

# UK experience with publishing CE description in json format

*Alastair Dewhurst, Alessandra Forti*

# Motivation

- Flat cash funding means real term cuts to effort.
  - Across the UK 18 Site BDIIIs and 2 Top BDIIIs probably at least 0.5FTE looking after it.
- BDII is not fit for purpose:
  - New users (e.g. DUNE, SKA) are getting their information elsewhere.
  - It cannot be easily extended for new services.
- Already know we can manage without BDII:
  - New services (e.g. Echo) are not publishing to the BDII
  - US stopped supporting it.



# Types of information

- **Static**
  - Describes how a service is configured.
  - Changes rare, normally because of significant change.
- **Semi-Static**
  - Accounting information.
  - Space usage.
- **Dynamic**
  - Job status - This should be got directly from the CEs which should have been designed to cope with the number of jobs being run.





# Types of usage

- Functional
  - E.g. Test jobs - 1 an hour.
  - Reasonable to expect json to provide sufficient information to do this!
  - Site admin shouldn't have to be involved.
- Scale
  - E.g. Production work - 10 000 an hour.
  - Optimal setup will require dialogue between Site admin and experiment.
- We are not trying to build an information system that can describe all the nuances of every system.
  - Major changes to systems running at scale should always have human oversight! (e.g. SL6 → Centos?)





# New Information System

OSG



EGI



A link to a webserver hosting the json is stored in the GOCDDB.



# New Information System

OSG



- Both ATLAS and CMS are consuming information from OSG sites (via github) for over a year using new Information system.
- ALICE only ever used Dynamic information from the BDII.
  - This is now provided by CEs directly.
- LHCb DIRAC is able to consume information from json format using their GOCDDB2CSAgent.

HT-Condor  
CE

CREAM CE

ARC CE

ATLAS

CMS

LHCb

ALICE

Other

EGI



GOCDDB



# Security implications

- Main security monitoring EGI CSIRT uses relies upon Argo.
  - Migration to Argo ongoing.
  - Argo developers need to be able to understand new format.
- Other things to consider:
  - Make sure up to date information on sites for incident response.
  - Need to propagate suspension information.
  - BDII is used to identify endpoints for security service challenges.



# Manual vs Automated

- What if CE scripts to create json are not available?
  - Time taken to manually edit json file ~ 5 minutes.
  - The information is static, so only needs updating when site make a significant change.
- **Doing more than one update a month seems unlikely...**

HOW LONG CAN YOU WORK ON MAKING A ROUTINE TASK MORE EFFICIENT BEFORE YOU'RE SPENDING MORE TIME THAN YOU SAVE? (ACROSS FIVE YEARS)

	HOW OFTEN YOU DO THE TASK					
	50/DAY	5/DAY	DAILY	WEEKLY	MONTHLY	YEARLY
1 SECOND	1 DAY	2 HOURS	30 MINUTES	4 MINUTES	1 MINUTE	5 SECONDS
5 SECONDS	5 DAYS	12 HOURS	2 HOURS	21 MINUTES	5 MINUTES	25 SECONDS
30 SECONDS	4 WEEKS	3 DAYS	12 HOURS	2 HOURS	30 MINUTES	2 MINUTES
1 MINUTE	8 WEEKS	6 DAYS	1 DAY	4 HOURS	1 HOUR	5 MINUTES
5 MINUTES	9 MONTHS	4 WEEKS	6 DAYS	21 HOURS	5 HOURS	25 MINUTES
30 MINUTES		6 MONTHS	5 WEEKS	5 DAYS	1 DAY	2 HOURS
1 HOUR		10 MONTHS	2 MONTHS	10 DAYS	2 DAYS	5 HOURS
6 HOURS				2 MONTHS	2 WEEKS	1 DAY
1 DAY					8 WEEKS	5 DAYS

Not automating the json update will cost between 30 minutes and half a day of work for each site admin over the next 5 years.

Keeping the automated script running may use more time.

Running a BDII is significantly more effort than either.





# Info System updates

- ATLAS Grid Information System polls BDII every 2 hours.
- All updates are logged.
- Comparison between new OSG system and BDII has been performed over last 1.8 years of updates.
- New OSG system updates each Panda Queue on average once every 5 years.
- The vast majority of BDII changes are flip-flopping between two values.
- We assume a correct one and a “default” one when the automated updates breaks.

	<b>EGI (BDII)</b>	<b>OSG (New InfoSys)</b>
Number of Panda Queues	722	349
Maxwallclock time updates	32643	75
Status updates	92925	45
Maxcputime updates	31761	1
<b>Total updates per year</b>	<b>120 / PQ</b>	<b>0.2 / PQ</b>



# Practical experiences



Alastair Dewhurst, 11<sup>th</sup> September 2018

# GOCDDB

- Link to json is stored in GOCDDB in a “Property Extension”
- We need to agree on a name.
- Suggestion: “InformationSystem”
- All json could be extracted from GOCDDB via:

```
$ wget https://goc.egi.eu/gocdbpi/public/?method=get\_service\_endpoint&extensions=\(InformationSystem=\)
```

- Currently only RAL available.

Extension Properties			Export all properties
Name	Value	Edit	<input type="checkbox"/> Select All
InformationSystem	<a href="https://www.gridpp.rl.ac.uk/RAL-LCG2/RAL-LCG2_CE.json">https://www.gridpp.rl.ac.uk/RAL-LCG2/RAL-LCG2_CE.json</a>		<input type="checkbox"/>
 Add Properties			Select action... <input type="button" value="Submit"/>

[https://goc.egi.eu/portal/index.php?Page\\_Type=Service&id=782](https://goc.egi.eu/portal/index.php?Page_Type=Service&id=782)



# Json

First attempt looked like this:

```
"computingservices" [  
  {  
    "ce_id": 1234,  
    "ce_name": "arc-ce01.gridpp.rl.ac.uk",  
    "cs_endpointurl": "https://arc-ce01.gridpp.rl.ac.uk:2811/",  
    "cs_flavour": "ARC-CE",  
    "cs_version": "5.64",  
    "cs_jobmanager": "condor",  
    "cs_jobmanager_version": "2",  
    "cs_status": "production",  
    "cs_state": "production",  
    "cs_queue_name": "grid3000M",  
    "cs_queue_maxcputime": 0,  
    "cs_queue_maxwalltime": 345600,  
    "cs_queue_maxmainmemory": "8048",  
    "cs_queue_maxrunningjobs": "25000",  
    "cs_assigned_vo": ["ATLAS", "CMS", "LHCb", "ALICE"],  
    "cs_message" : "some free form text" }  
  ...  
],
```



# Initial feedback

- We can simplify the names of JSON since all the data is inside the "computingservices".
  - The "cs\_" or "ce\_" prefixes can be dropped.
- The "id" and "name" are unnecessary, we can identify it via the endpoint.
- We need to add "site" field to represent GOCDB/OIM site name.



# Json take 2

```
"computingservices": [  
  {  
    "site": "RAL-LCG2",  
    "endpointurl": "https://arc-ce01.gridpp.rl.ac.uk:2811/",  
    "flavour": "ARC-CE",  
    "version": "5.0.5",  
    "jobmanager": "condor",  
    "jobmanager_version": "8.6.9",  
    "status": "production",  
    "state": "production",  
    "queue_name": "grid3000M",  
    "queue_maxcputime": 0,  
    "queue_maxwalltime": 345600,  
    "queue_maxmainmemory": 8048,  
    "queue_maxrunningjobs": 25000,  
    "assigned_vo": ["ATLAS", "CMS", "LHCb", "ALICE"],  
    "message" : "some free form text"  
  },  
]
```



# More feedback

- Its better but the structure is still flat.
- We have several CEs, which all have access to the same batch resources - How do we publish this?
- This is not a new problem and Brian B / OSG have made a proposal.
- [https://docs.google.com/document/d/1pg\\_5Kibc\\_-Z4JF4\\_HJyW5xL6GVYKwXxOU7DXf2QP9Ag/edit](https://docs.google.com/document/d/1pg_5Kibc_-Z4JF4_HJyW5xL6GVYKwXxOU7DXf2QP9Ag/edit)



# Json proposal

```
"computingservices": [  
  {  
    "site": "RAL-LCG2",  
    "jobmanager": "condor",  
    "jobmanager_version": "8.6.9",  
    "maxcputime": 0,  
    "maxwalltime": 345600,  
    "maxmainmemory": 8048,  
    "maxrunningjobs": 25000,  
    "assigned_vo": ["ATLAS", "CMS", "LHCb", "ALICE"],  
    "message" : "some free form text",  
    "endpoints": [  
      {  
        "url": "https://arc-ce01.gridpp.rl.ac.uk:2811/",  
        "flavour": "ARC-CE",  
        "version": "5.0.5",  
        "status": "production",  
        "state": "production"  
      },  
      ....  
    ],  
  },  
  ....  
],
```





# Conclusions

- Discussion of json format has prevented a new information system for being rolled out.
  - We need to find a way forward.
  - Json does allow things to evolve.
- The information system needs to provide enough information to allow the automated submission of test jobs.
  - It should provide guidance for those running at scale.



# Backup



# UK BDI instances

Site	Site Bdi	Top Bdi
ECDF	info4.glite.ecdf.ed.ac.uk	
BHAM	epgr09.ph.bham.ac.uk	
Liverpool	hepgrid4.ph.liv.ac.uk	
Manchester	site-bdi.tier2.hep.manchester.ac.uk	top-bdi.tier2.hep.manchester.ac.uk
Sheffield	lcg.shef.ac.uk	
Durham	site-bdi.dur.scotgrid.ac.uk	
Brunel	dc2-grid-68.brunel.ac.uk	
RHUL	sbdi2.ppgrid1.rhul.ac.uk	
Cambridge	vserv02.hep.phy.cam.ac.uk	
Bristol	lgbdi02.phy.bris.ac.uk	
Sussex	grid-bdi-02.hpc.susx.ac.uk	
Lancaster	py-fjalar.hec.lancs.ac.uk	
QMUL	bdi02.esc.qmul.ac.uk	
Oxford	t2bdi06.physics.ox.ac.uk	
Imperial	bdi.grid.hep.ph.ic.ac.uk (2)	topbdi.grid.hep.ph.ic.ac.uk (2)
Glasgow	svr030.gla.scotgrid.ac.uk	
RALPP	site-bdi.pp.rl.ac.uk	
RAL Tier-1	site-bdi.gridpp.rl.ac.uk (2)	lgbdi.gridpp.rl.ac.uk (2)
<b>TOTAL hosts</b>	<b>20</b>	<b>5</b>

