ProdSys2

Kaushik De, Alexei Klimentov

ATLAS TIM, Chicago

October 28, 2014

Introduction

- Many talks and discussions about ProdSy2 already at TIM
  - Almost every talk so far touched on ProdSys2 (and Rucio)
  - And many talks at S&C week, ATLAS week…
  - These discussions will continue with the rest of TIM agenda
  - I will not repeat or summarize what we heard already
- This talk will focus on the future
  - Short term – goals and priorities
  - Medium term – new features and capabilities
  - Long term – what is left to do?
  - Following tradition of past TIM's, I will mostly focus on the long term

ProdSys2 Status

- Startup was a little late but in good shape for Run 2
- ProdSys1 will be switched off by the end of 2014
- Very active development and commissioning team
- User communities are growing rapidly
- JEDI
  - Fully deployed for all workflows, mostly tuning now
  - New features being added continuously
- DEFT
  - Being tested for most Run 1 workloads
  - Many new Run 2 use cases being commissioned
- Monitoring
  - BigPandaMon is finally here, thanks to heroic effort by Torre, Jarka
  - New developers are being added – they should be active soon

Short Term Goals

- While JIRA tickets continue to be extremely valuable for tracking issues, they also consume all developers
- Highest priority is to support commissioning
- Most Run 1 workloads already supported
  - We will wait for requests for rarely used workflows
- Many new Run 2 workloads (which would have been impossible with ProdSys1) are being commissioned
  - Train production is perhaps most active/important
- Most JEDI tuning will come during commissioning
- DEFT UI has long list of requests
- Need some free time to work on missing DEFT features

Priorities

- Priorities are evaluated every week in ProdSys2 meetings attended by development team, commissioning team, users
- But we risk developers getting swamped with commissioning
- Need to set milestones, and derive priorities from them
- It would be useful to get timescale for each remaining task from Jose, Nurcan, David and others
- Merge these with medium term development priorities

Medium Term

- New features and capabilities
  - Continue tuning workflows – next few sessions
  - More monitoring – list keeps growing
  - Easy Task cloning
  - Stop use of lists and external scripts
  - Meta-Tasks
  - Basket of Tasks

Long Term Planning

- Not much discussed so far
- Focus has been on developing and commissioning new systems
- Steady state maintenance will continue
  - Workflows change, requirements are driven by operations
  - Past experience suggests this will keep developers busy
- We also need new ideas
  - It takes a long time to nurture new concepts into reality (Rucio started 3 TIM's ago, ProdSys2 2 TIM's past)

D2LM

- What is D2LM? Another bad name, continuing on PD2P?
- Dynamic Data Lifetime Management
  - We will get policies from management – I do not expect such fixed policies to work much beyond 1-2 months of Run 2 (remember how quickly we needed PD2P after Run 1 started)
  - Rucio will enforce lifetime rules
  - Large space between management and Rucio to implement automated and dynamic lifetime management
  - Workload management is the best place to implement D2LM
- I promise we will change name – after working hard till 2am last night, and 3am the night before, no resolution yet, but will find good name for this project soon

DNP

- Sky is the limit – when it comes to bad acronym
- Dynamic Network Provisioning
  - PanDA successfully using network metrics already for workload management (see recent talks by Artem)
  - But perfSONAR is so Old Style – though reliable operations is still missing
  - SDN is the new rage
  - We want to put network provisioning into ProdSys2/PanDA workflows – supported by ANSE and ASCR BigPanDA projects
  - Looking into DaTRi/PD2P as use cases

DCEA

- This acronym outshines all previous ones!
- Dynamic Compute Engine Allocations
  - We will have discussions this morning
  - We heard from Ale yesterday how too much fragmentation of resources leads to unstable load brokering
  - Pre-assigning resources is also bad
  - CE resources continue to become more AGILE, providing new opportunities and new challenges for workload management
  - How to automatically balance resource size (many dimensions – not just number of CPU's) to match constantly changing workloads and their priorities