ACAT 2017

21-25 August 2017
University of Washington, Seattle

Supercomputers, Clouds and Grids powered by BigPanDA for Brain studies

A. Beche[1], K. De[2], F. Delalondre[1], A. Klimentov[3,4], R. Mashinistov[2,4], F. Schuermann[1], I. Tertychnyy[4]

[1]École polytechnique fédérale de Lausanne (EPFL) (Lausanne, Switzerland)
[2]University of Texas at Arlington (UTA) (Arlington, TX US)
[3]Brookhaven National Laboratory (BNL) (Brookhaven, NY US)
[4]National Research Center "Kurchatov Institute" (NRC KI) (Moscow, Russia)

Outline

- Introduction
- PanDA - Production and Distributed Analysis Workload Management System
- Extending PanDA to Supercomputers
- BBP Computing Model
- PanDA portal at BBP
- Conclusions

Introduction

- In 2017 the pilot project between BigPanDA and Blue Brain Project (BBP) (Swiss Federal Institute of Technology in Lausanne) teams began.
- The proof of concept project aimed to demonstrate efficient application of the PanDA Workload Management System, initially developed for HEP applications, for the supercomputer-based reconstructions and simulations offering a radically new approach for understanding the multilevel structure and function of the brain.
- In the first phase, the goal of this joint project is to support the execution of BBP software on a variety of distributed computing systems powered by PanDA.

PanDA - Production and Distributed Analysis System

- PanDA - Production and Distributed Analysis Workload Management System has been developed for ATLAS experiment at LHC.
- PanDA was able to cope with increasing LHC luminosity, ATLAS data taking rate, processing and analysis challenges
- Recently PanDA has been extended to run HEP scientific applications on Supercomputers
- PanDA beyond ATLAS
  - HEP and astro-particle experiments COMPASS and AMS has chosen PanDA as WMS for data processing and analysis. nEDM with LSST will evaluate PanDA. ALICE is interested in PanDA evaluation for OLCF.
  - JINR (Russia) is considering PanDA as main WMS for NICA collider
  - Several PanDA instances beyond ATLAS : OLCF, Taiwan, Amazon EC2, Russia (Moscow, Dubna)

PanDA WMS basics

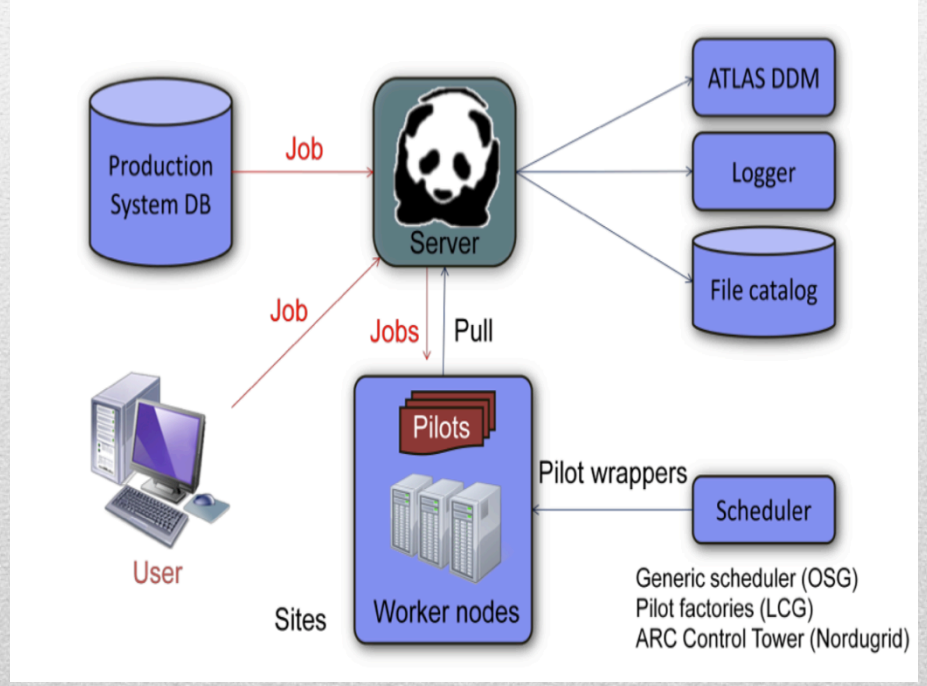

- Basic components are
  - Server – provides mapping of all jobs in the system to all available resources
  - Client – allows to submit jobs, performs actions on jobs and Server, retrieves some information from the Server
  - Pilot – launched on resource retrieves jobs from the Server and handles payload execution
  - Pilot schedulers – responsible to launch the pilots on resources
  - Monitor – provides overall information about the jobs, resources and the system
- External to the PanDA components:
  - Distributed Data Management system
  - File catalog

Extending PanDA to Supercomputers

- BigPanDA is an extension of the PanDA WMS to run ATLAS and non-ATLAS applications on Leadership Class Facilities and supercomputers, as well as traditional grid and cloud resources.
- Use of supercomputers in ATLAS
  - Supercomputing centers in USA, Europe, and Asia, in particular the Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF), the National Energy Research Supercomputing Center (NERSC) in USA, Ostrava supercomputing center in Czech Republic, and "Kurchatov Institute" in Russia (NRC KI) are now integrated within the ATLAS workflow via the PanDA WMS.
- PanDA is a pilot based WMS. Using a pilot and pilot submission system customized for actual SC allows to integrate any opportunistic resources
  - Each supercomputer is unique: Unique architecture and hardware, Specialized Operating System, "weak" worker nodes, limited memory per worker node, different job submission systems, Unique security environment, Own usage policy etc.

Blue Brain Project

- The goal of the Blue Brain Project is to build biologically detailed digital reconstructions and simulations of the rodent, and ultimately the human brain.
- The supercomputer-based reconstructions and simulations built by the project offer a radically new approach for understanding the multilevel structure and function of the brain.
- The novel research strategy allows the project to build computer models of the brain at an unprecedented level of biological detail.
- Supercomputer-based simulation turns understanding the brain into a tractable problem, providing a new tool to study the complex interactions within different levels of brain organization and to investigate the cross-level links leading from genes to cognition.

BBP Target resources

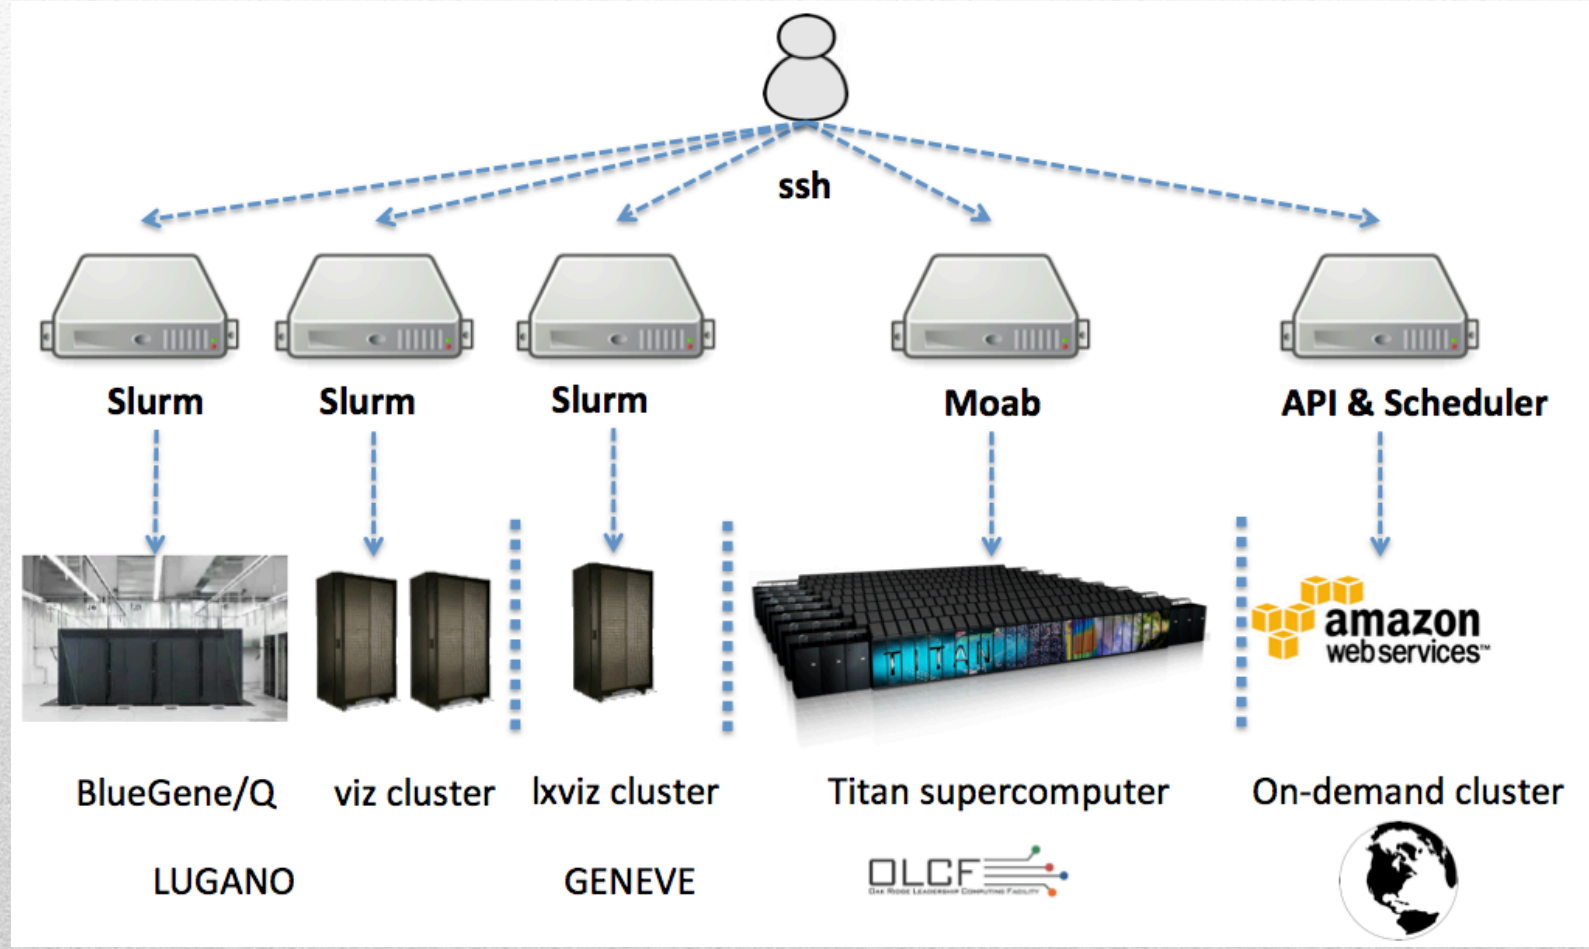

BBP resources metrics

- Intel x86-NVIDIA GPU based BBP clusters located in Geneva (47 TFlops) and Lugano (81 TFlops)
- BBP IBM BlueGene/Q supercomputer (0.78 PFlops and 65 TB of DRAM memory) located in Lugano
- Titan Supercomputer with peak theoretical performance 27 PFlops operated by the Oak Ridge Leadership Computing Facility (OLCF)

Computing model of the Blue Brain Project

Moving from machine specific setup… to common functional building blocks

- Software deployment --> Modules
- Data distribution --> Manual method (rsync)
- Job scheduling --> site specific scheduler

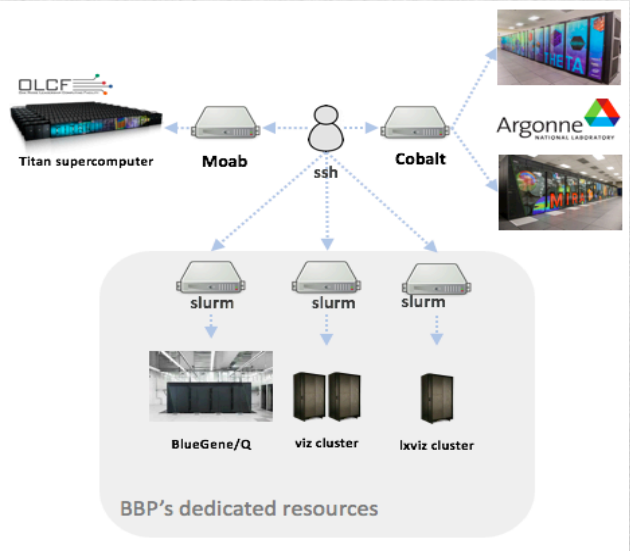

Computing model of the Blue Brain Project

Moving from machine specific setup… to common functional building blocks

- Software deployment --> Modules --> NIX
- Data distribution --> Manual method (rsync) --> Globus
- Job scheduling --> site specific scheduler --> BigPanDA

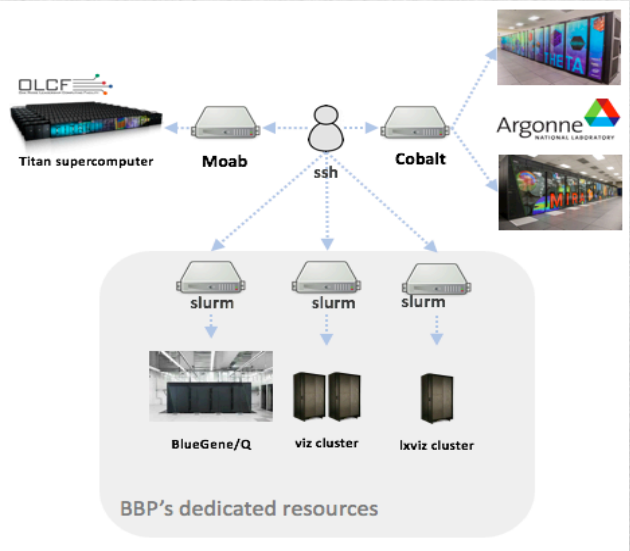

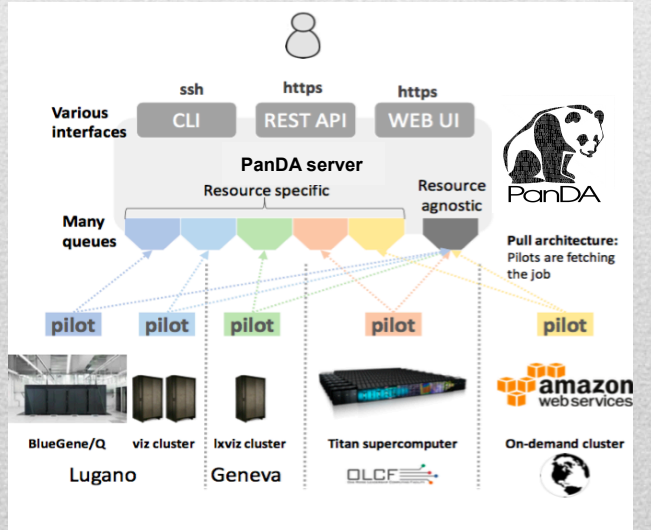

Initial PanDA test @ BBP

- For initial test in March 2017 PanDA Server + Client were installed to the VM at the Campus Biotech (Geneva).
- Job submission has been done via the python based PanDA Client. Pilot were started manually on Geneve and Lugano clusters including BlueGene/Q, Titan and Amazon resources.

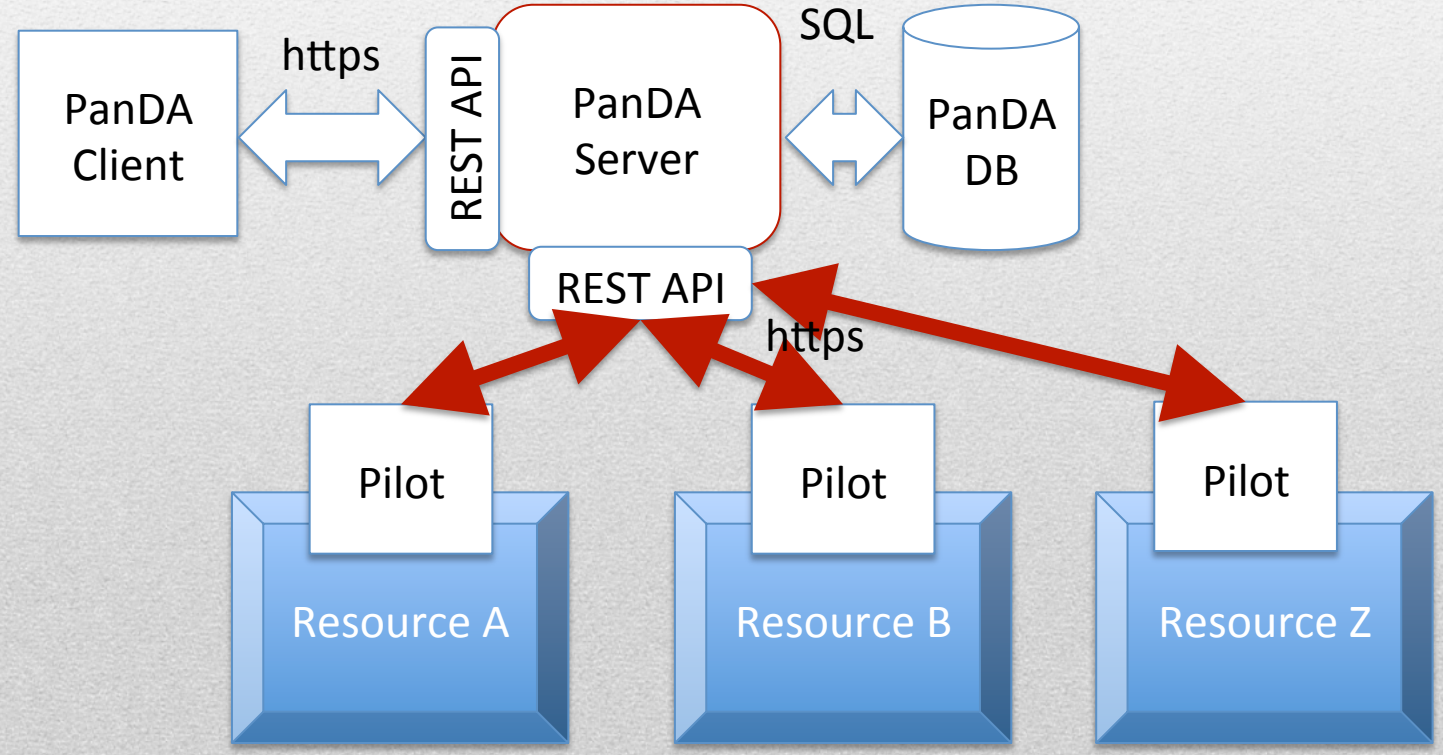

PanDA Portal @ BBP (1/2)

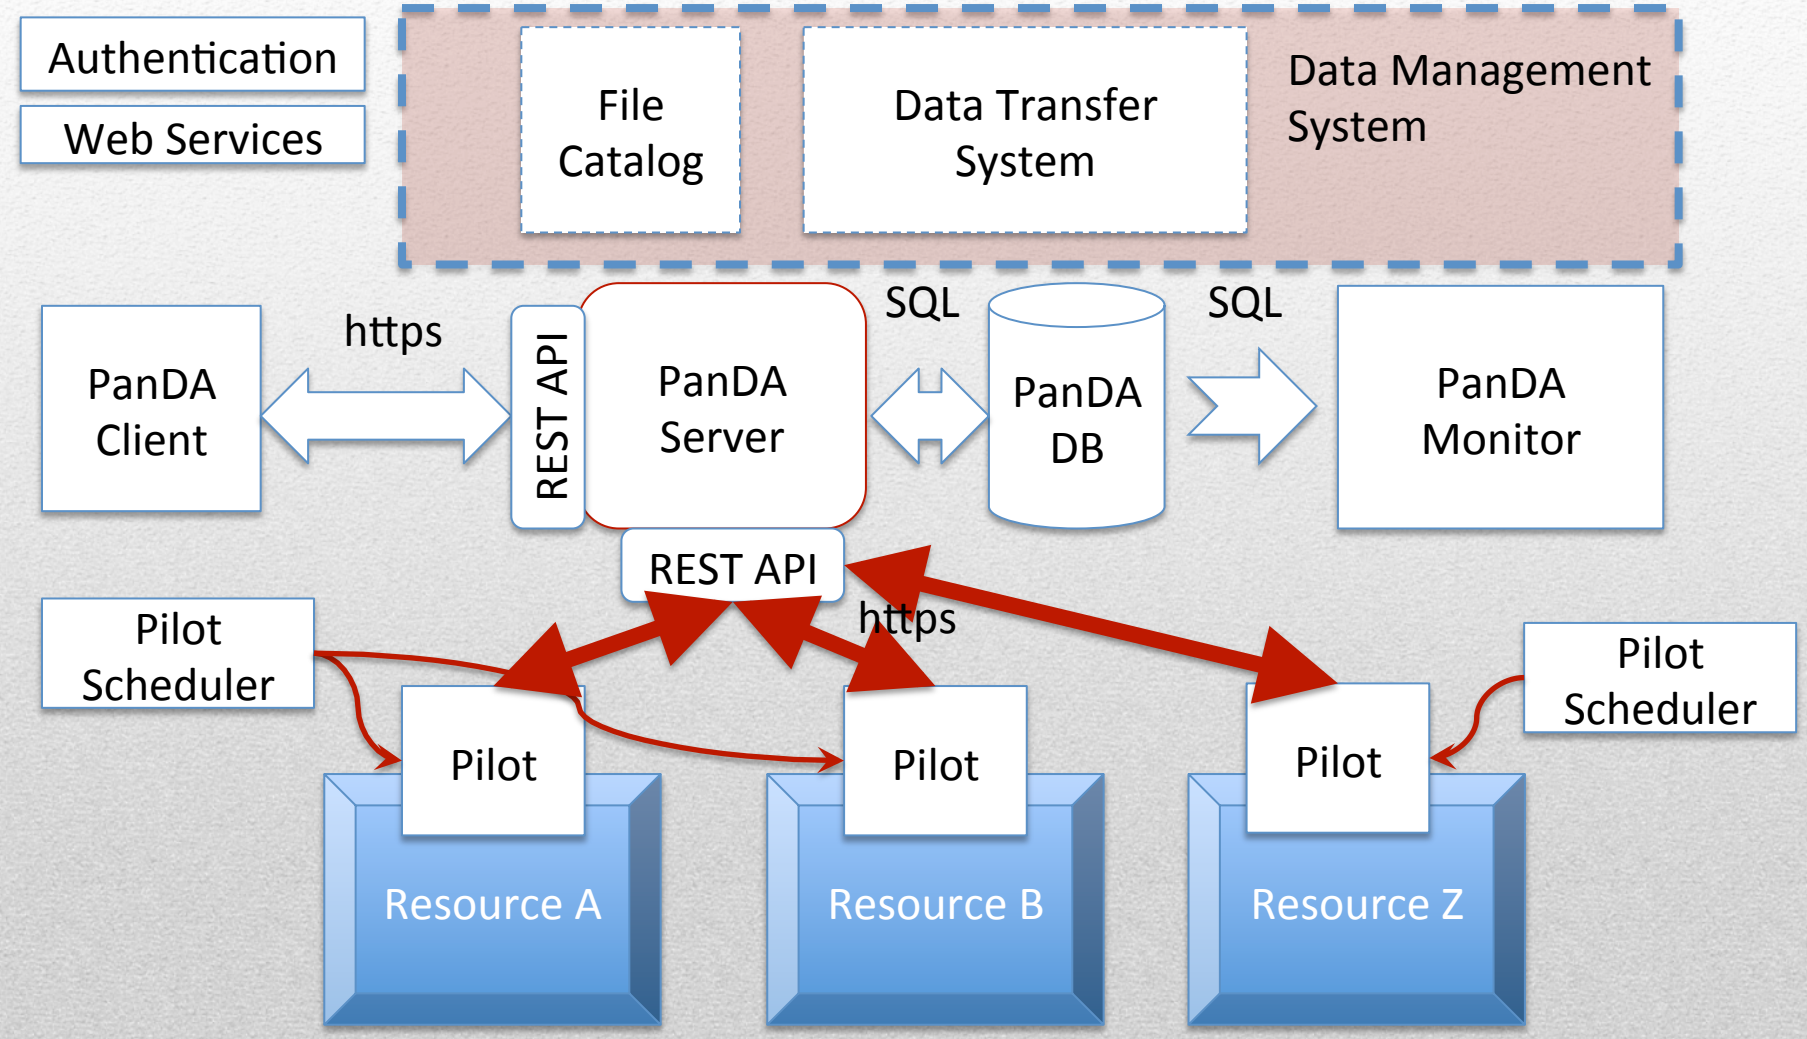

PanDA Portal @ BBP (2/2)

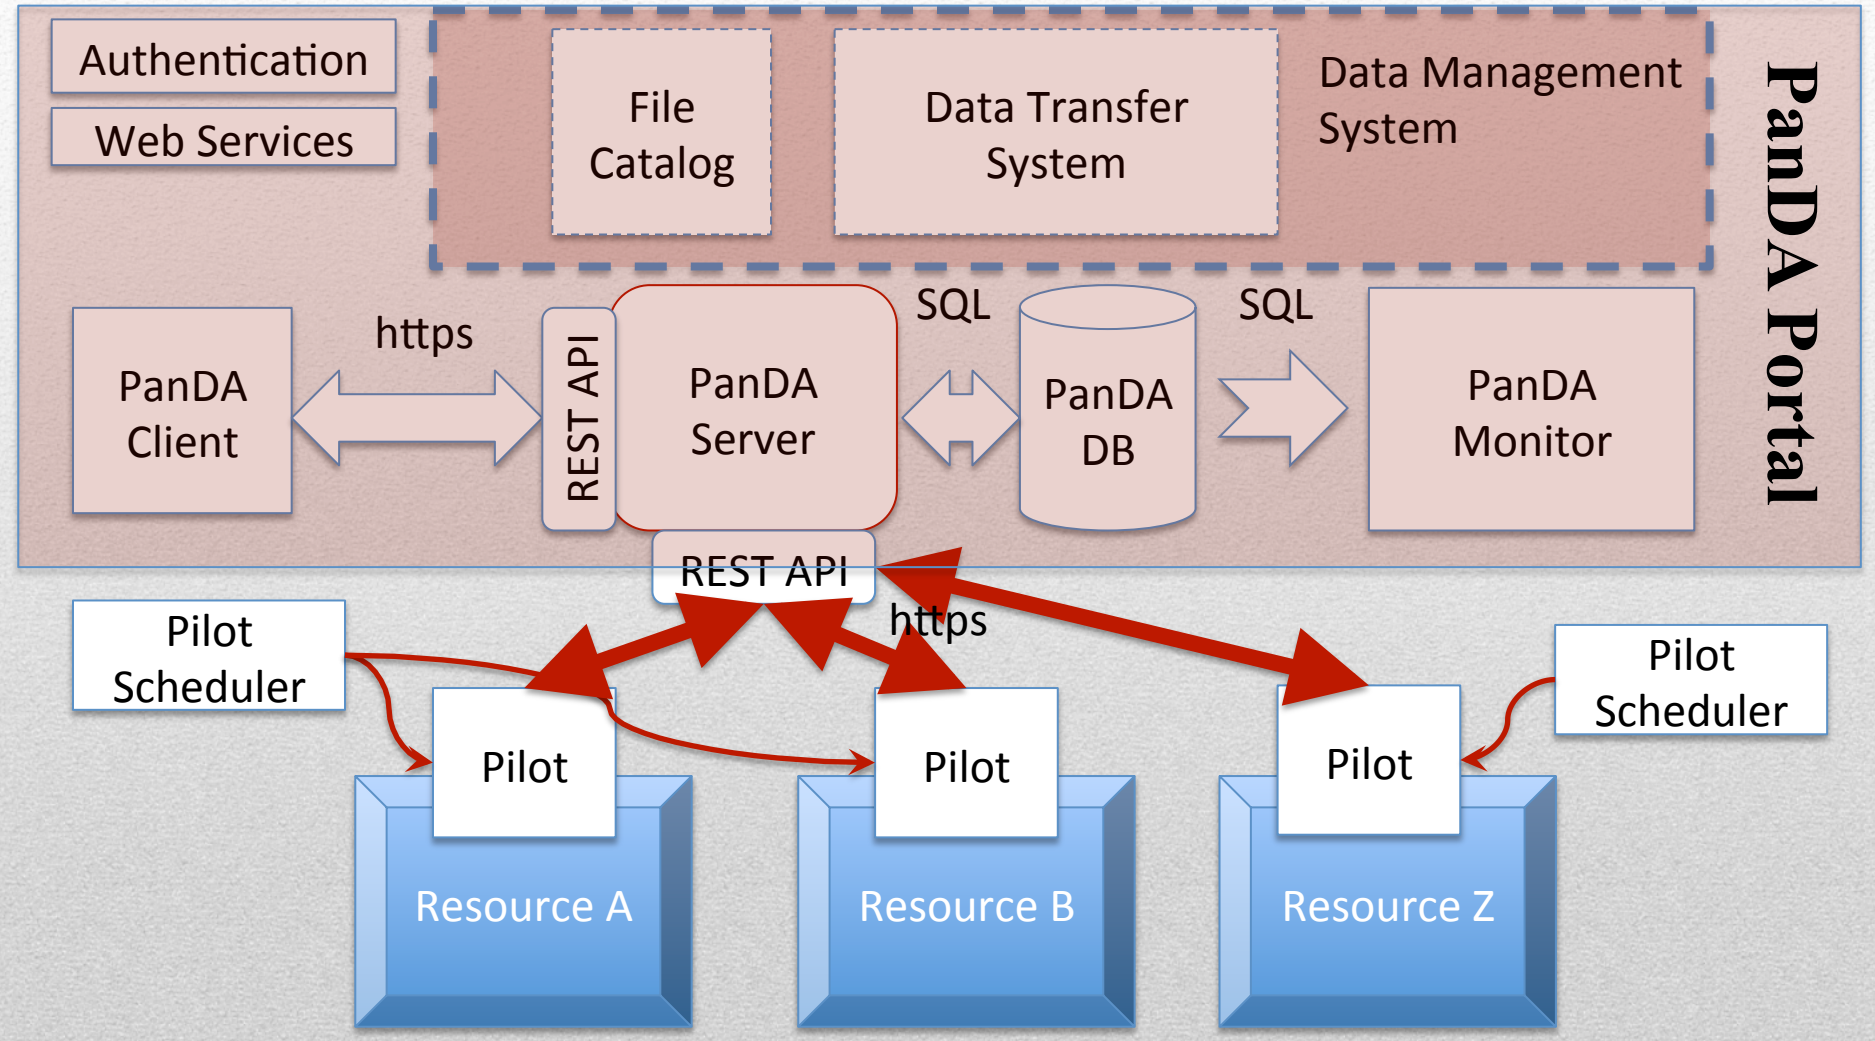

User Interfaces

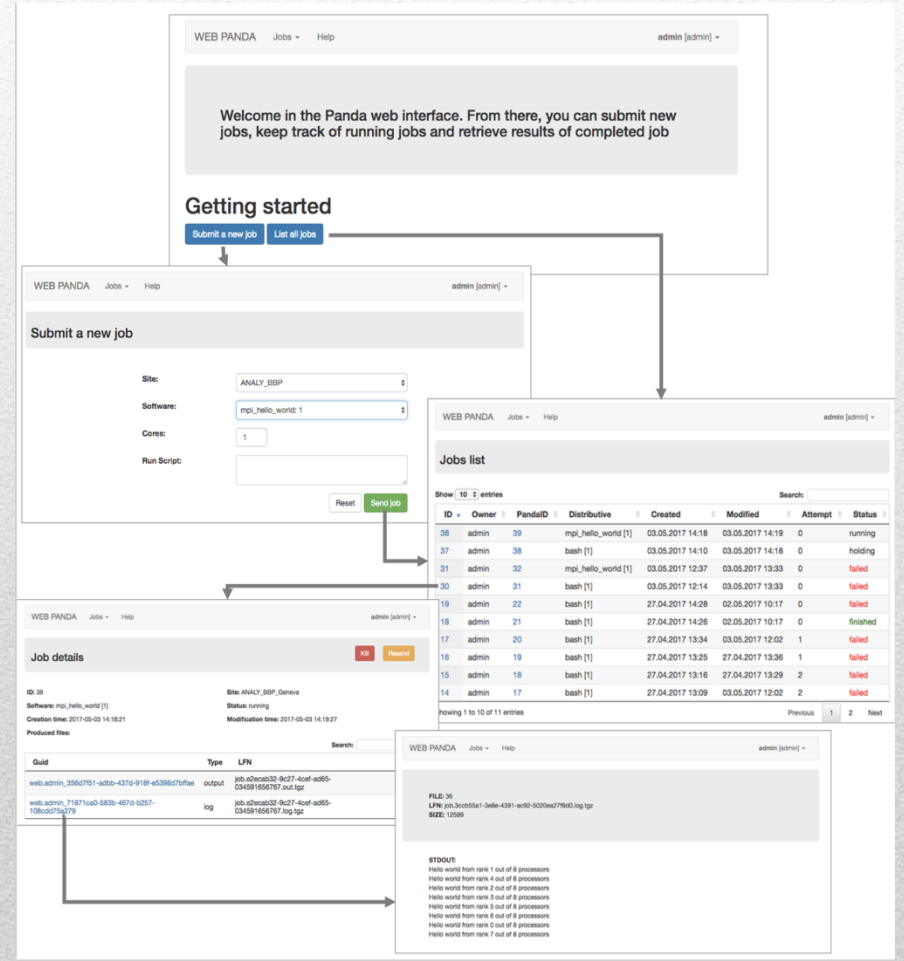

- Reduced intuitive feature set
  - submit job
  - monitor job
  - retrieve output
- Authentication, Authorization
  - integrated with BBP SSO
  - ldap groups aware
- Multiple available distributives
  - Pre-defined scientific apps
  - bash sandbox
- Multiple available sites (PanDA Queue)

Conclusion

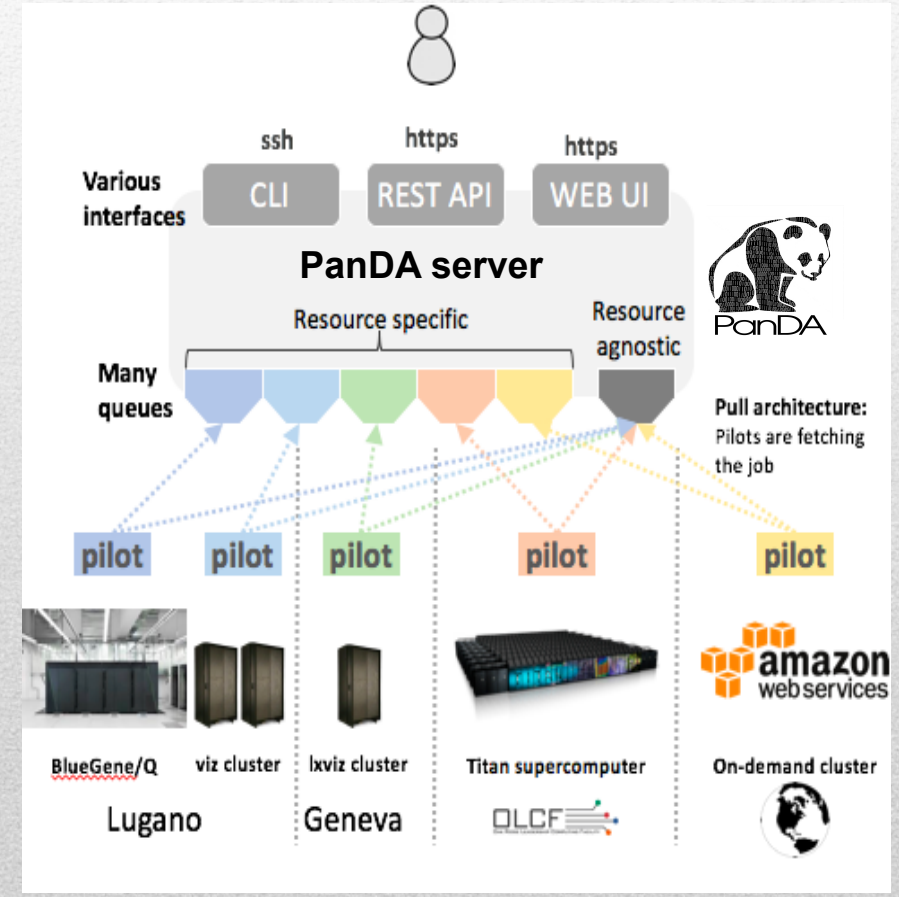

- Phase 1 of "Proof of the concept" was successfully finished
  - The test jobs were successfully submitted to the targeted resources via PanDA portal.
- The project demonstrated that the software tools and methods for processing large volumes of experimental data, which have been developed initially for experiments at the LHC, can be successfully applied to BBP.
- Phase 2 of "Pre-production" is under investigation currently

Thank you