



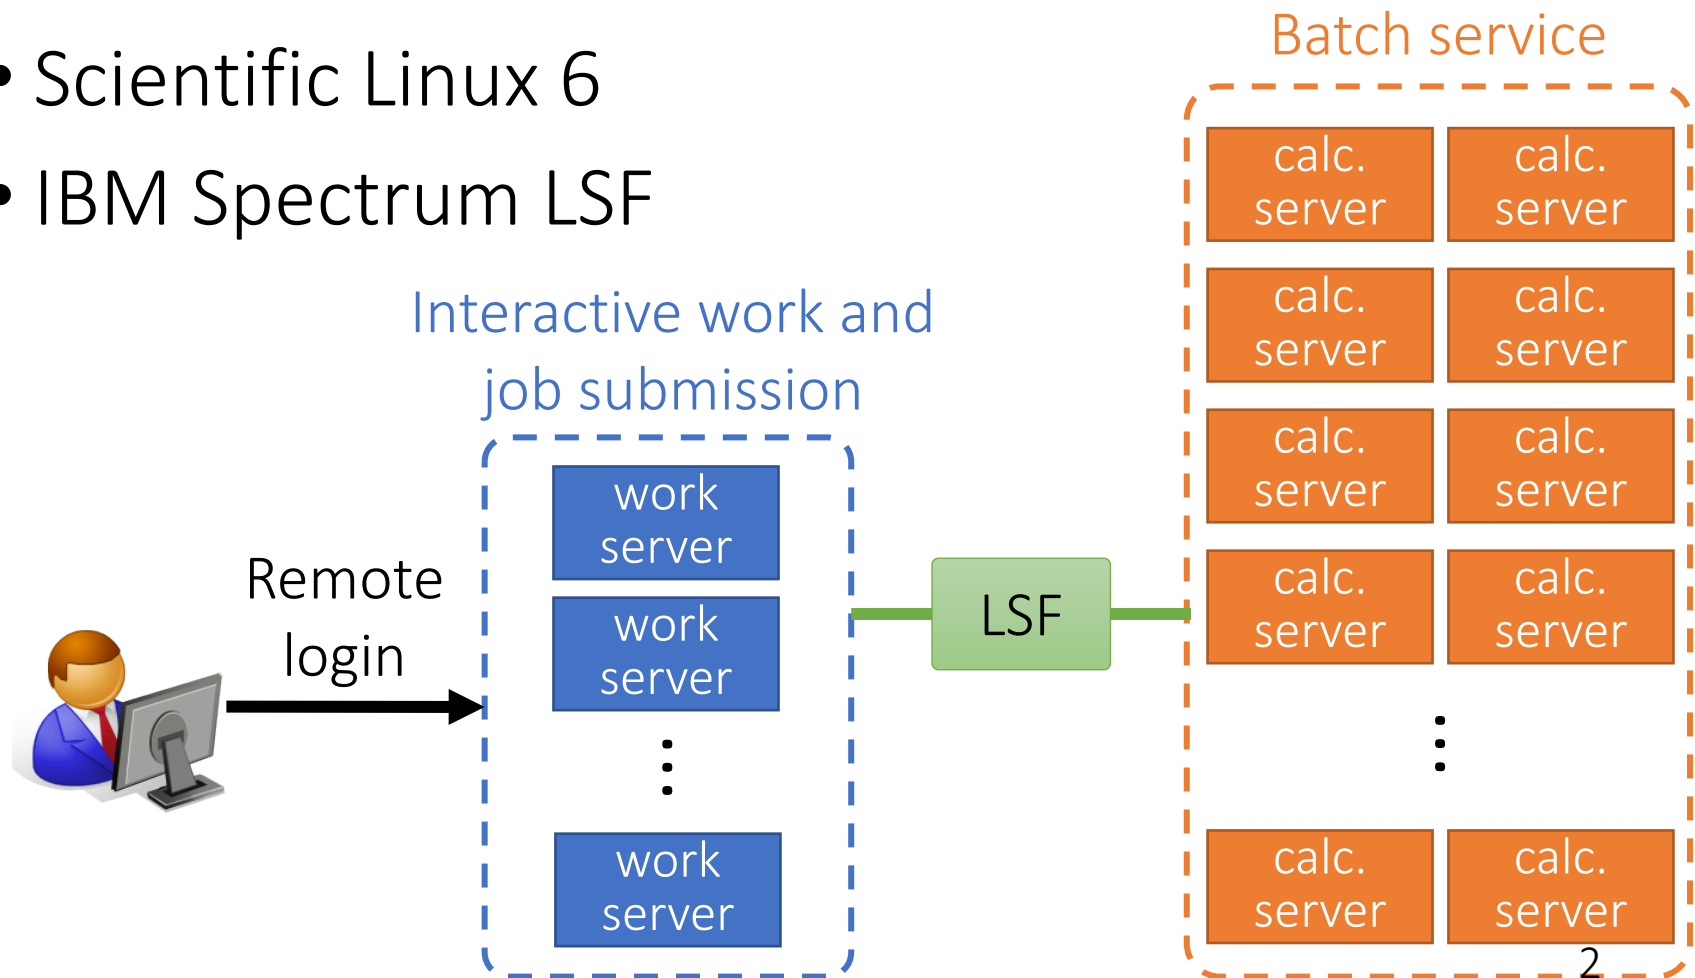
Cloud Deployment at KEK

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

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

The screenshot shows the 'IBM Cloud Manager - Self Service' interface. The left sidebar contains navigation options: 概要 (Overview), インスタンス (Instances), イメージ・テンプレート (Image Templates), and 要求 (Requests). The main content area is titled 'イメージ・テンプレート' and shows the path 'base-cent7-05' with a '起動' (Launch) button. Below this, it says 'インスタンスの起動: base-cent7-05'. There are input fields for '名前' (Name), 'プロジェクト' (Project) set to 'Public', 'インスタンス数' (Number of instances) set to '1', and 'フレーバー' (Flavor) set to 'c02-m008G'. A 'プロジェクト使用状況' (Project Usage) section shows three progress indicators: 'インスタンス' (Instances) at 6/10, '仮想 CPU' (Virtual CPU) at 12/20, and '/50 GB メモリー' (50 GB Memory) at 47.3 GB. Below these are resource requirements: '仮想 CPU' (Virtual CPU) 2, 'メモリー' (Memory) 7.9 GB, 'ストレージ' (Storage) 50 GB, and 'スワップ' (Swap) --. There is also a section for 'アクセスおよびセキュリティ' (Access and Security) with a 'セキュリティ・グループ' (Security Group) dropdown set to '使用可能' (Available).

OpenStack dashboard for cloud admins

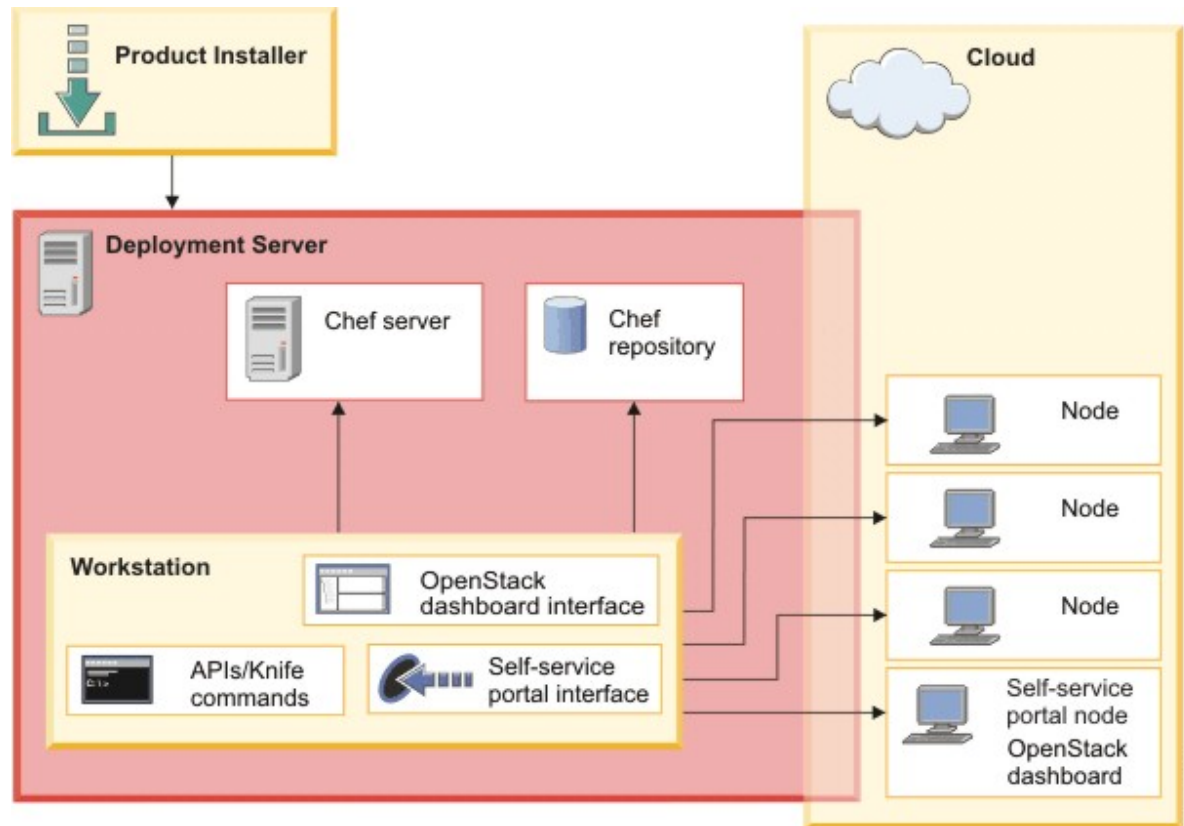
The screenshot shows the 'IBM Cloud Manager - Dashboard' interface for an administrator. The left sidebar has a navigation menu with categories: PROJECT, ADMIN, System, Overview, Resource Usage, Hypervisors, Host Aggregates, Instances, Volumes, Flavors, Images, Networks, Routers, Defaults, Metadata Definitions, System Information, IDENTITY, and RESOURCE SCHEDULER. The main content area is titled 'Overview' and 'Usage Summary'. It includes a 'Select a period of time to query its usage:' section with 'From' and 'To' date pickers (both set to 2017-02-01) and a 'Submit' button. Below this, it displays usage statistics: 'Active Instances: 9 Active RAM: 70.9GB This Period's VCPU-Hours: 2675.63 This Period's GB-Hours: 66890.71 This Period's GB-Hours: 10788133.58'. A 'Usage' table is shown with a 'Download CSV Summary' link. The table has columns for Project Name, VCPUs, Disk, RAM, VCPU Hours, Disk GB Hours, and Memo. The data rows are: admin (4 VCPUs, 100GB Disk, 15.8GB RAM, 297.29 VCPU Hours, 14864.60 Disk GB Hours, 23973 Memo), geant4 (2 VCPUs, 50GB Disk, 7.9GB RAM, 148.65 VCPU Hours, 7432.30 Disk GB Hours, 11986 Memo), crc (2 VCPUs, 50GB Disk, 7.9GB RAM, 148.65 VCPU Hours, 7432.30 Disk GB Hours, 11986 Memo), and Public (10 VCPUs, 250GB Disk, 39.4GB RAM, 743.23 VCPU Hours, 37161.51 Disk GB Hours, 59934 Memo). The footer of the table says 'Displaying 4 items'.

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.

CMO: Simplified Cloud deployment by Chef

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



CMO Deployment at KEK

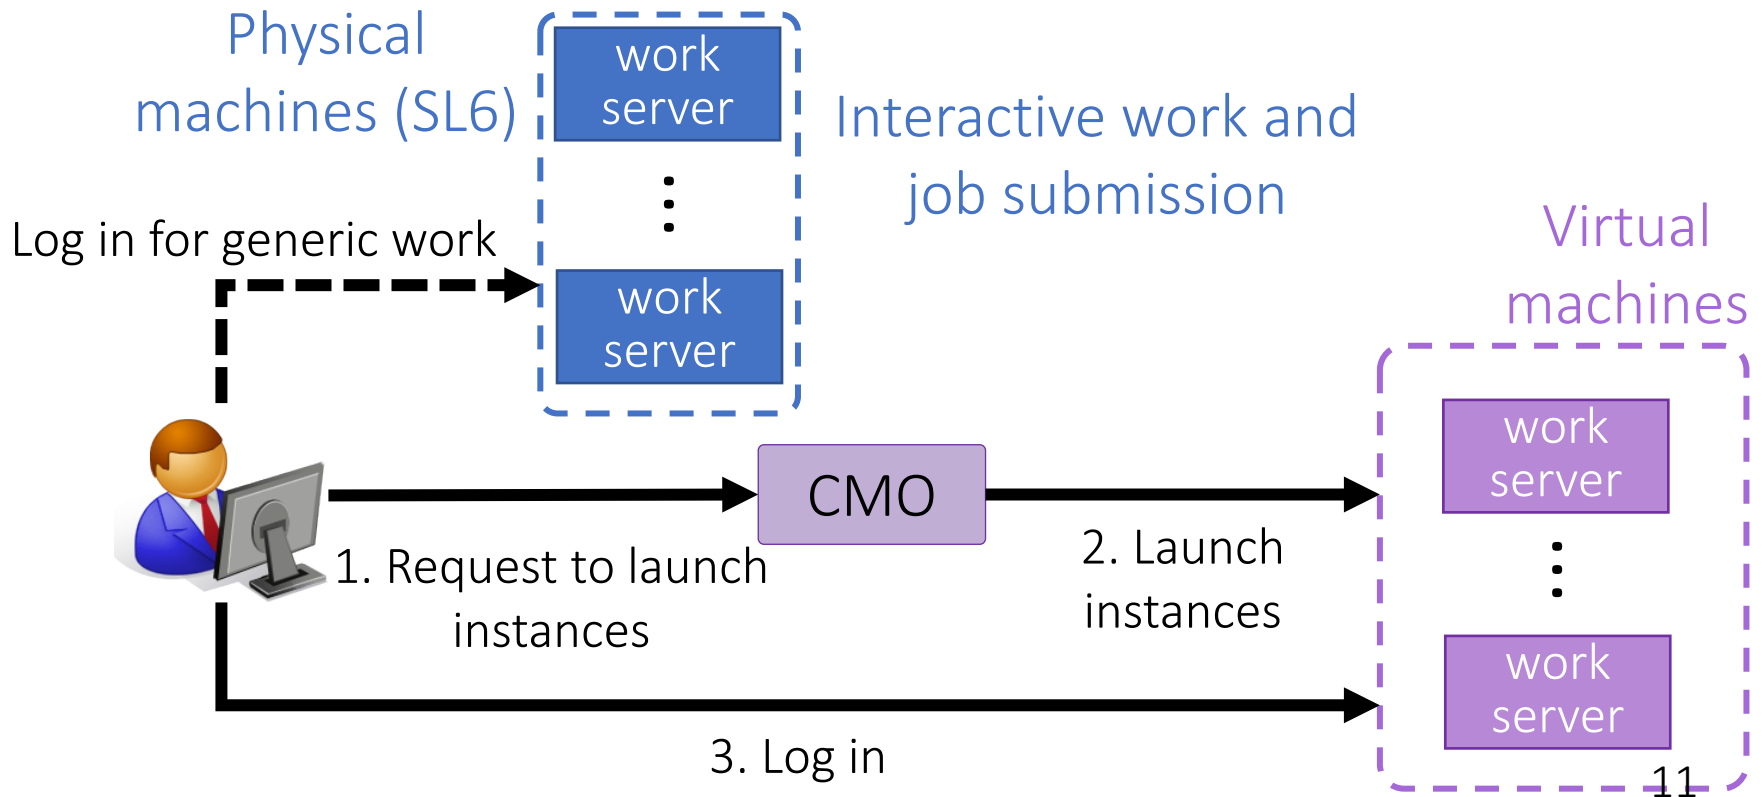
- OpenStack version
 - Kilo: The latest version supported by CMO
- Used OpenStack components
 - Keystone (Identity)
 - Nova (Compute)
 - Glance (Image Storage)
 - Horizon (Dashboard)
 - Neutron (Networking)
- Compute nodes
 - Scientific Linux 7
 - KVM
 - 75 CPU cores

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.


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

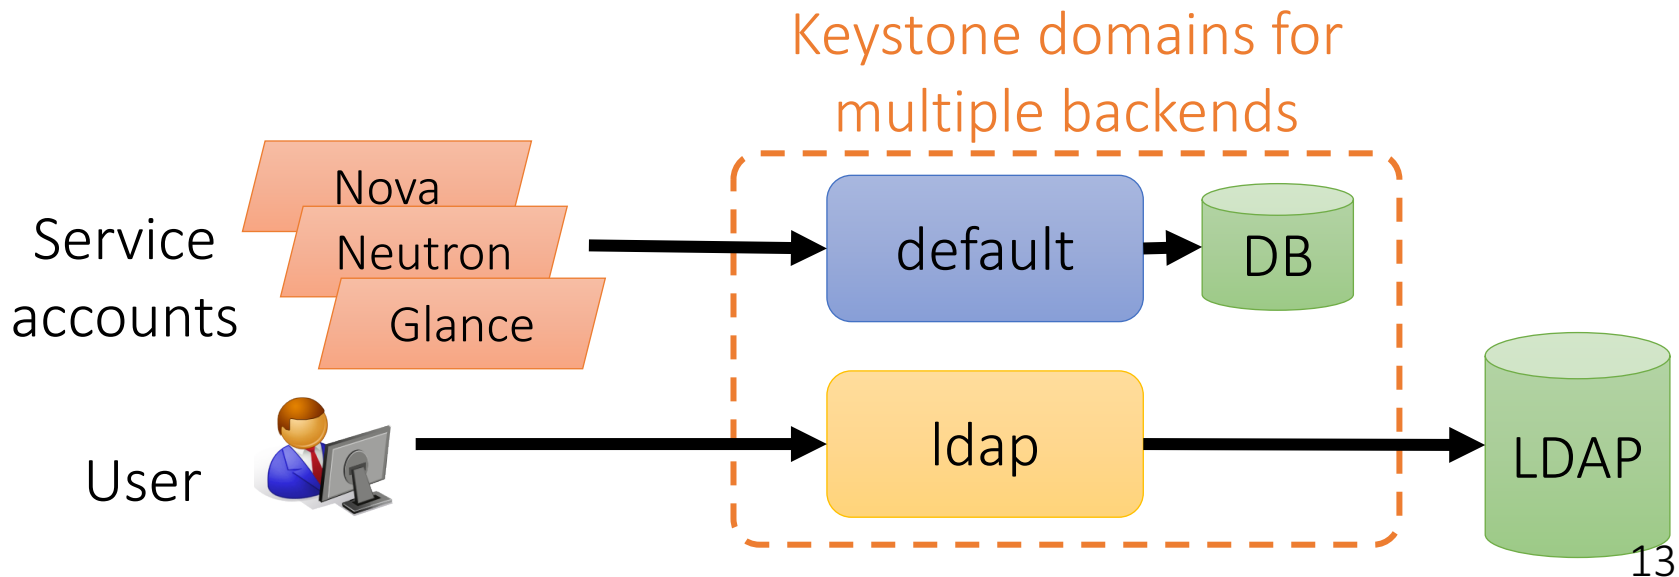


- Batch integrated cloud




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.



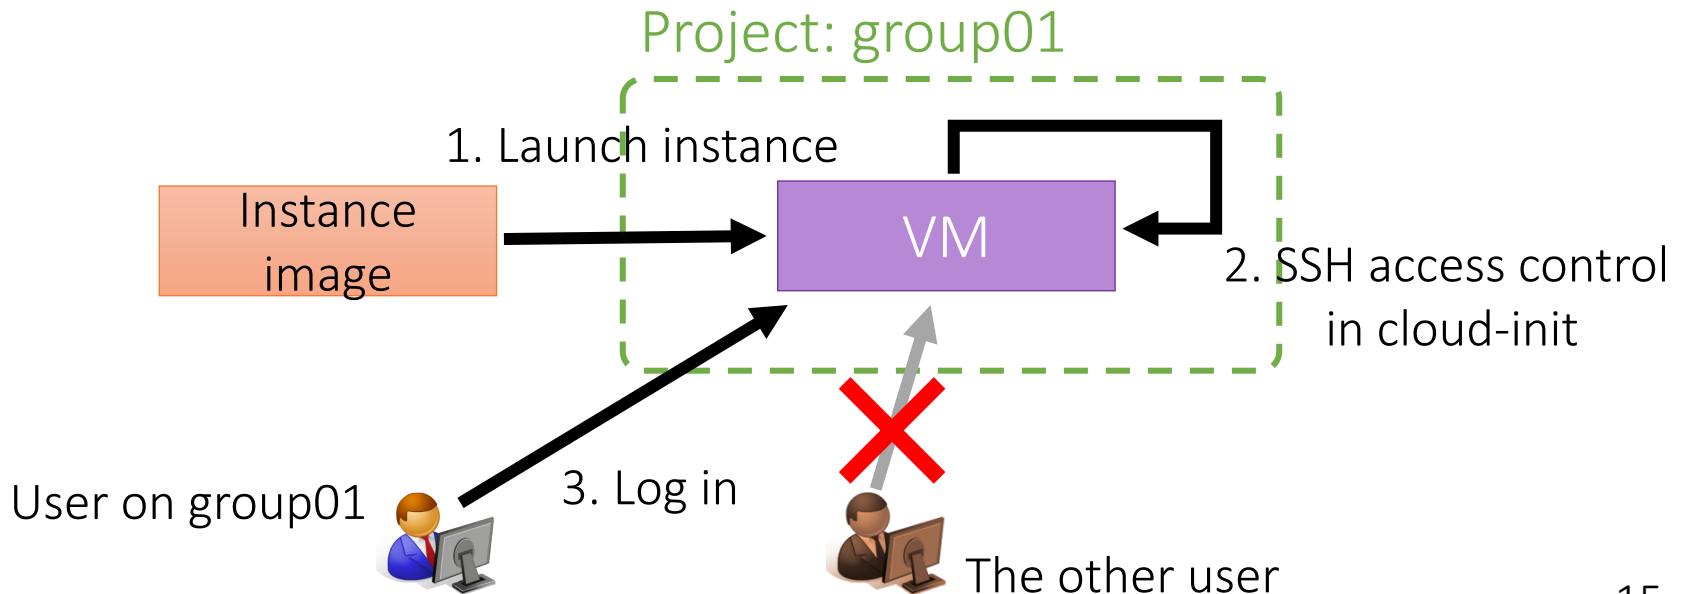
Integration with Existing System: GPFS

- 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.



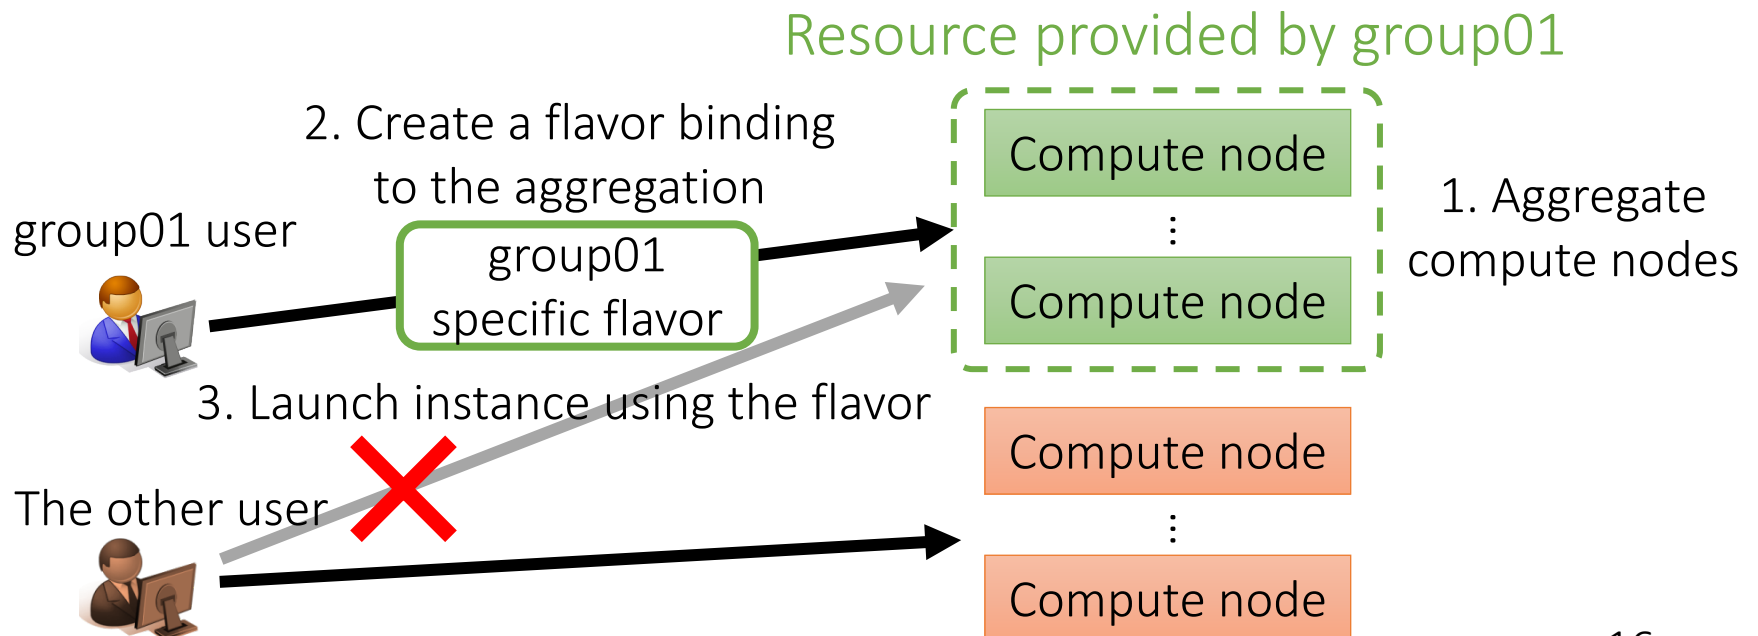
Group based Management: SSH Access Control

- 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.

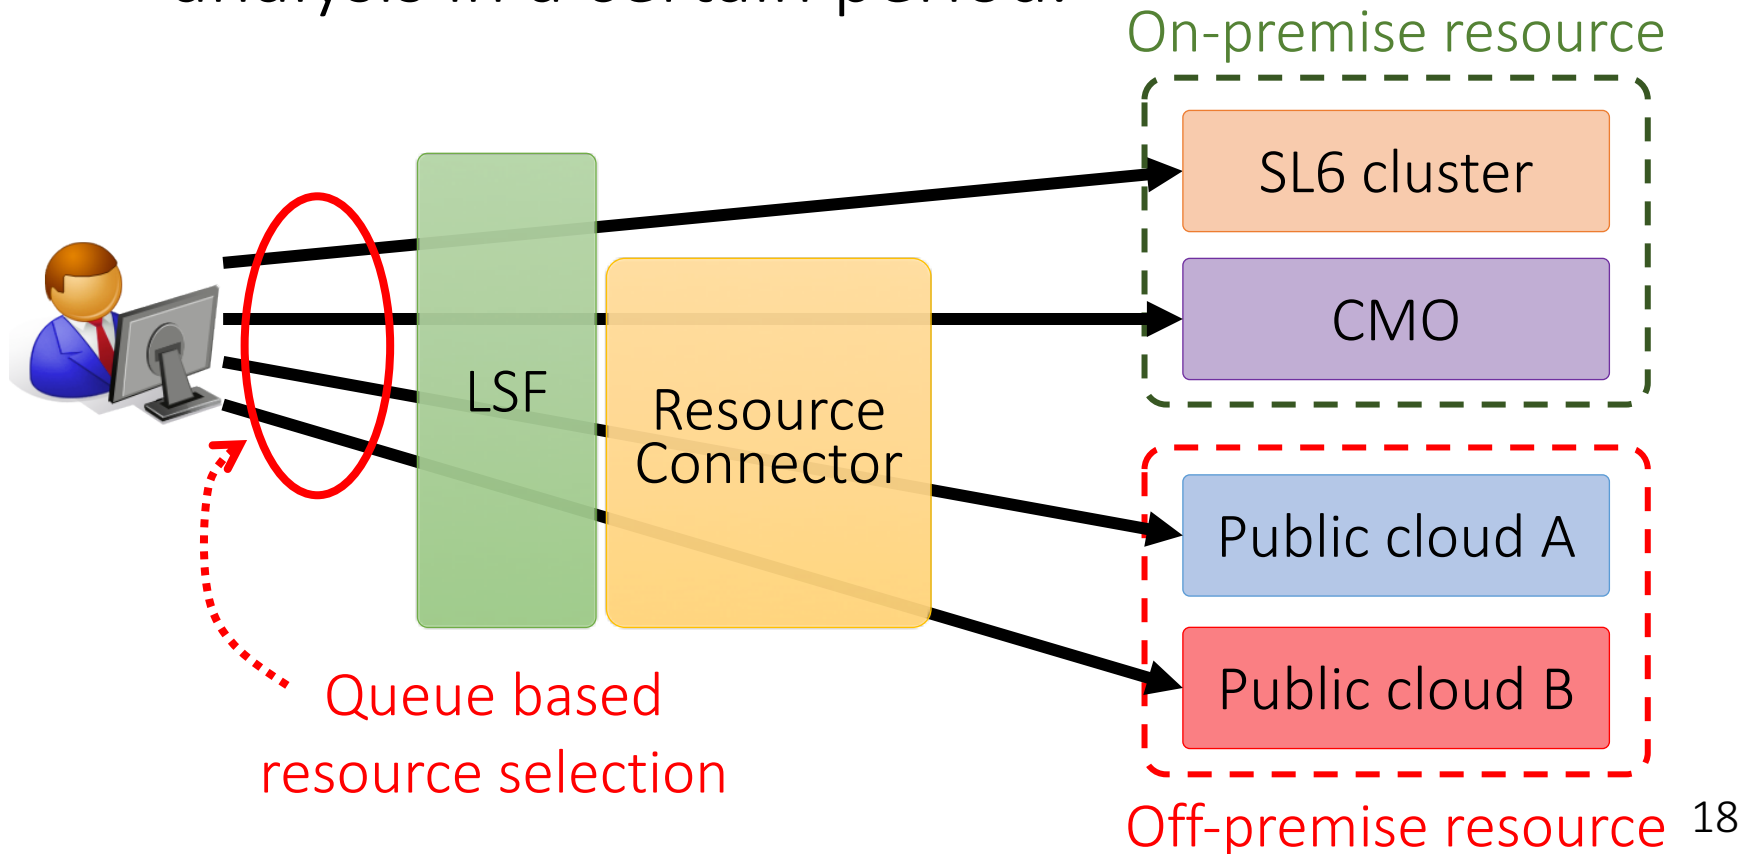


Current Status

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.