

CMO Deployment at KEK

- OpenStack version
  - Kilo: The latest version supported by CMO
- Used OpenStack components
  - Keystone (Identity)
  - Nova (Compute)
  - Glance (Image Storage)
  - Horizon (Dashboard)
  - Neutron (Networking)
- Compute nodes
  - Scientific Linux 7
  - KVM
  - 75 CPU cores

The Cloud Covers 2 Use Cases

- Batch integration
  - LSF + OpenStack
  - Prepares requested data analysis environment triggered by job submission.
- Self-service provisioning
  - Provides customizable servers for experimental groups.

Batch Integration

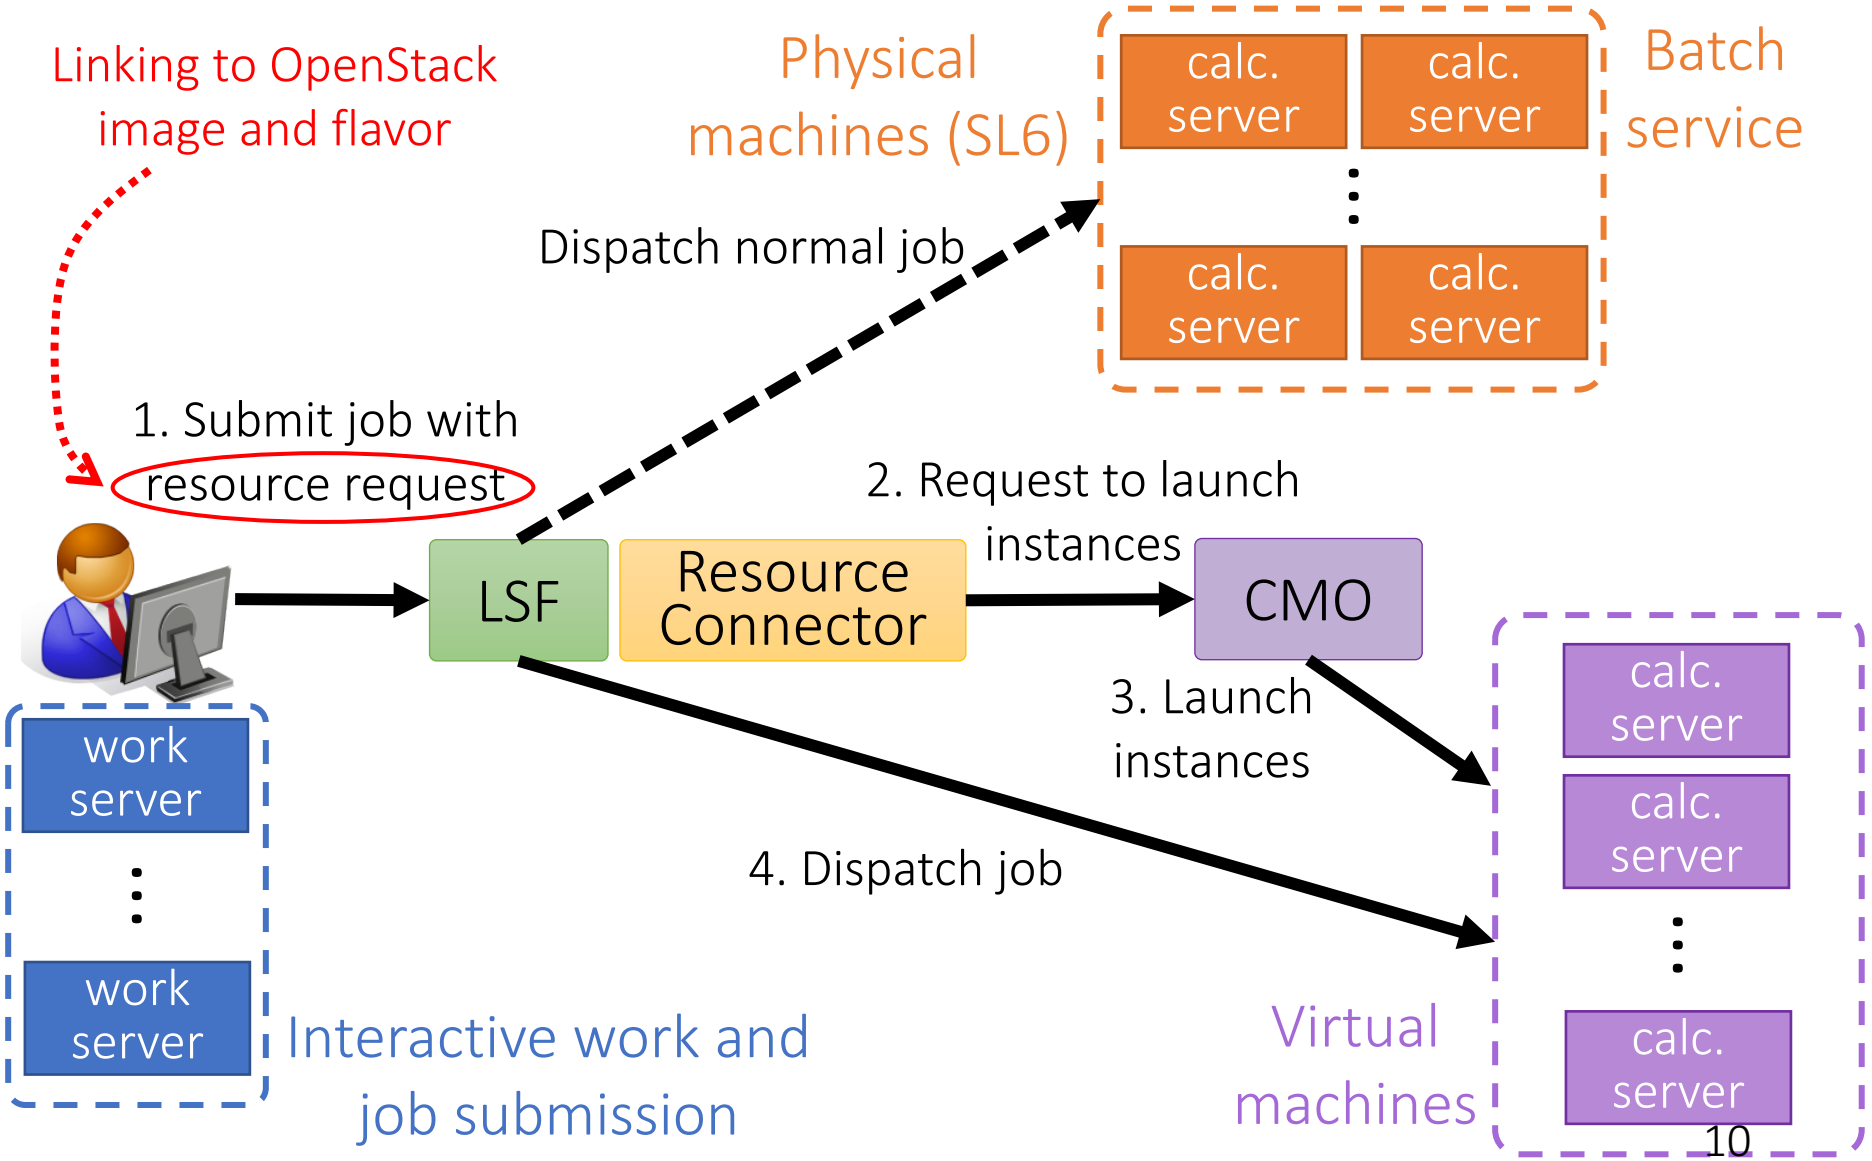

Self-Service Provisioning

- Provide Simplified Service Portal.
- Control allowed actions by OpenStack role.

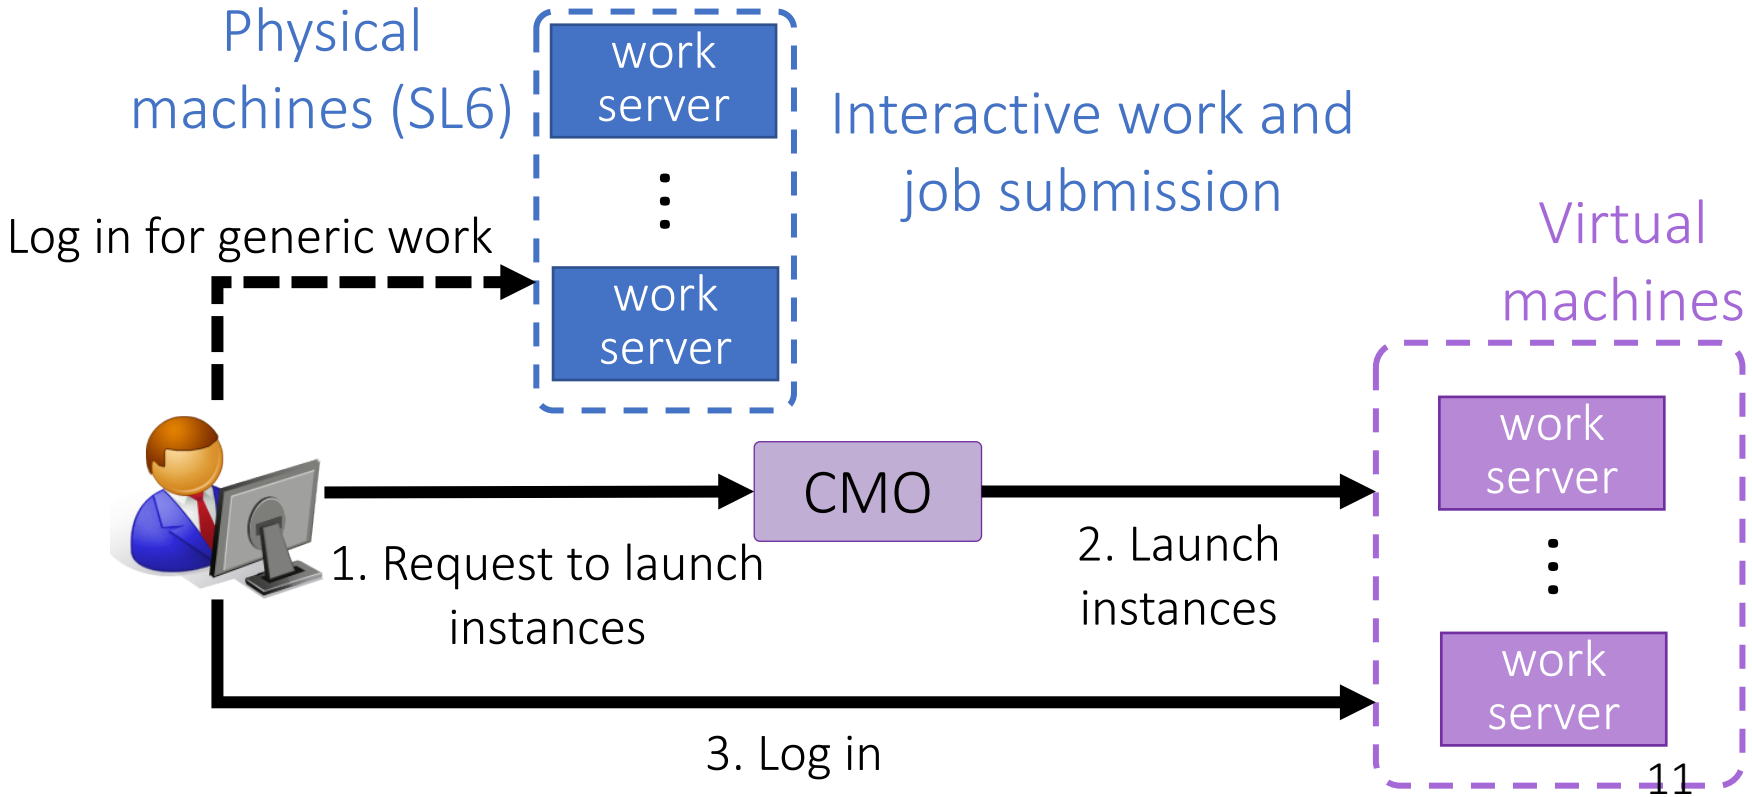

Planned Usage of the Cloud

- Cloud admin prepares base images
  - SL6, CentOS7, Ubuntu16
- Self-service

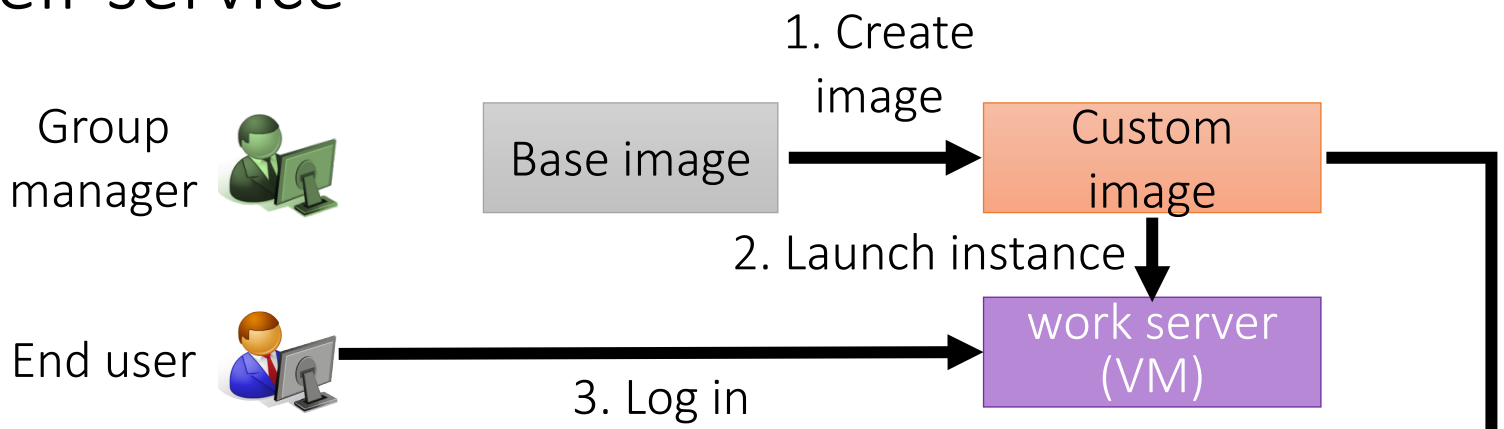

- Batch integrated cloud

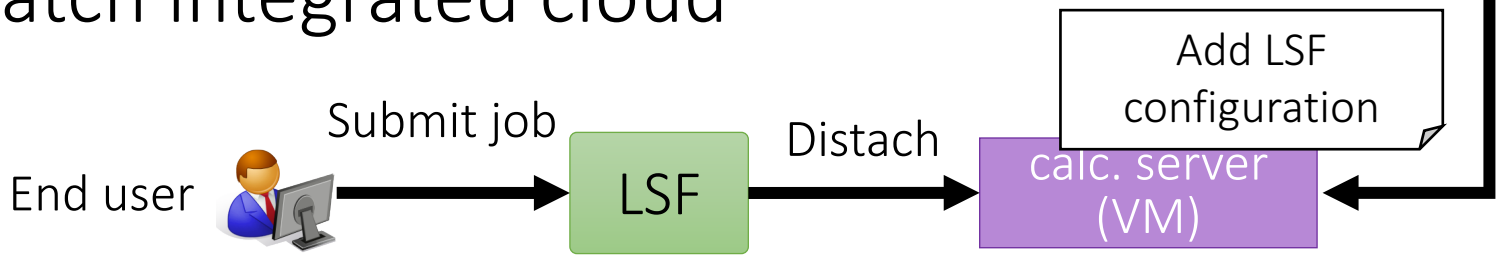

Integration with Existing System: LDAP

- LDAP authentication is used for the cluster.
- Use the LDAP service as OpenStack authentication backend.
- Use the LDAP for Linux accounts on a VM.

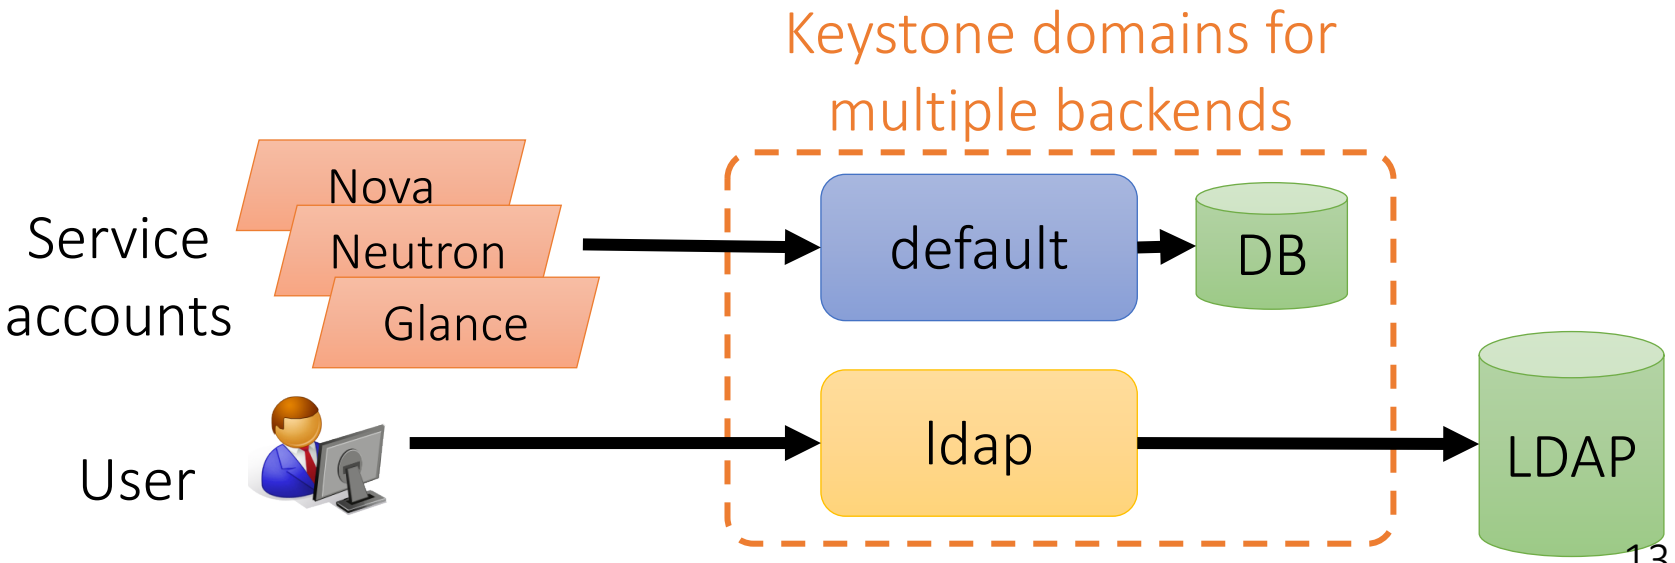

Integration with Existing System: GPFS

- GPFS is used for home directories and group shared directories.
- Don’t GPFS mount from VM to avoid additional GPFS operation.
  - GPFS requires node registration to the cluster.
- Each compute node mounts GPFS and exposes the directories to VM via NFS.

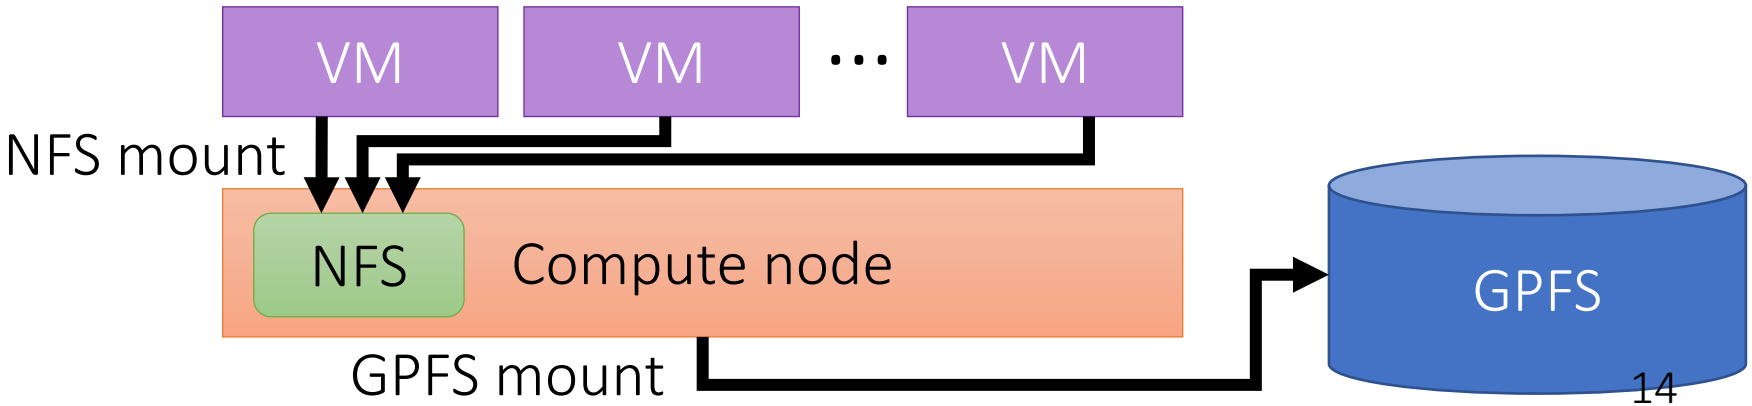

Group based Management: SSH Access Control

- OpenStack project is mapped to Linux group.
- Instance in a project only allows SSH access from users on the project.

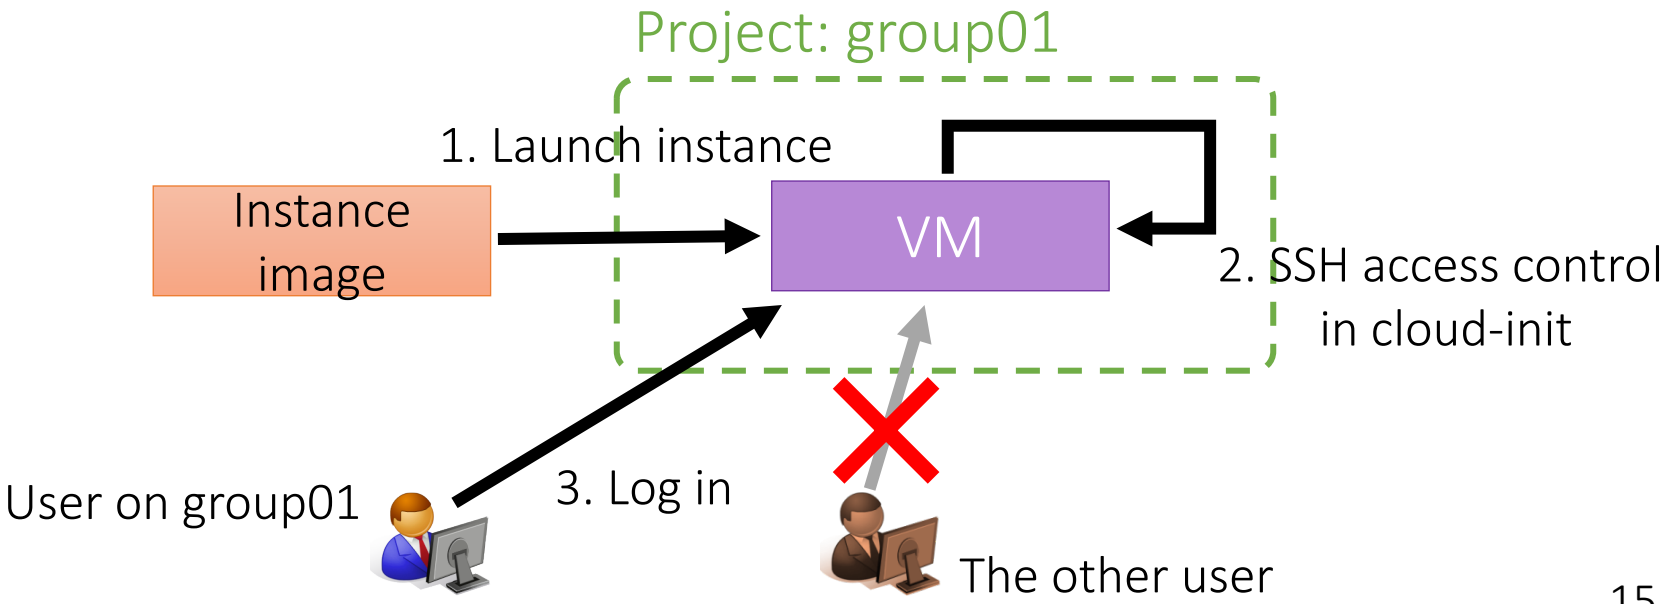

Group based Management: Launch Instances on Specific Nodes

- Plan to partially deploy physical servers provided by a group as compute nodes.
  - The resource is only for the group.
- Use "Host Aggregate" feature.

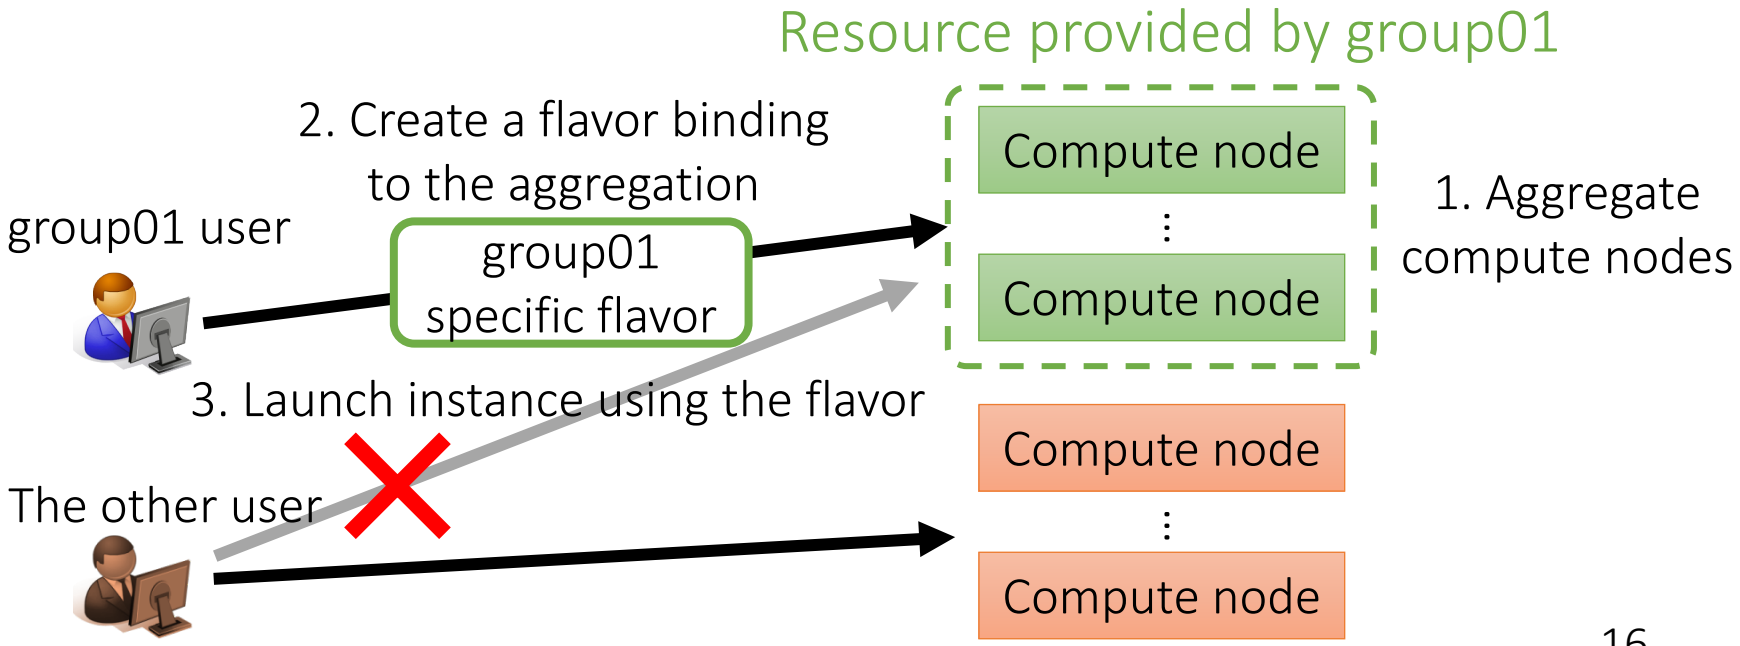

Current Status

Work item	Status
Create common base images for work and calc. instances	Done
Create LSF configuration script for the common images	WIP
Integrate with existing systems (LDAP, GPFS)	WIP
Implement group based management (SSH access & resource)	Done
Test by cloud admins	WIP
Test by group managers	Being started

Near-future Plan: Public Clouds Integration

- Supplement the shortage of existing resource.
- Add temporal resource for intensive data analysis in a certain period.

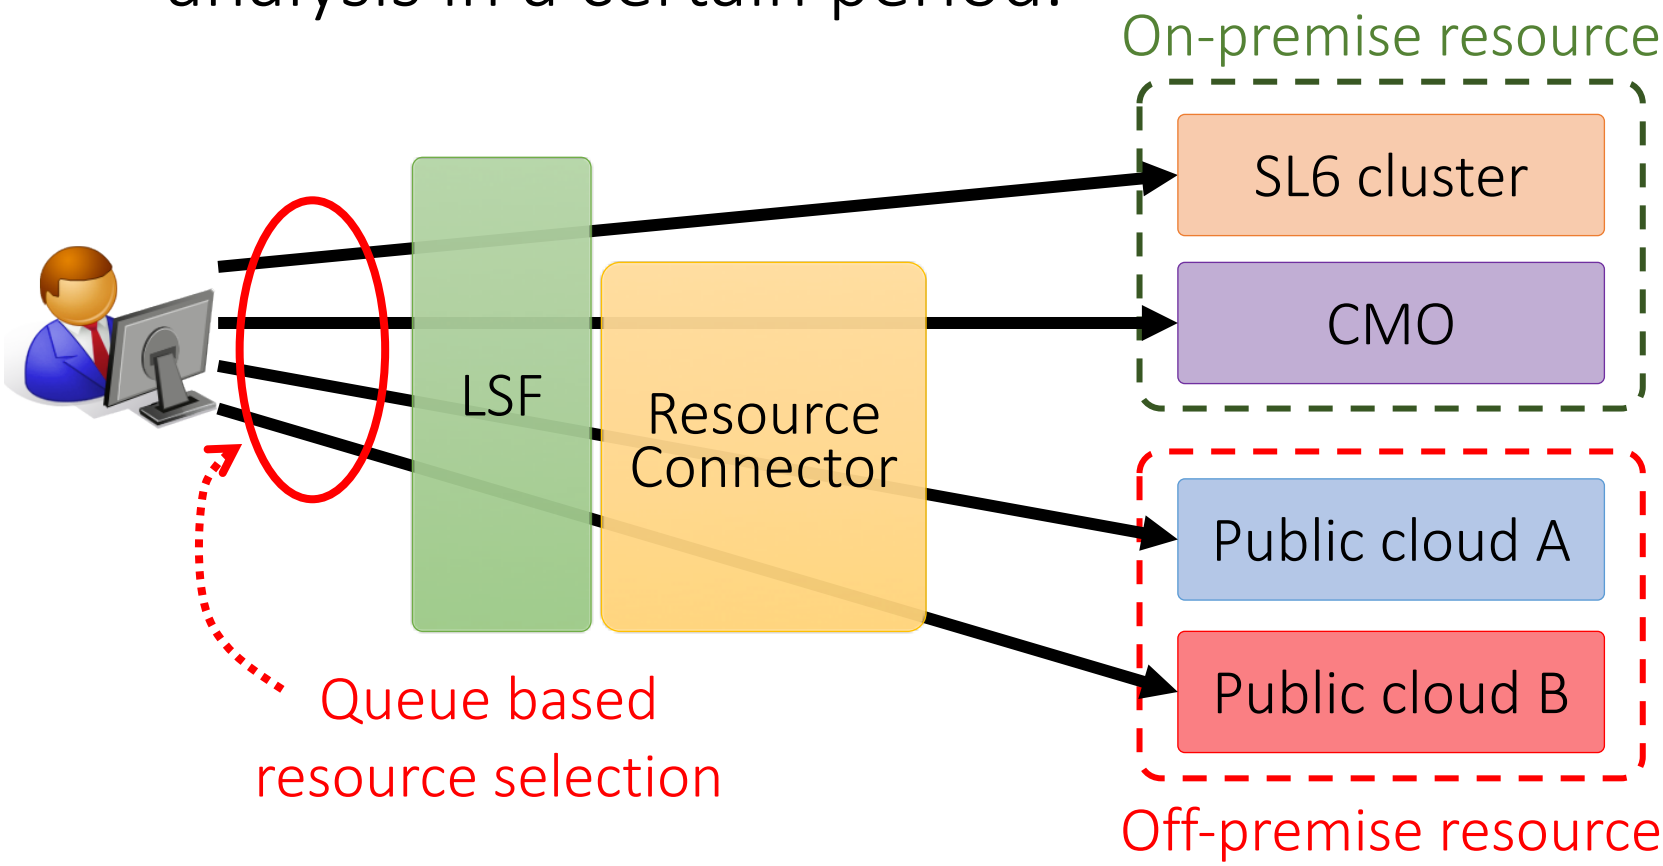

Summary

- CMO based Cloud service at KEK is now test phase.
- Our Cloud will cover 2 use cases:
  - Batch integration
  - Self-service provisioning
- We are integrating the Cloud with existing LDAP and GPFS.
- Group based SSH access control and resource allocation have been implemented.
- We are investigating the way to integrate batch service with public clouds for more flexible workload management.