

Building and operating a HPC Cloud

XIAO Haili (haili@sccas.cn)

CNGrid Operation Center
CNIC, CAS

ISGC 2017, March 9, Taipei

Contents

- (HPC) grid vs. (HPC) cloud
- CNGrid services and operations
- Summary

(HPC) grid vs. (HPC) cloud

- Public grids provide huge computing resources to scientific users all around the world since 20 years ago.
 - WLCG, OSG, XSEDE2, etc.
 - (HPC) Clusters are connected.
- Public clouds in the last 10 years: IaaS/PaaS/SaaS.
 - Amazon, Microsoft, Google and Ali, etc.
 - Application infrastructures in the cloud
 - VMs, storage, (No)SQL, CDN, LB, security, analytics...
- Public clouds show their great interests in HPC area besides traditional markets recently.
 - Building a cluster with (a lot) VMs

AWS HPC Cloud

- (AWS defined) High Performance Computing
 - allows scientists and engineers to solve *complex, compute-intensive* and *data-intensive* problems.
 - often requires *high network performance, fast storage, large amounts of memory, very high compute capabilities*, or all of these.
- High performance computing in the (AWS) cloud
 - Compute-optimized (C family) instance
 - Memory-optimized (R and X families) instances
 - GPU (P family) and FPGA (F Family) instances
 - General-purpose (M family) instance

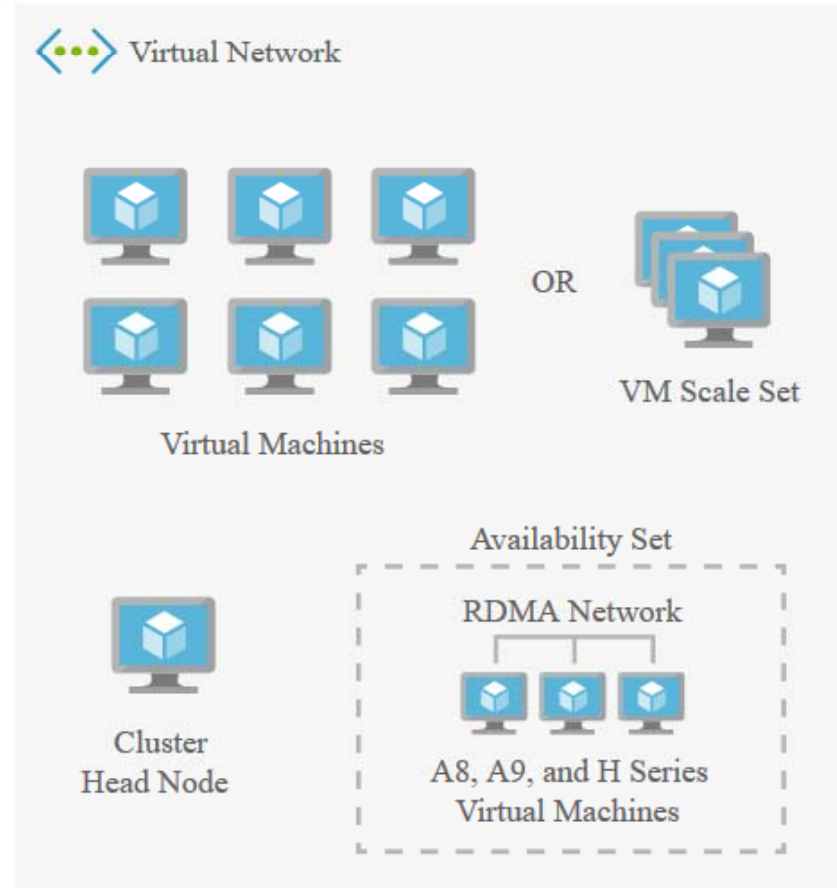
- Azure



ARM template



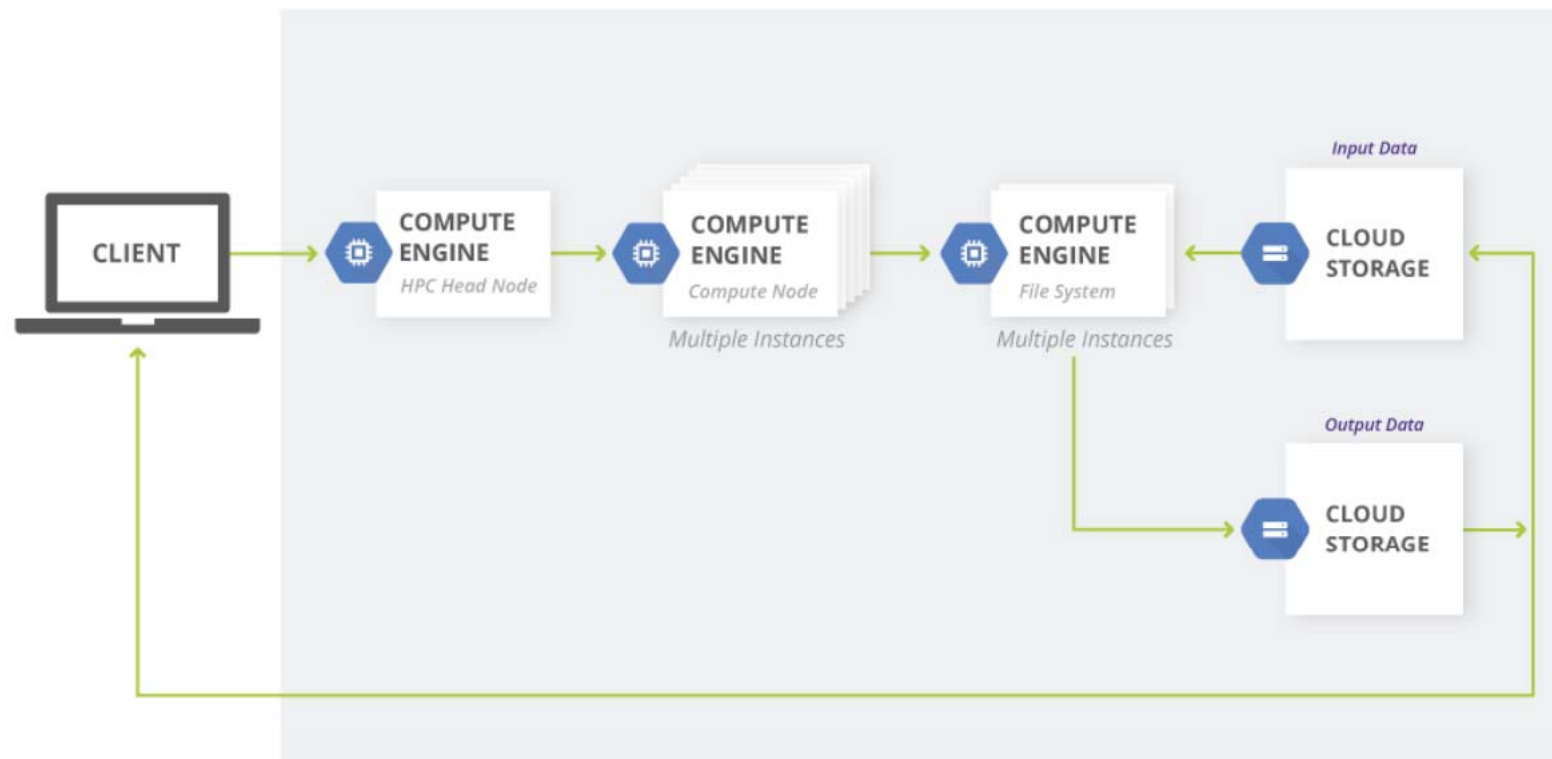
Script file



Google Cloud Platform



- Architecture: High Performance Computing
 - Compute Engine VMs + Cloud Storage = running HPC workloads



<https://cloud.google.com/solutions/architecture/highperformancecomputing>

Alibaba

- Alibaba Cloud HPC
 - G4 instance
 - CPU: Intel Xeon E5 v4 CPU 32 cores
 - GPU: Nvidia Tesla M40 x2
 - G2 instance
 - CPU: Intel Xeon E5 v2 CPU 16 cores
 - GPU: Nvidia Tesla K40 x2



| Rank | Site | System | Cores | Rmax (TFlop/s) | Rpeak (TFlop/s) | Power (kW) |
|------|--|--|------------|-------------------|--------------------|---------------|
| 1 | National Supercomputing Center in Wuxi China | Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRPCPC | 10,649,600 | 93,014.6 | 125,435.9 | 15,371 |
| 2 | National Super Computer Center in Guangzhou China | Tianhe-2 (MilkyWay-2) - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT | 3,120,000 | 33,862.7 | 54,902.4 | 17,808 |
| 3 | DOE/SC/Oak Ridge National Laboratory United States | Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc. | 560,640 | 17,590.0 | 27,112.5 | 8,209 |
| 4 | DOE/NNSA/LLNL United States | Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM | 1,572,864 | 17,173.2 | 20,132.7 | 7,890 |
| 5 | DOE/SC/LBNL/NERSC United States | Cori - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect Cray Inc. | 622,336 | 14,014.7 | 27,880.7 | 3,939 |
| 6 | Joint Center for Advanced High Performance Computing Japan | Oakforest-PACS - PRIMERGY CX1640 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path Fujitsu | 556,104 | 13,554.6 | 24,913.5 | 2,719 |
| 7 | RIKEN Advanced Institute for Computational Science (AICS) Japan | K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu | 705,024 | 10,510.0 | 11,280.4 | 12,660 |
| 8 | Swiss National Supercomputing Centre (CSCS) Switzerland | Piz Daint - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 Cray Inc. | 206,720 | 9,779.0 | 15,988.0 | 1,312 |
| 9 | DOE/SC/Argonne National Laboratory United States | Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM | 786,432 | 8,586.6 | 10,066.3 | 3,945 |
| 10 | DOE/NNSA/LANL/SNL United States | Trinity - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc. | 301,056 | 8,100.9 | 11,078.9 | 4,233 |

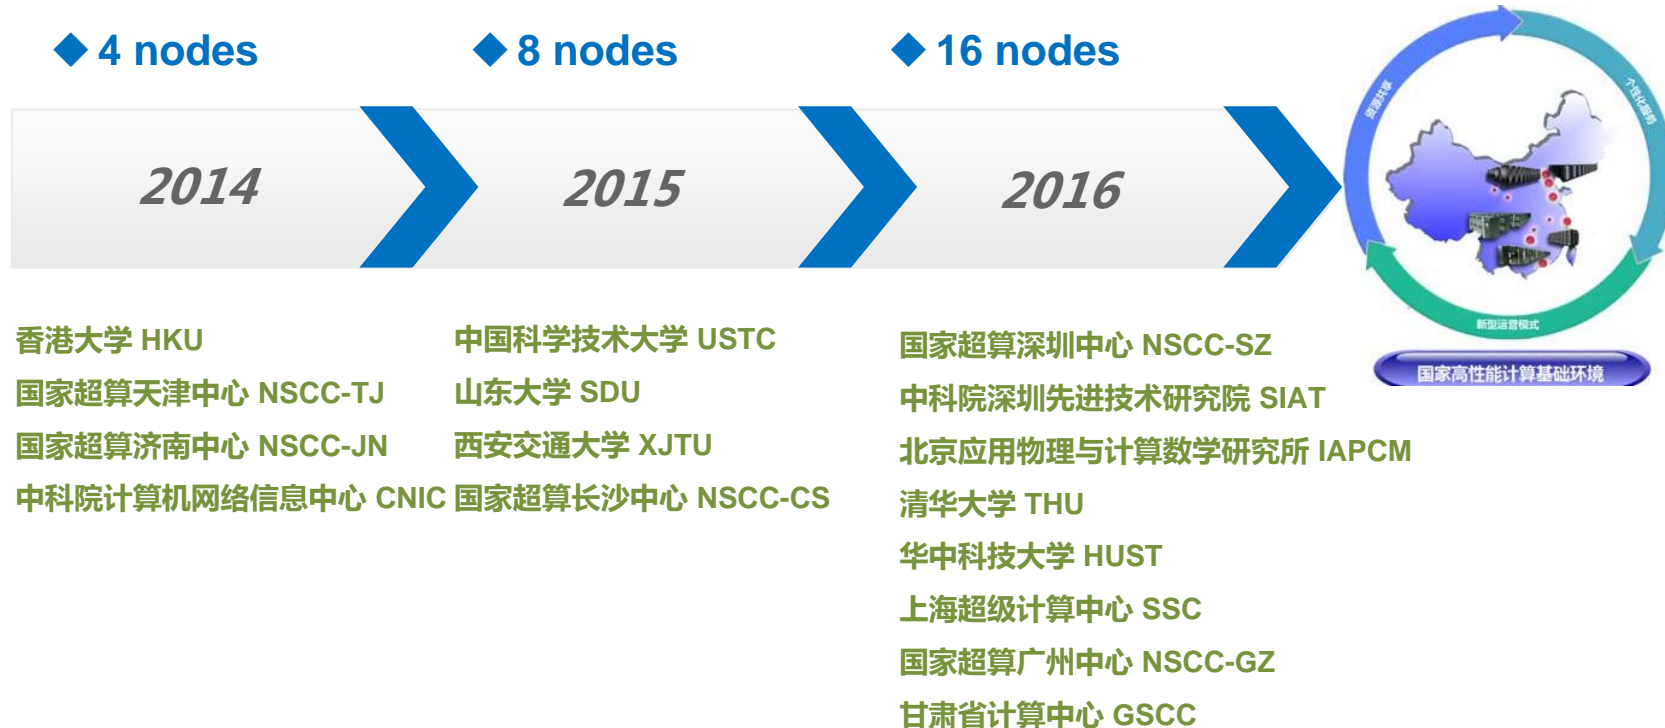
CNGrid

- One of the biggest computing grid in the world.
- It has been providing grid computing services for scientific users for more than 15 years.



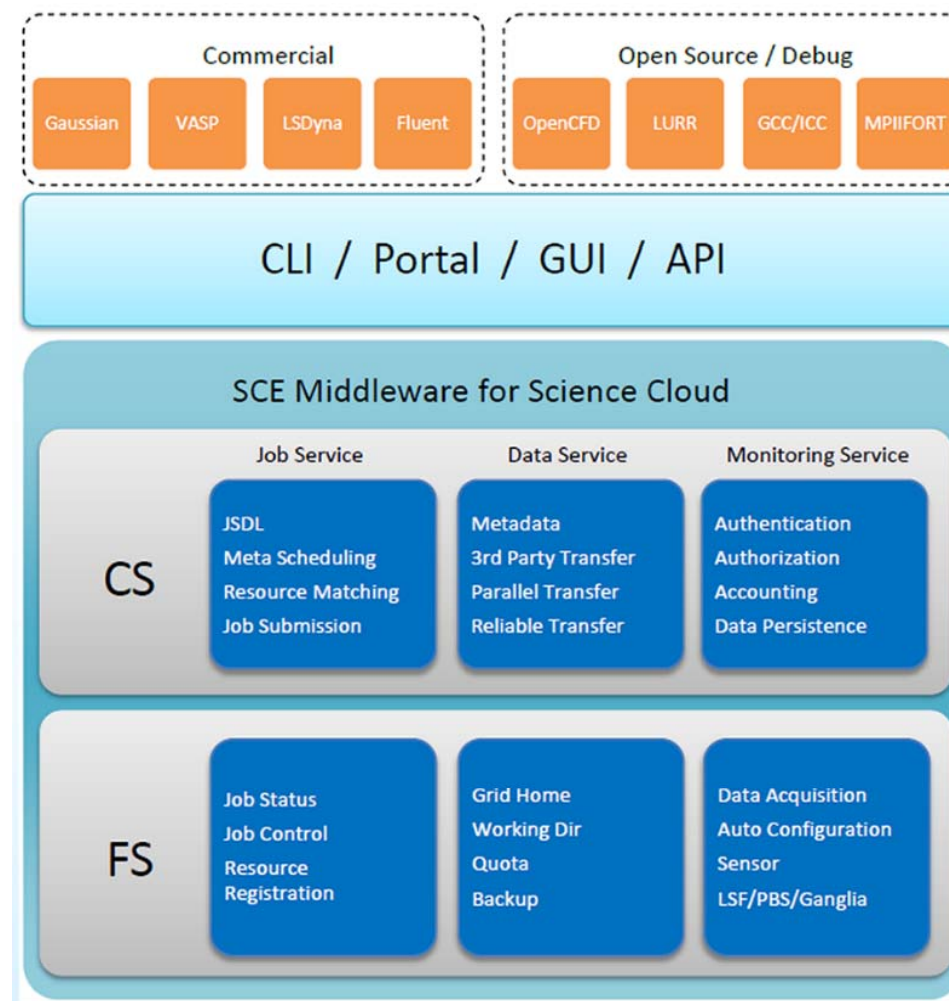
CNGrid Nodes

- 16 grid nodes
 - 8 HPC centers, 6 universities, 2 institutes
 - Aggregating resources:
 - Until Dec 2016: 900+ users, 700k jobs, 150m cpu hrs



SCE - Middleware for HPC Cloud

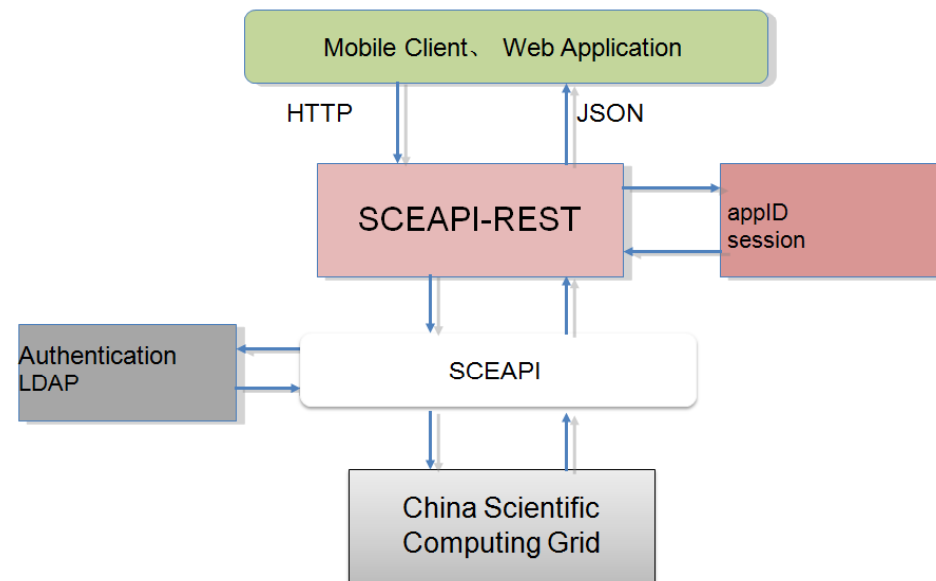
- Developed by SCCAS
- SCE
 - Scientific computing
 - Lightweight
 - Stable
- Diveristy
 - CLI
 - Portal
 - GUI
 - API



International Patent
(PCT/CN2011/071640)

SCEAPI - HPC Cloud API based on SCE


- RESTful API
 - Lightweight Web Service
 - OS independent
 - Windows, Linux
 - iOS, Android
 - Language independent
 - Java, C/C++
 - PHP, Python, Ruby
 - ...
 - Support App. Community
 - Support mobile APPs



SCEAPI - HPC Cloud API



Software
as a service



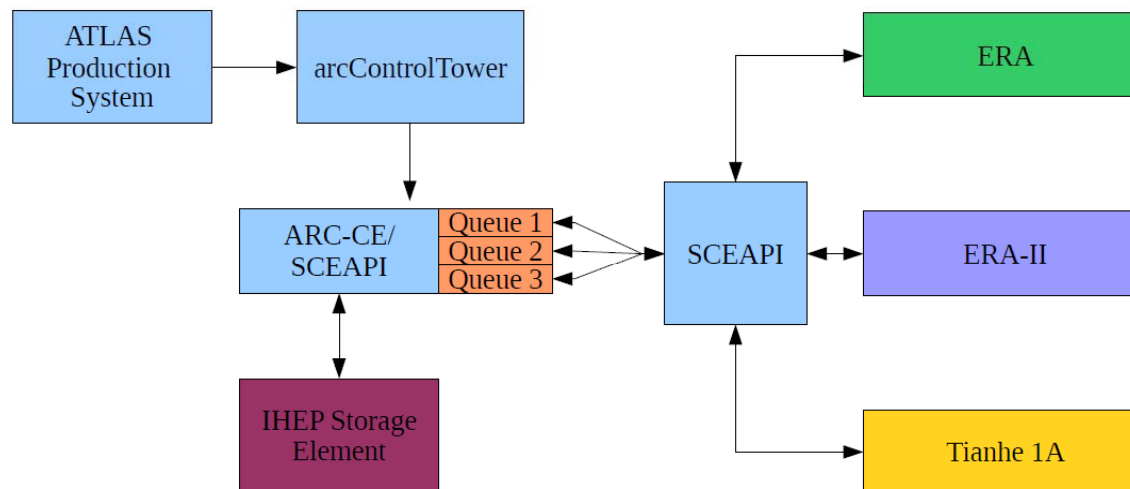
Platform
as a service



Infrastructure as a
service

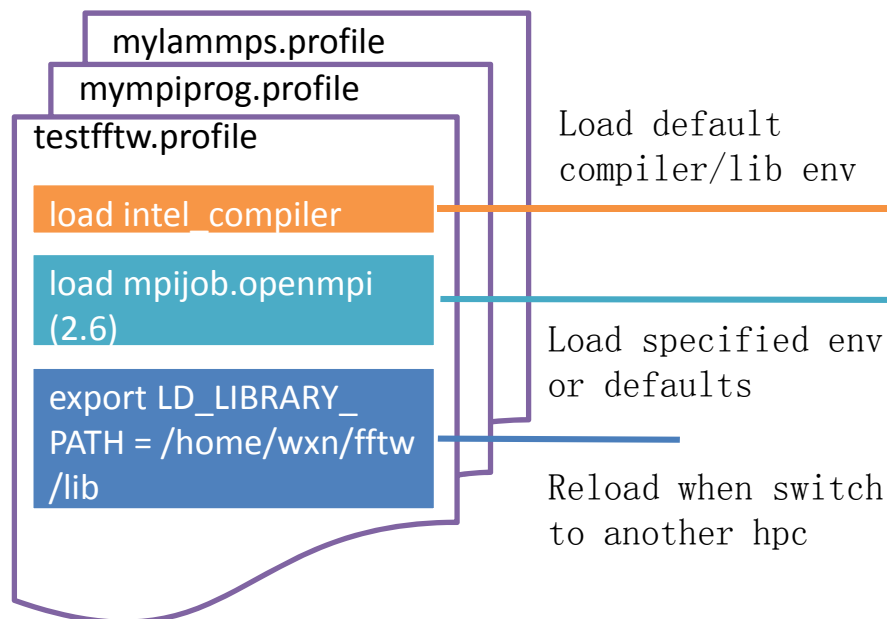
CNGrid running ATLAS jobs

Architecture

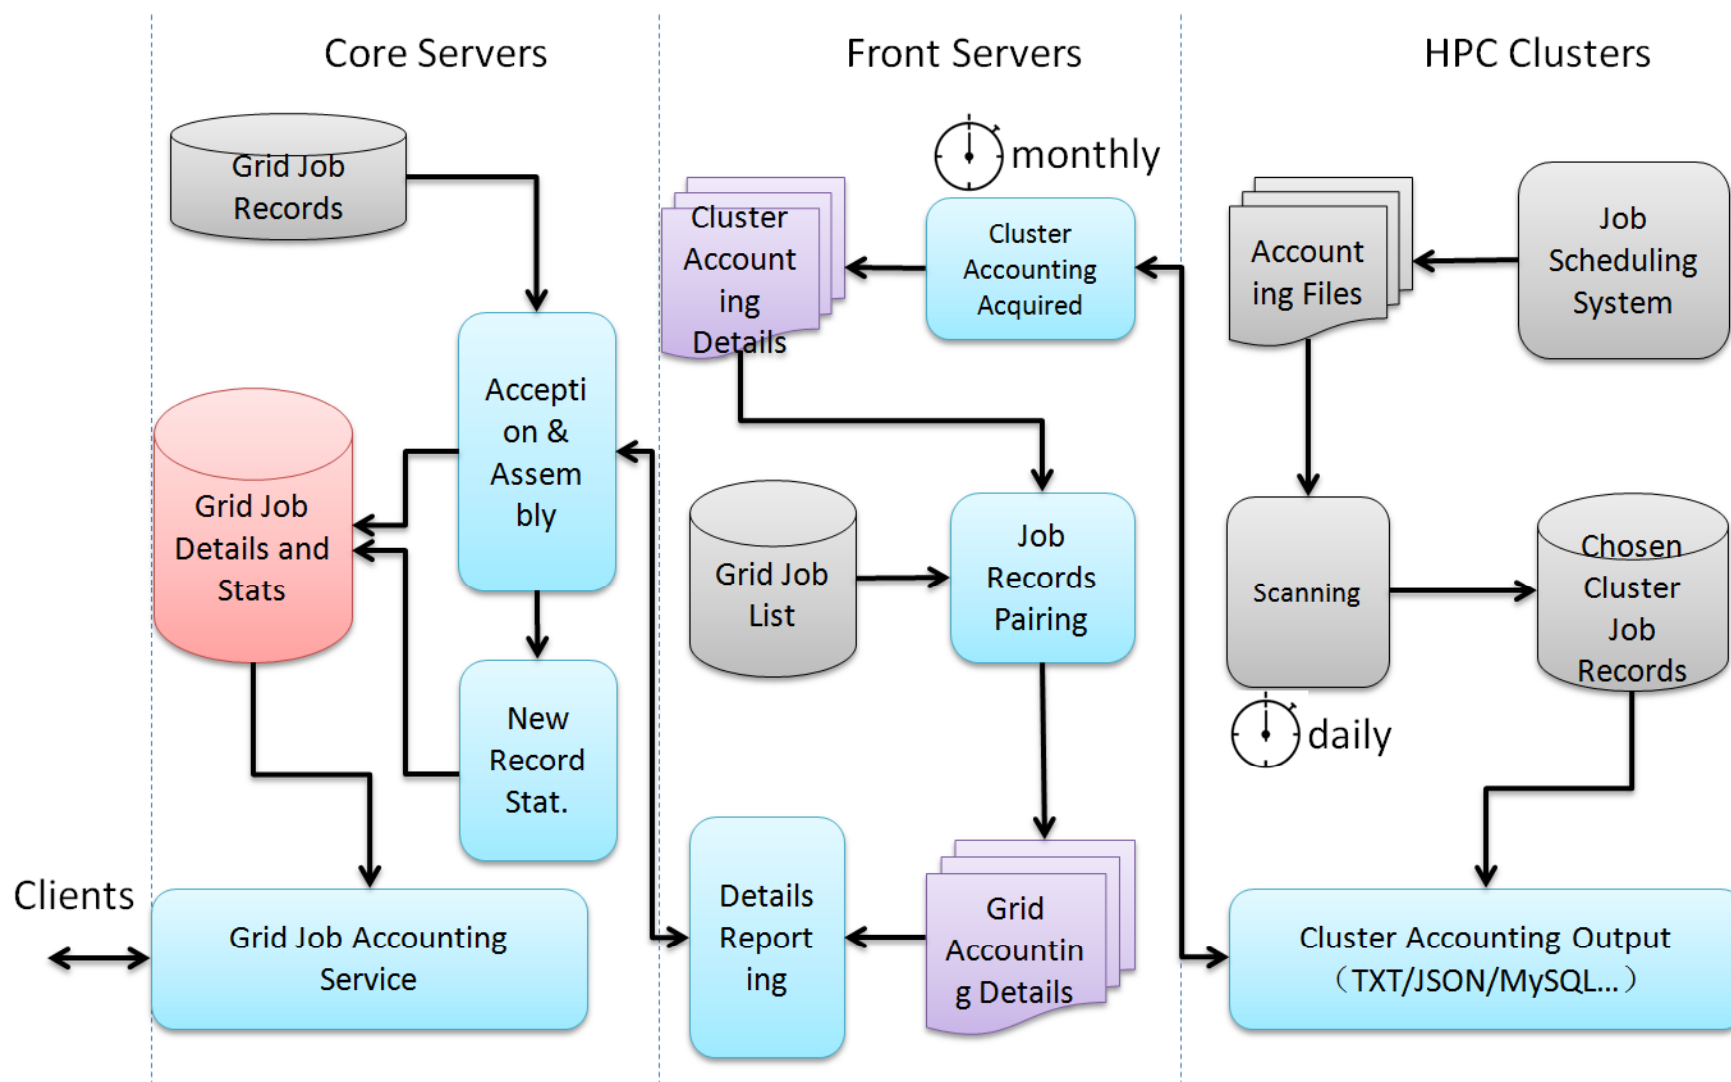


Command Line Users

- maintain & load env. variables
 - Compiler settings
 - Application specified env.
 - Self-defined variables
 - Effective after setting

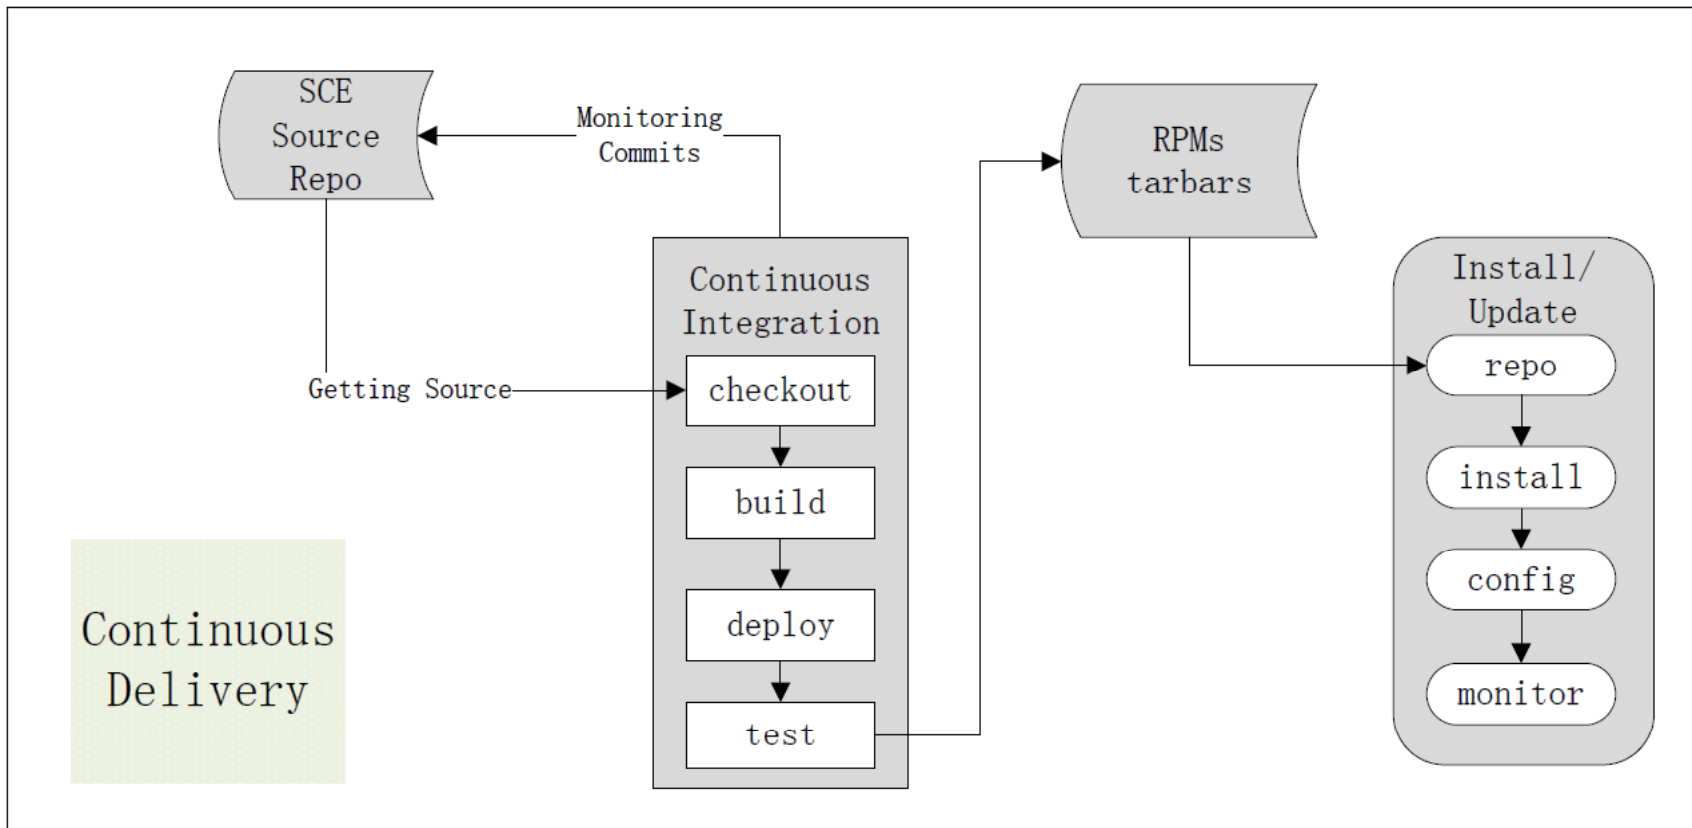


Grid Job Accounting Service



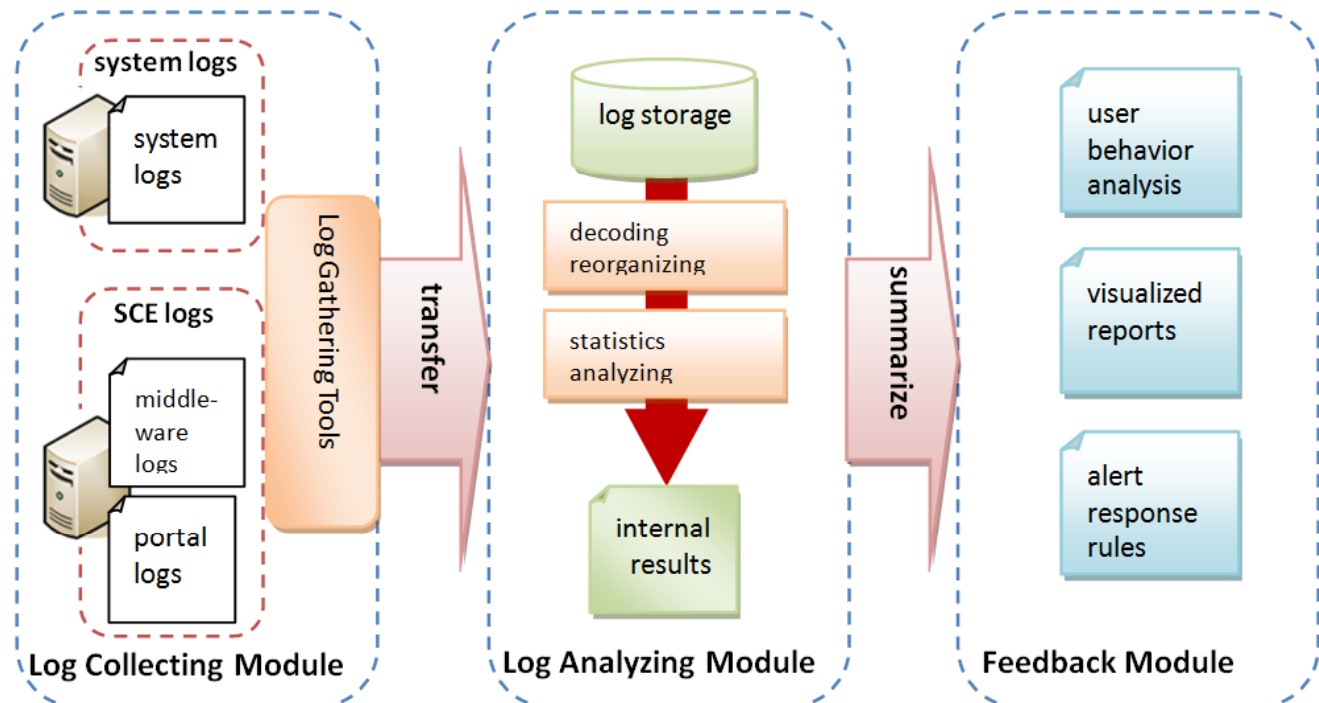
Continuous Delivery

- SCE production line

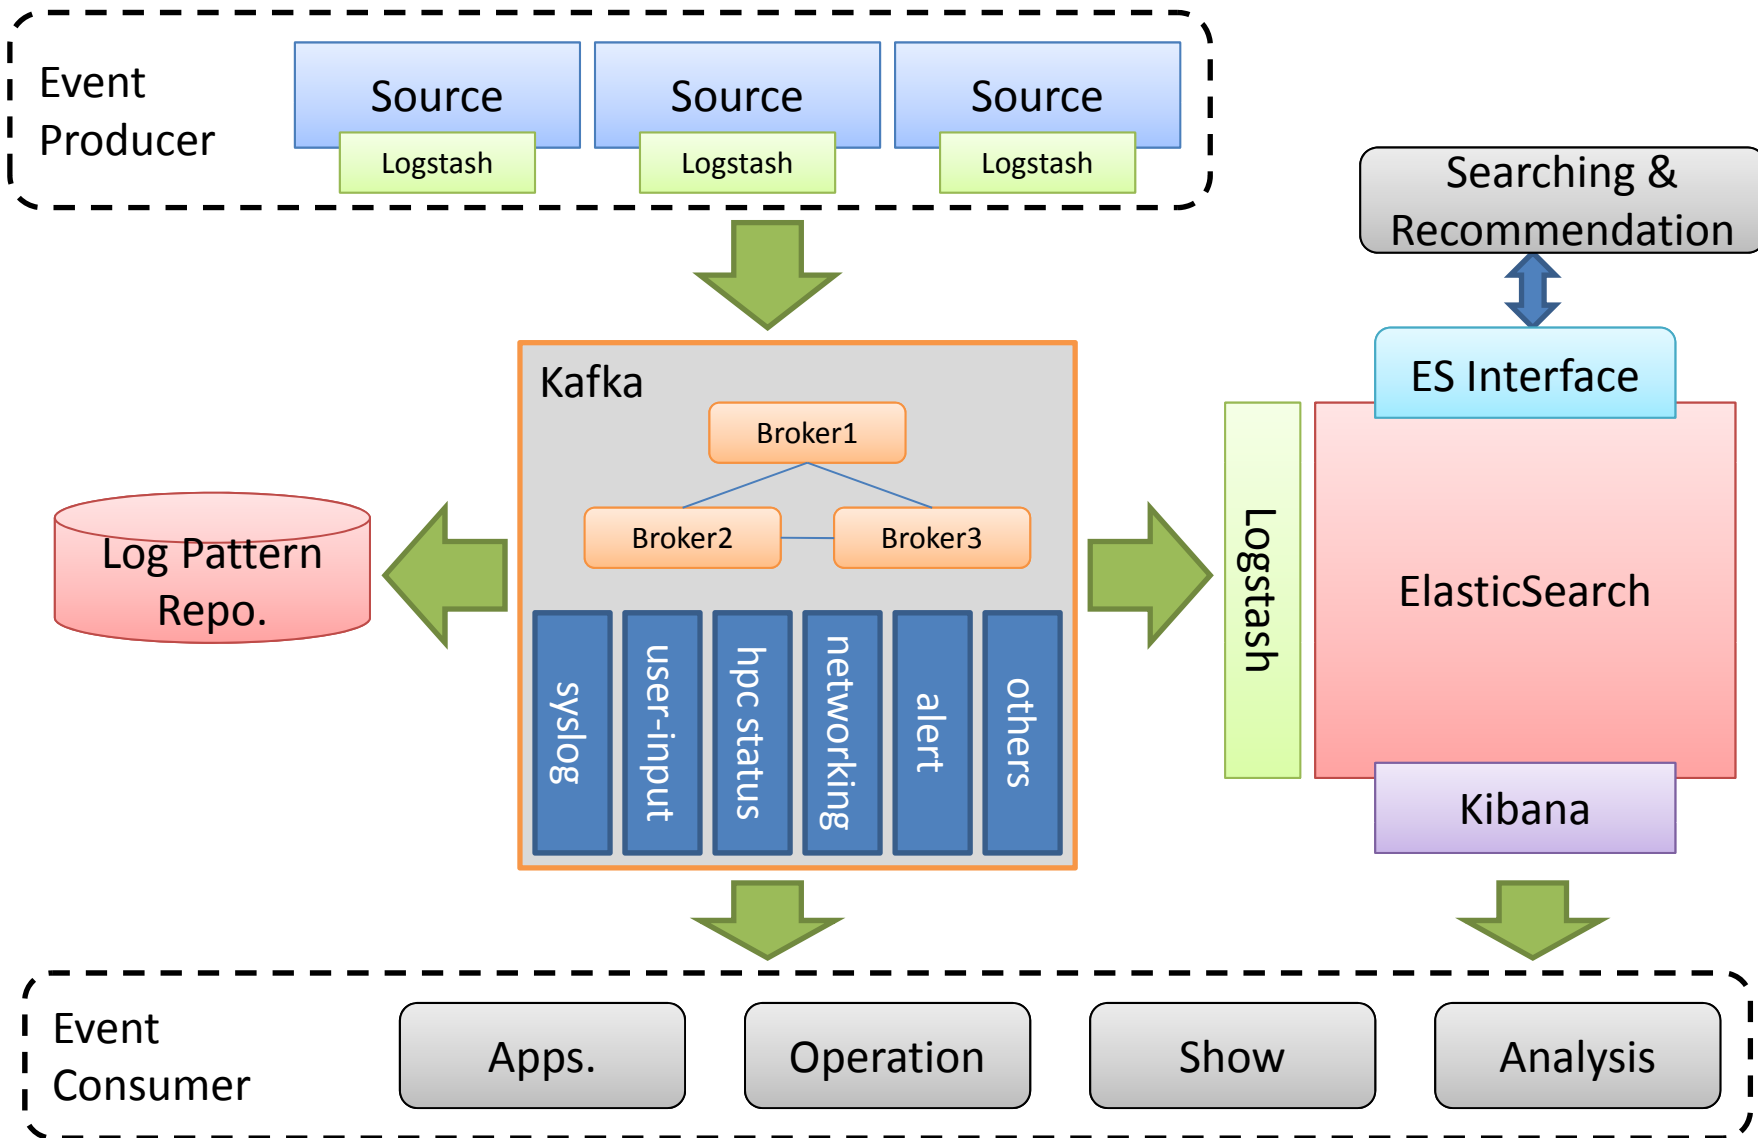


LARGE

- LARGE (Log Analysing fRamework in Grid Environment)
- Grid operation logs
 - Grid servers & services
 - Collecting / Analyzing / Feedback
- Useful
 - Monitoring
 - User patterns
 - Active response



LARGE - event distributing



Summary

- CNGrid: HPC grid + cloud services
- More on <http://www.cngrid.org>

Thank You!

