
Brookhaven
National Laboratory
U.S. DEPARTMENT OF
ENERGY
PanDA @ BNL
Xin Zhao (on behalf of the PanDA team)
PanDA Community Forum, January 9th, 2025
@BrookhavenLab

Overview

- A general purpose PanDA service at BNL
  - Drawing on PanDA Rubin work to support PanDA outside CERN
    - non-Oracle PanDA (PostgreSQL)
    - OpenSearch analytics independent of ATLAS/CERN infrastructure
- Serving EIC projects like AID2E (AI-assisted detector design for the EIC)
- Supporting EIC/ePIC's coming evaluation of PanDA for their workflow management system (probably in first half of 2025)
- Also can be used by other BNL local scientific communities, and other PanDA evaluators as they arise

The Setup

- A standard PanDA instance with all essential components
  - PanDA server/Jedi host
    - one host – 8 Core, 16GB RAM, 50GB local disk
  - harvester (2 instances)
    - one dedicated host – 4 core, 8GB RAM, 50GB local disk
    - another OSG submit host – 4 core, 16GB RAM, 100GB local disk
  - iDDS
    - one host – 8 core, 16GB RAM, 50GB local disk
  - PanDA monitor (bigmon)
    - one host – 4 core, 16GB RAM, 50GB local disk
  - postgresql Database for the PanDA and iDDS DB
    - one host – 8 core, 24GB RAM, 50GB local disk
  - messaging service (activemq)
    - one host – 4 core, 8GB RAM, 50GB local disk
- Analytics
  - OpenSearch cluster
    - One controller/data node – 8 core, 16GB RAM, 150GB local disk
    - One data node – 8 core, 16GB RAM, 150GB local disk
    - OpenSearch dashboard – 8 core, 16GB RAM, 50GB local disk
  - logstash
    - one host – 8 core, 16GB RAM, 50GB local disk
  - filebeat agent runs on the PanDA hosts
- External services
  - PanDA IAM as token issuer
- Manually installed on RHEL8 nodes (VMs)

Computing Resources

- The shared pool at BNL SDCC
  - A harvester instance submits PanDA pilots/jobs to the shared pool via BNL OSG CEs (SDCC), using token based authentication
    - A vault server (SDCC) fetches/renews CE access tokens from the BNL CILogon on behalf of the harvester operator
- OSG OSPool (Open Science Pool)
  - A distributed computing resource operated by the OSG, with contributions from many institutes. We are using the ePIC VO shares of it.
  - BNL SDCC runs a dedicated ePIC OSG submit host, where users log on and submit HTCondor jobs as if it were a local HTCondor cluster — behind the scene OSG GlideinWMS provisions resources. Users don't need any other credentials like tokens or x509 certificates
  - In order to allow PanDA jobs to run on the OSPool, a harvester instance is installed on the same OSG submit host, submitting vanilla universe HTCondor jobs.
- PanDA and OSG/GlideinWMS
  - From the PanDA point of view, OSPool/GlideinWMS acts as regular PanDA queue(s)
  - From the OSPool/GlideinWMS point of view, PanDA acts as a regular user

Using PanDA@BNL

- PanDA client tools allow users to submit jobs to the PanDA system.
- User authentication and authorization using tokens
  - PanDA server functions as an OIDC client, authenticating users through OIDC providers such as PanDA IAM and CILogon.
  - Users register with PanDA IAM, which, upon successful authentication via their chosen identity provider (usually their home institute) through CILogon, issues ID tokens. These tokens are then used to authenticate users to the PanDA server.
  - PanDA client tools manage the token retrieval process automatically on behalf of the user (except for the manual logon step to their IdP).

Analytics

- Similarly to what is used in the ATLAS land, filebeat and custom scripts gather logs and statistics from various PanDA hosts and database, forwarding them to the OpenSearch cluster through logstash. OpenSearch dashboard is used for visualization.
- OpenSearch dashboard supports two sign-in options : username/password, and OIDC (using PanDA IAM as the token issuer)

Brookhaven National Laboratory

Next Steps

- Adding a new Rucio service at BNL, integrated with PanDA@BNL for job input/output
  - Cooperative effort with SDCC (Hiro Ito) and other NPPS experts (Cedric Serfon, a Rucio core team member)
  - Timeline – first quarter of 2025
- BNL PanDA+Rucio will support ePIC's evaluation of the two of them working together orchestrating workload/data management, as they were designed to do (ePIC has adopted Rucio)
- And they will be used together for the distributed workflows of AID2E