

# Monitoring HTCondor with the BigPanDA monitoring package

J. Schovancová<sup>1</sup>, P. Love<sup>2</sup>, T. Miller<sup>3</sup>,  
T. Tannenbaum<sup>3</sup>, T. Wenaus<sup>1</sup>

<sup>1</sup> Brookhaven National Laboratory

<sup>2</sup> Lancaster University

<sup>3</sup> UW-Madison, Department of Computer Science

HTCondor Week 2014

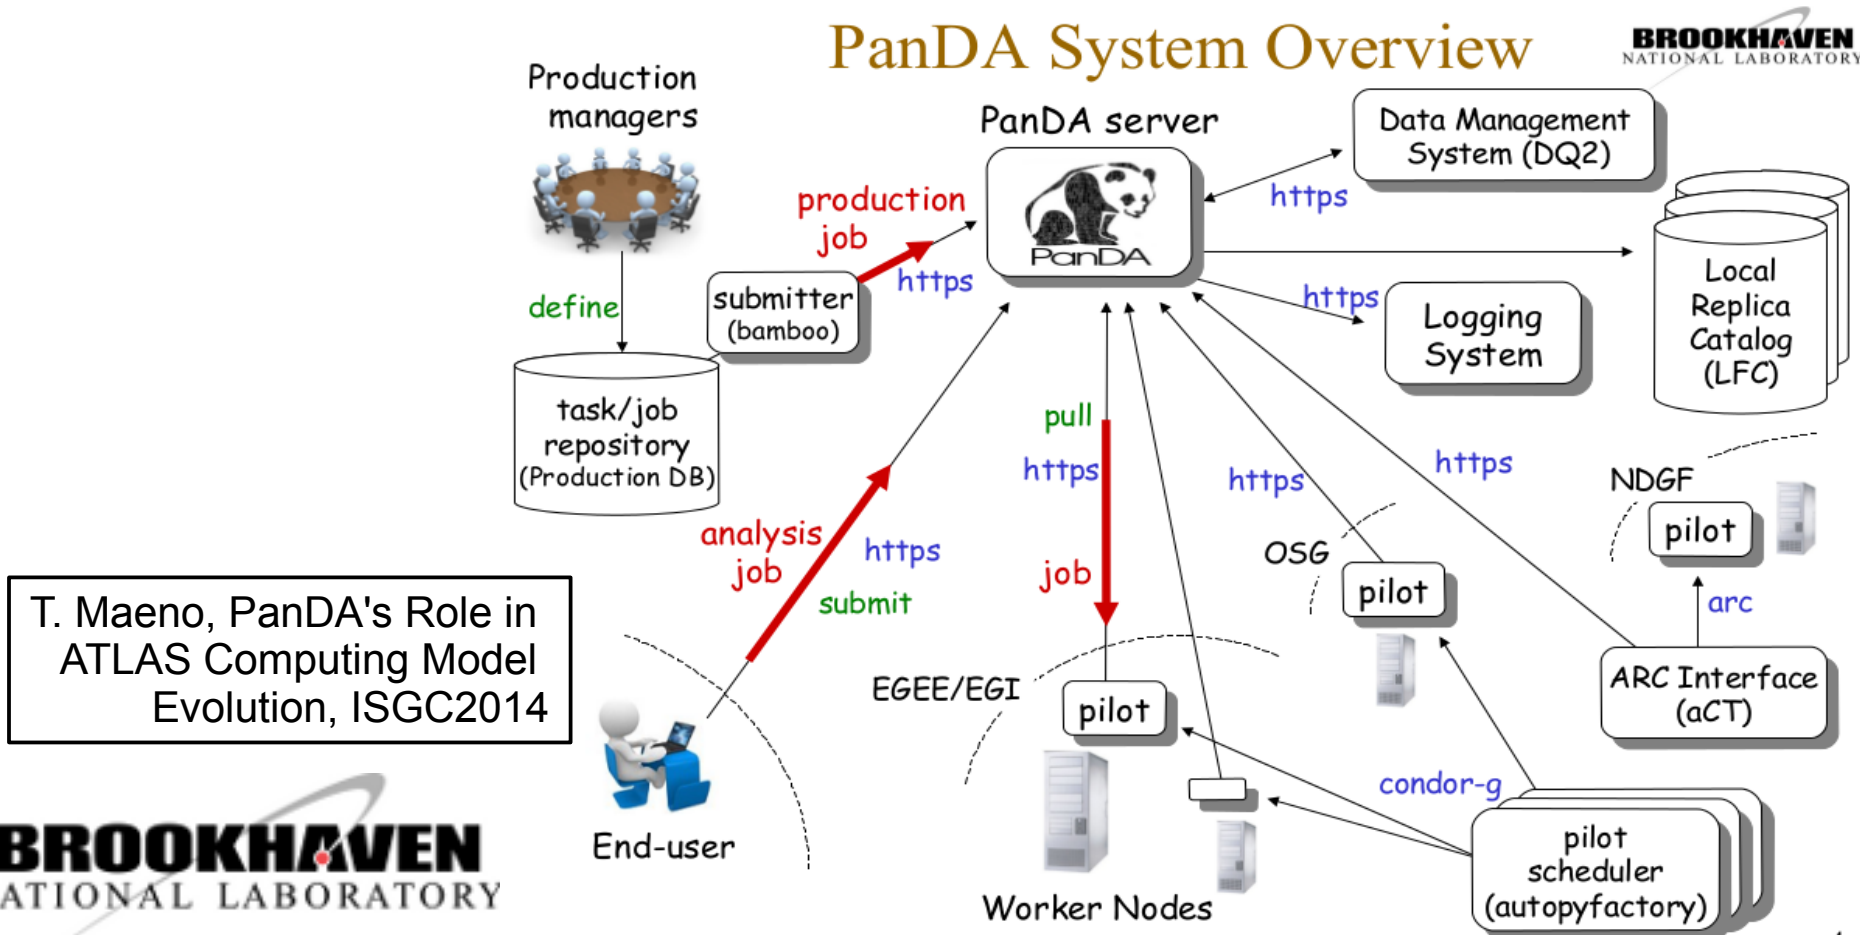
28 – 30 April 2014, University of Wisconsin, Madison, WI, USA



# Introduction

- PanDA = **P**roduction **an**d **D**istributed **A**nalysis
- Start: Aug 2005, in production for US ATLAS: Dec 2005, since 2008 WMS for the whole ATLAS Collaboration

## PanDA System Overview



# Evolution: The Next Generation

- US DOE ASCR and HEP funded project: *Next Generation Workload Management and Analysis System for Big Data*, code name: **BigPanDA**
  - 3 year project since September 2012 to evolve PanDA Workload Management System beyond ATLAS and LHC
    - Factorize core
    - Leverage intelligent networking
    - Extend scope
    - **Monitoring** and usability

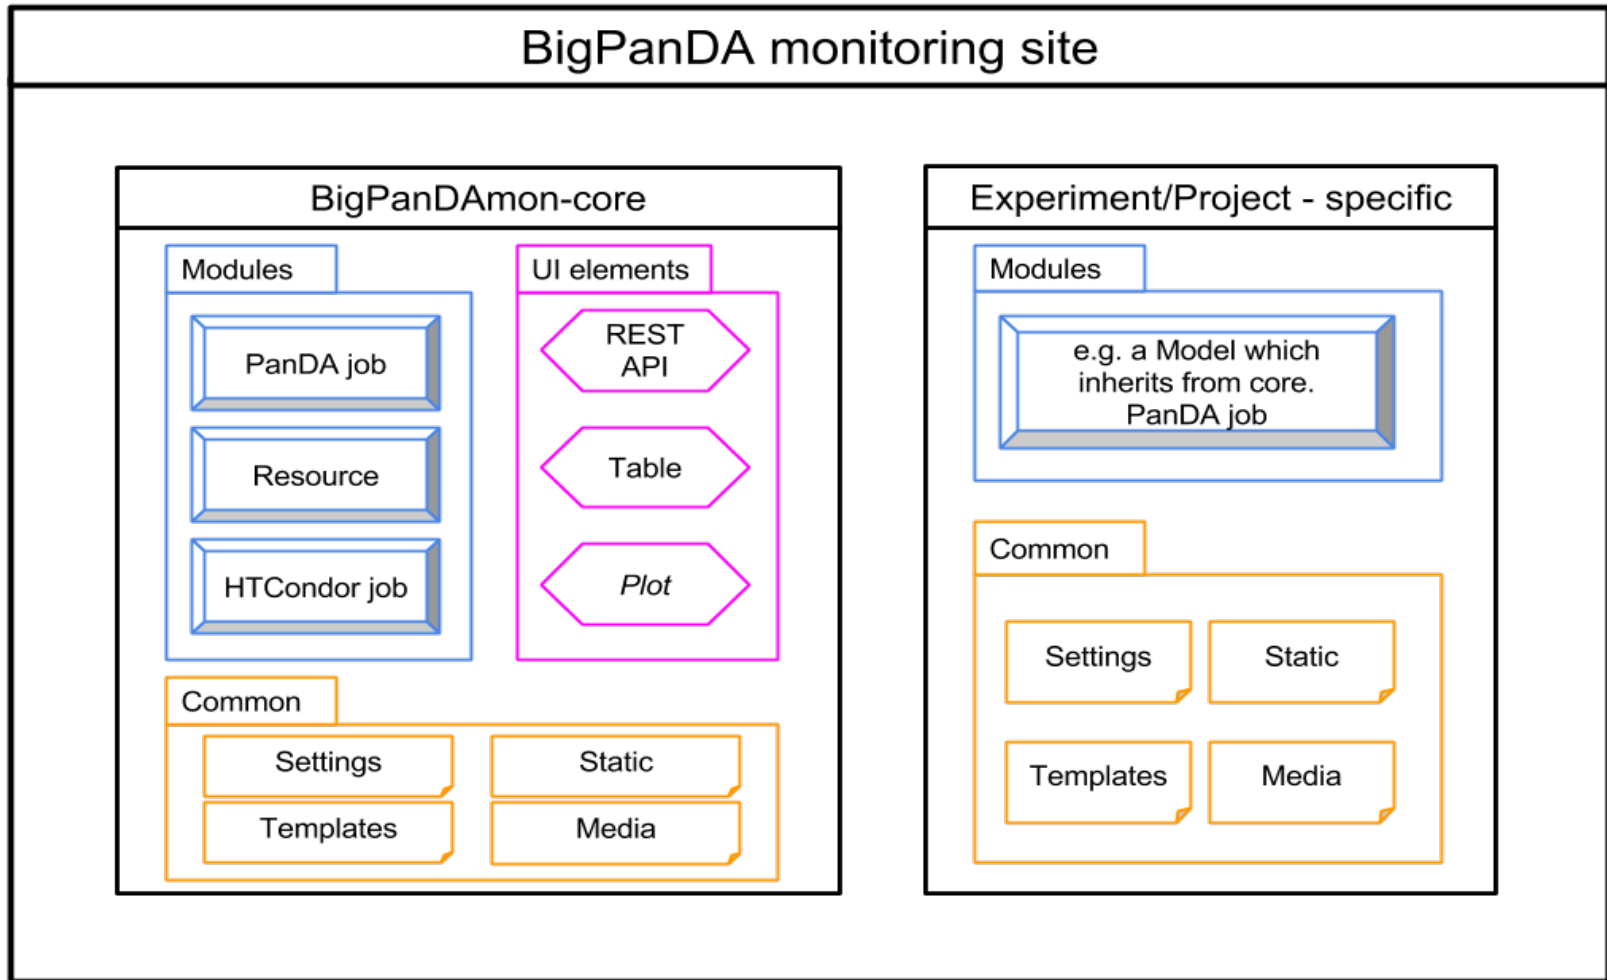
# BigPanDA Monitoring

- BigPanDAmon package based on **django** framework
- Modular, easy to bring up a new project/VO
- Clear separation between data access and visualization
  - Provide REST APIs to access object information
- Runs on top of Oracle or MySQL DB backends
- Documentation for developers: describes configuration and modules
- Deployed with RPMs

# HTCondor Monitoring

- HTCondor is an essential foundation technology for PanDA
  - the engine for PanDA's resource provisioning system (pilot factory)
- Integrating HTCondor monitoring into PanDA's successful monitoring system can provide value for both
- HTCondor monitoring has been implemented as a first application of BigPanDAmon's modular architecture

# BigPanDA Monitoring Modules



# Monitoring Traffic

- Condor sched lib plugin to send updates to PanDA mon
  - Asynchronous w.r.t. scheduler
  - More technical details in Todd Tannenbaum's talk
- **Connection to PanDA API** with frequency **1 Hz**
  - Bulk Update, Add, Flag deleted
  - **Avg. 2-5 records/s, Peak: 100s records/s**
  - Currently HTTPS (X509) authentication
- Updated information available in the monitoring

# PanDA – HTCondor REST API

- Bulk-operation API resource

---

Resource	<code>/v2/api-auth/htcondor/jobs/</code>	
HTTP verb	Purpose	Description
<b>POST</b>	create	Bulk create new HTCondor job.
<b>GET</b>	read	Bulk list HTCondor jobs.
<b>PUT</b>	update	Bulk update HTCondor jobs.
<b>DELETE</b>	delete	Bulk flag of records which were deleted from the HTCondor.

- Single job API resource is not needed for now.
  - Perhaps we will need a single-job GET later on.

Nice and useful reading: <http://apigee.com/about/api-best-practices/api-design/ebook>



# HTCondor Mon – list jobs

Show 10 entries

Details	Owner	WMS ID	Global Job ID	Submitted	Run time	Status	Guessed status	Priority
	apf	1792	aipanda002#177326.0#1385140808	2013-11-22T12:20:08Z		R		0
	apf	1792	aipanda002#160997.29#1383646907	2013-11-05T05:21:47Z		R		0
	apf	1792	aipanda002#177377.17#1385143058	2013-11-22T12:57:38Z		I		0
	apf	1793	aipanda002#177315.9#1385140494	2013-11-22T12:14:54Z		I		0
	apf	1793	aipanda002#177321.2#1385140682	2013-11-22T12:18:02Z		I		0
	apf	1793	aipanda002#177354.2#1385142023	2013-11-22T12:40:23Z		I		0
	apf	1793	aipanda002#177402.0#1385143926	2013-11-22T13:12:06Z		R		0
	apf	1793	aipanda002#160999.25#1383647010	2013-11-05T05:23:30Z		R		0
	apf	1793	aipanda002#177315.14#1385140494	2013-11-22T12:14:54Z		I		0
	apf	1793	aipanda002#177348.27#1385141833	2013-11-22T12:37:13Z		I		0

Showing 1 to 10 of 2,739 entries

First Previous 1 2 3 4 5 Next Last

# HTCondor Mon – job filter

**Wildcard \*** available for fields Owner, Global Job ID, Guessed status. **Interval mark :** available for fields WMS ID, Run time, Priority. Use e.g. **1776** for exact match, or **:1750** for everything  $\leq 1750$  , or **1780:** for everything  $\geq 1780$  , or **1720:1730** for interval including boundaries.

<b>Owner:</b>	<input type="text"/>	<b>WMS ID:</b>	<input type="text"/>	<b>Global Job ID:</b>	<input type="text" value="*aipanda002*"/>
<b>Submitted:</b>	<input type="text"/> from: <input type="text"/>	<b>Status:</b>	<input type="text" value="H = on hold&lt;br/&gt;R = running&lt;br/&gt;I = idle (waiting for a machine to execute on)&lt;br/&gt;C = completed"/>		
	to: <input type="text"/>				
<b>Run time:</b>	[s] <input type="text"/>	<b>Guessed status:</b>	<input type="text"/>	<b>Priority:</b>	<input type="text"/>

## Filter table!

- Filter by all visible columns
- Bookmark filtered view, send to a collaborator

# HTCondor Mon – job details

+	apf	1792	aipanda002#177326.0#1385140808	2013-11-22T12:20:08Z	R	0
-	apf	1792	aipanda002#160997.29#1383646907	2013-11-05T05:21:47Z	R	0

<b>Owner</b>	apf	<b>WMS ID</b>	1792	<b>Global Job ID</b>	aipanda002#160997.29#1383646907
<b>Submitted</b>	2013-11-05T05:21:47Z	<b>Priority</b>	0	<b>Condor ID</b>	160997.29
<b>Status</b>	R	<b>Run time</b>	null	<b>size</b>	1.0
<b>Guessed status</b>		<b>cpu_time</b>	0	<b>manager</b>	
<b>p_start_time</b>	null	<b>p_end_time</b>	null	<b>p_modif_time</b>	null
<b>host</b>	null	<b>p_factory</b>	aipanda002	<b>p_schedd</b>	aipanda002
<b>Command</b>	/etc/apf/runpilot3-wrapper-oct02.sh			<b>executable</b>	/etc/apf/runpilot3-wrapper-oct02.sh
<b>p_stdout</b>	<a href="http://aipanda002.cern.ch:25880/2013-11-05/CERN-P1_OpenStack/160997.29.out">http://aipanda002.cern.ch:25880/2013-11-05/CERN-P1_OpenStack/160997.29.out</a>			<b>p_stderr</b>	<a href="http://aipanda002.cern.ch:25880/2013-11-05/CERN-P1_OpenStack/160997.29.err">http://aipanda002.cern.ch:25880/2013-11-05/CERN-P1_OpenStack/160997.29.err</a>
<b>goodput</b>	-1	<b>cpu_util</b>	-1	<b>mbps</b>	null
<b>read_</b>	null	<b>write_</b>	null	<b>seek</b>	null
<b>xput</b>	null	<b>bufsize</b>	524288	<b>blocksize</b>	32768
<b>p_description</b>					

+	apf	1792	aipanda002#177377.17#1385143058	2013-11-22T12:57:38Z	I	0
---	-----	------	---------------------------------	----------------------	---	---

# Monitoring Wish list

- HTTPS → HTTP with IP/host restriction of agents
- Schema
  - PanDA mon side: use classads attributes (schema-less)
  - Condor side: minimize data translation
- RPCs
  - Provide list of compulsory/available attributes in API
- Scalability
  - Leverage experience with Redis
  - Webserver performance tuning (Apache2/WSGI, nginx/gunicorn)
- Summaries
  - Provide parameterized summaries/plots beside job lists

# Summary

- Fruitful collaboration between PanDA and HTCondor teams!
- First version of the HTCondor monitor developed
  - ✓ API v1 used on a production machine at UW cluster
- More updates yet to come.

# Monitoring HTCondor with the BigPanDA monitoring package

J. Schovancová<sup>1</sup>, P. Love<sup>2</sup>, T. Miller<sup>3</sup>,  
T. Tannenbaum<sup>3</sup>, T. Wenaus<sup>1</sup>

<sup>1</sup> Brookhaven National Laboratory

<sup>2</sup> Lancaster University

<sup>3</sup> UW-Madison, Department of Computer Science

HTCondor Week 2014

28 – 30 April 2014, University of Wisconsin, Madison, WI, USA

