



UNIVERSITY
OF OSLO

ATLAS

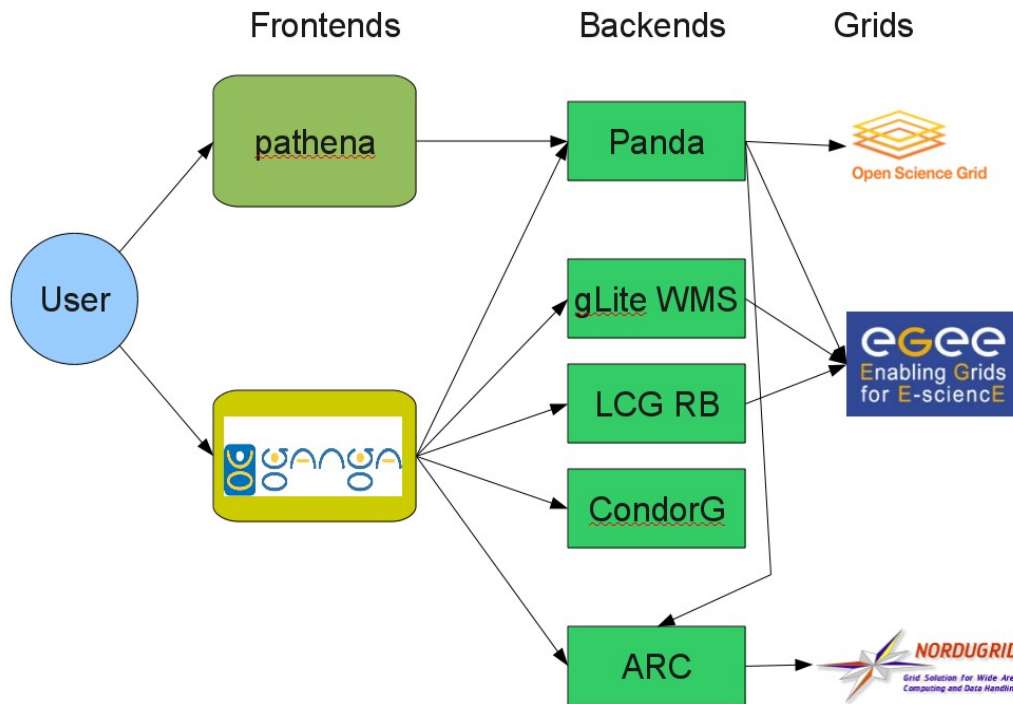
Distributed Analysis Support in the ATLAS experiment

EGEE09, 24.09.2009

Bjørn H. Samset, on behalf of the ATLAS DAST team



ATLAS analysis is complex



- ~2800 physicists
- Several hundred active grid users, more to come
- Large, official software kits
- Private code
- MC simulations
- Real data
- Databases
- Security
- ...

How do we give ATLAS physicists smooth and painless access to data and resources on the grid(s)? Tutorials, documentation, **support!**

DAST – the Distributed Analysis Support Team for the ATLAS VO

- DAST is the main ATLAS support hotline for distributed analysis questions
 - Expertise on ganga and pAthena
 - General knowledge of All Things Grid
 - Entry point for forwarding complex issues to relevant experts
- Mailing list is main point of contact
 - ~5 new threads per day, many taking days to follow up
 - Highest volume mailing list in ATLAS
- Work pattern
 - Two shifters each day, one per EU and US time zones
 - Shifters on for one week at a time
 - Mailing list (egroups) is the contact point
 - Common gmail account for tracking and answering issues
 - Weekly status meetings
 - Skype chatroom for live communication
 - Goal: 24h maximum wait for replies



Dick DASTardly
Unofficial DAST mascot



DAST gmail snapshot

Gmail [Calendar](#) [Documents](#) [Photos](#) [Reader](#) [Sites](#) [Web](#) [more](#) ▾

atlasdast@gmail.com | [Settings](#) | [Older version](#) | [Help](#) | [Sign out](#)


 Search Mail Search the Web [Show search options](#)
[Create a filter](#)

Compose Mail

Inbox

[Sent Mail](#)

[Drafts \(10\)](#)

[All Mail](#)

ACTIVE

[Cloud support](#)

[dCache problem](#)

[ESCALATED](#)

[NA investigating](#)

[Offline](#)

WAITING

[33 more](#) ▾

[Contacts](#)

[Tasks](#)

+ ● Atlas DAST

Search, add, or invite

+ Invite a friend

Archive Report spam Delete Move to ▾ Labels ▾ More actions ▾ Refresh 1 - 17 of 17

Select: All, None, Read, Unread, Starred, Unstarred

<input type="checkbox"/>	★	Helenka .. me, Sebastian (6)	ACTIVE	root files: pathena job - Hi again, so I did some source code research and I think I found my problem. As far as I can tell ...	10:17 am
<input type="checkbox"/>	★	Lauren, me (2)	ACTIVE	GGUS WAITING Cache Install Request? - Hi Lauren, for this a GGUS ticket with CC to atlas-grid-install has to be filed. I've done	Sep 20
<input type="checkbox"/>	★	Maria Moreno Llacer	ACTIVE	Submitted to aborted jobs for cosmic MC with release 14.5.2.12 - Dear all, I am trying to run over a cosmic MC dataset (valid2	Sep 19
<input type="checkbox"/>	★	Johannes, me (2)	ACTIVE	WAITING [Fwd: Ganga] - Hi Yahya, Have you tried: j.backend.sandboxcache = DQ2SandboxCache() Cheers, Dan On Sat, Sep 1	Sep 19
<input type="checkbox"/>	★	Fabrizio, me (12)	ACTIVE	WAITING job killed by Panda server : upstream job failed - Hi Jarka, I hope this is the command you wanted me to issue: [pfs2!	Sep 18
<input type="checkbox"/>	★	Mark, Jan, me (11)	ACTIVE	user2user Output files from jobs missing - Atlas DAST wrote: > Hi Mark, > apologies for delay... > > I found something weird ...	Sep 18
<input type="checkbox"/>	★	Eve, Bjrn, Andrej (6)	ACTIVE	some ROOT problem at ANALY_ARC - This jobs use 15.4.0 which I am not sure it works at all. Was never requested for installat	Sep 18
<input type="checkbox"/>	★	Lauren, Sergey (4)	ACTIVE	ANALY_LONG_BNL_ATLAS:lfc_creatg failed - Hi Lauren, I contacted Hiro about the LFC problem. Cheers, Sergey Lauren Alexand	Sep 17
<input type="checkbox"/>	★	Tulay, me (3)	ACTIVE	ESCALATED Ganga Expert WAITING exit 244 - Hi Jarka, The script to submit the job in ~tcuhadar/public/myTAGgangajob.py Cheer:	Sep 17
<input type="checkbox"/>	★	Wolfgang, Ulrik, me (7)	ACTIVE	ESCALATED Ganga Expert Ganga: job.copy() problem - Dear Ulrik, savannah bug report has been filed: https://savannah.cern.ch/	Sep 17
<input type="checkbox"/>	★	Wolfgang, me (2)	ACTIVE	WAITING Panda build job fails - Hi Wolfgang, I've seen more recent jobs succeeding. Is this still a problem for you? Cheers ...	Sep 17
<input type="checkbox"/>	★	Lauren, Sergey (2)	ACTIVE	BNL dccp fail: '\int\` object is not callable - Hi Lauren, it looks like it was a site problem. I talked to dCache people at BNL. There	Sep 16
<input type="checkbox"/>	★	Massimiliano, me, Bjrn (4)	ACTIVE	ESCALATED Ganga Expert How to peek job stdout runtime with Ganga - Hi Bjorn, >> Hi Max, >> you've found a rather old servic	Sep 16
<input type="checkbox"/>	★	George .. Di, Sergey (24)	ACTIVE	ERROR : failed to access LFC lfc.triumf.ca - Hi Di, Thanks for looking into it. So, it's a time out. I wonder why the first direct ...	Sep 15
<input type="checkbox"/>	★	Helenka, me, Johannes (7)	ACTIVE	(no subject) - Hi Helenka, Helenka Przysieznik Frey wrote: > Hi Johannes, > thanks for your reply. > I ...	Sep 14
<input type="checkbox"/>	★	Fabrizio, me (5)	ACTIVE	WAITING pilot: Put error: Copy command returned error code 256 - Hi, one of my jobs finished fine at ANALY_OX, but the pilot	Sep 14
<input type="checkbox"/>	★	Paul, me (4)	ACTIVE	EU investigating WAITING Wrong python version in athena job? - Hi Bjorn, I'm still getting the same problems with the new job. I r	Sep 8

Select: All, None, Read, Unread, Starred, Unstarred

Archive Report spam Delete Move to ▾ Labels ▾ More actions ▾ Refresh 1 - 17 of 17

Gmail works great for us!

Get **Gmail on your phone**. It's super-fast. Visit <http://mobile.google.com/> on your phone's web browser. [Learn more](#)

You are currently using 68 MB (0%) of your 7372 MB.

This account is open in 2 other locations. Last account activity: 1 minute ago at IP 147.188.47.199. [Details](#)

Gmail view: standard | [turn off chat](#) | [basic HTML](#) [Learn more](#)

©2009 Google - [Terms](#) - [Privacy Policy](#) - [Google Home](#)

Present DAST team

EU		US	
Daniel van der Ster	<i>CERN</i>	Nurcan Ozturk	<i>UT Arlington</i>
Mark Slater	<i>Birmingham</i>	Alden Stradling	<i>UT Arlington</i>
Hurung-Chun Lee	<i>NIKHEF</i>	Jacob Searcy	<i>Oregon</i>
Bjørn H. Samset	<i>Oslo</i>	Sergey Panitkin	<i>BNL</i>
Christian Kumner	<i>München</i>	Anirvan Sircar	<i>Louisiana</i>
Maria Shiyakova	<i>JINR</i>	Bill Edson	<i>SUNY Albany</i>
Manoj Jha	<i>INFN</i>	Karthik Arunachalam	<i>Oklahoma</i>
Jaroslava Schovancova	<i>Prague</i>		

- Highly international team, EU and US time zones
- First formed in summer 2008, mainly consisted of distributed analysis software developers
- Today: still some developers, but also many others who have either been self-taught or trained by the team
- Full credit as an ATLAS shift activity, though not regarded as “mission critical” shifts

Activities for a DAST shifter

- **Follow the mailing list, answer new items, follow up old ones**
- Follow the SAM and GangaRobot site-by-site analysis testing systems
- Report site problems of any kind that show up for users (file catalogs down, permission problems, bad files...)
- Forward complex problems to relevant experts
- Send test jobs to sites if requested by the worldwide computing operations team
- Maintain analysis FAQs and some documentation
- Encourage user2user support
- ...

Typical DAST questions

1) USE case

Q: Hi, I have a problem with my analysis.

A: *Please provide your logs and some info on your job*

Q: Oh, yes... Here they are.

A: *OK, you've filled the data member outputdata wrongly*

Q: Cool, thanks.

User education

2) Site problem(s)

Q: Waaah! My analysis ran fine yesterday, not today!

A: *Please provide your logs and some info on your job*

Q: Oh, yes... Here they are.

A: *You're forcing the job to INFN, they're in downtime. Try FZK?*

Q: OK, trying... No, still doesn't work. New error, though.

A: *Yes, FZK can't access your input files either. No downtime. I'll ticket the site. If you need the data quickly, please request a transfer to a third site.*

Site feedback

3) Feature request in disguise

Q: Hi, I'm using ganga like this, but it doesn't do what I want. What gives?

A: *Hmm, seems ganga isn't equipped to do just that. I'm adding a feature request to their bug tracker.*

Improved tools

Experiences with DAST in ATLAS

- DAST seems to be a successful activity
 - Good team, managed to attract nice and very competent people
 - Many questions, large knowledge base built
 - Very good response from users – almost too good
 - Users seem to get more competent
- We need to encourage user2user support
 - Many issues are “easy”, most experienced users can answer them quickly
 - Some users to help out, but not enough
- The level of knowledge needed to this kind of activity is quite high
 - User tools, underlying tools, sites, data management, databases, ...
- Knowledge of the technical aspects of grid computing is still essential for understanding problems on the analysis side

Future plans for ATLAS DAST

- Double shifts
 - When we get data in ~November, we anticipate a large influx of (inexperienced) users
 - From this point on, we'll have two shifters in each time zone
- Better information base, search functionality
 - There are mountains of ready-made solutions in our archives
 - Need to sort this better, make it more easily searchable
- Keep helping our physicists
 - ...and thus ourselves, most of us also being physicists...
 - Distributed analysis is crucial for ATLAS physics, but it's still quite complex

Conclusions

- ATLAS distributed analysis user support through DAST seems to work well
 - at the present level of activity...
- All activities will have to be scaled up soon
 - Can we handle the increased number and complexity of inquiries that we can expect?
 - Can we find and train enough shifters?
- Technically, we have the tools we need to help users
 - though better grid monitoring, site status, ..., is always welcome
- Contact the team directly here:
 - atlas-project-adc-operations-analysis-shifts@cern.ch