


AWS-POC Demo

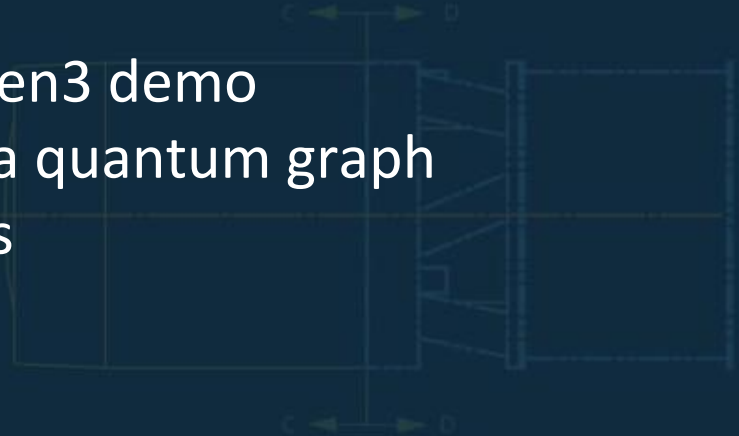


Amazon Web Services Proof of Concept

- Started late April with AWS credits
- DMTN-114
- Team:
 - Dino Bektesevic (UW)
 - Greg Thain, Todd Miller, Miron Livny (HTCondor team; UWisc)
 - Aaron Bucher, Lorena Costanza, Chris Mills, Sanjay Padhi (AWS)
 - Hsin-Fang Chiang, Tim Jenness, K-T Lim, Wil O'Mullane
 - Michelle Butler, Greg Daues, Michelle Gower, Steve Pietrowicz

Amazon Web Services Proof of Concept

- Cloud deployment of data release production
 - Ability to burst to the cloud
 - Scalability: Condor Annex + Spot > 60,000 instances
 - Shareability: across regions or countries
 - Pricing
-
- Use Gen3 middleware
 - "ci_hsc" workflow from January's Gen3 demo
 - Use Gen3 pipetask to generate a quantum graph
 - Translate workflow into Pegasus



Demo / Tools

- Pegasus + HTCondor
 - Submit and monitor the workflow via Pegasus
 - Execute in the HTCondor pool, with machines added via HTCondor Annex
- HTCondor Annex get EC2 (Elastic Compute Cloud) instances via AWS Spot Fleet
- Datastore: AWS S3 (Simple Storage Service)
- Registry: one shared gen3.sqlite3 on EFS (Elastic File System)
 - Work in progress → RDS (Relational Database Service)

Demo / Steps

1. Start an on-demand instance from an AMI. This is the HTCondor central manager.
2. Setup the w_2019_28 stack + the daf_butler S3 branch
3. Submit the Pegasus ci_hsc workflow
4. Use condor_annex to get Spot instances for the condor pool
5. Check Spot request fulfillment and the condor pool
6. Show Butler config
7. Check output files in the S3 bucket