



ASGC site report

Eric Yen, Felix Lee
Academia Sinica Grid Computing Centre (ASGC),
Taiwan

March GDB in Taipei



About ASGC

A little history re-cap



ASGC site report

- **2002:** ASGC was born and started to support WLCG since 2002.
- **2005:** Became APROC in Asia, and then became WLCG Asia T1 centre since Dec.
- **2007:** Completed first 10Gbp direct link between Asia and EU.
- **2008:** Completed the second CCRC'08
 - ASGC was rated as one of the most efficient and stable Tier-1 Center.
- **2009:** DC fire accident due to UPS.
 - We then stop using big UPS system ever since this accident



DC fire incident





DC fire incident

Me





ASGC site report

- **2010:** Started investing cloud computing and volunteer computing.
 - Hadoop, Opennebula, Openstack, BOINC
- **2011:** Started investigating Ceph.
 - Buggy as heck.
- **2013:** Re-investing Openstack and Ceph
- **2014:** [REDACTED]
- **2015:** [REDACTED]
- **2016:** Keep investing SDN and extending Grid/Cloud/HPC infrastructure to support Bioinfo, Neural science (e.g. brain image processing) applications.



ASGC site report

- **2017: Here we are, still alive**





Report summary

- DC introduction
- Network/ipv6 deployment
- Resources for WLCG
 - EGI, Cloud, HPC...
- Monitoring.
- Challenges.



ASGC Computer center

- **Total Capacity**
- 2MW, 400 tons AHUs
- 93 racks
- DC area takes ~ 800 m²
- **Resources**
- 20,000+ CPU Cores
- 12.5 PB Disk
- **Rack Space Usage percentage (Racks)**
- WLCG: 32.9%
- E-Science: 43.2%
- RCEC: 11.9%
- IPAS: 5.8%
- ASGC: 3.6%
- IES: 2.6%



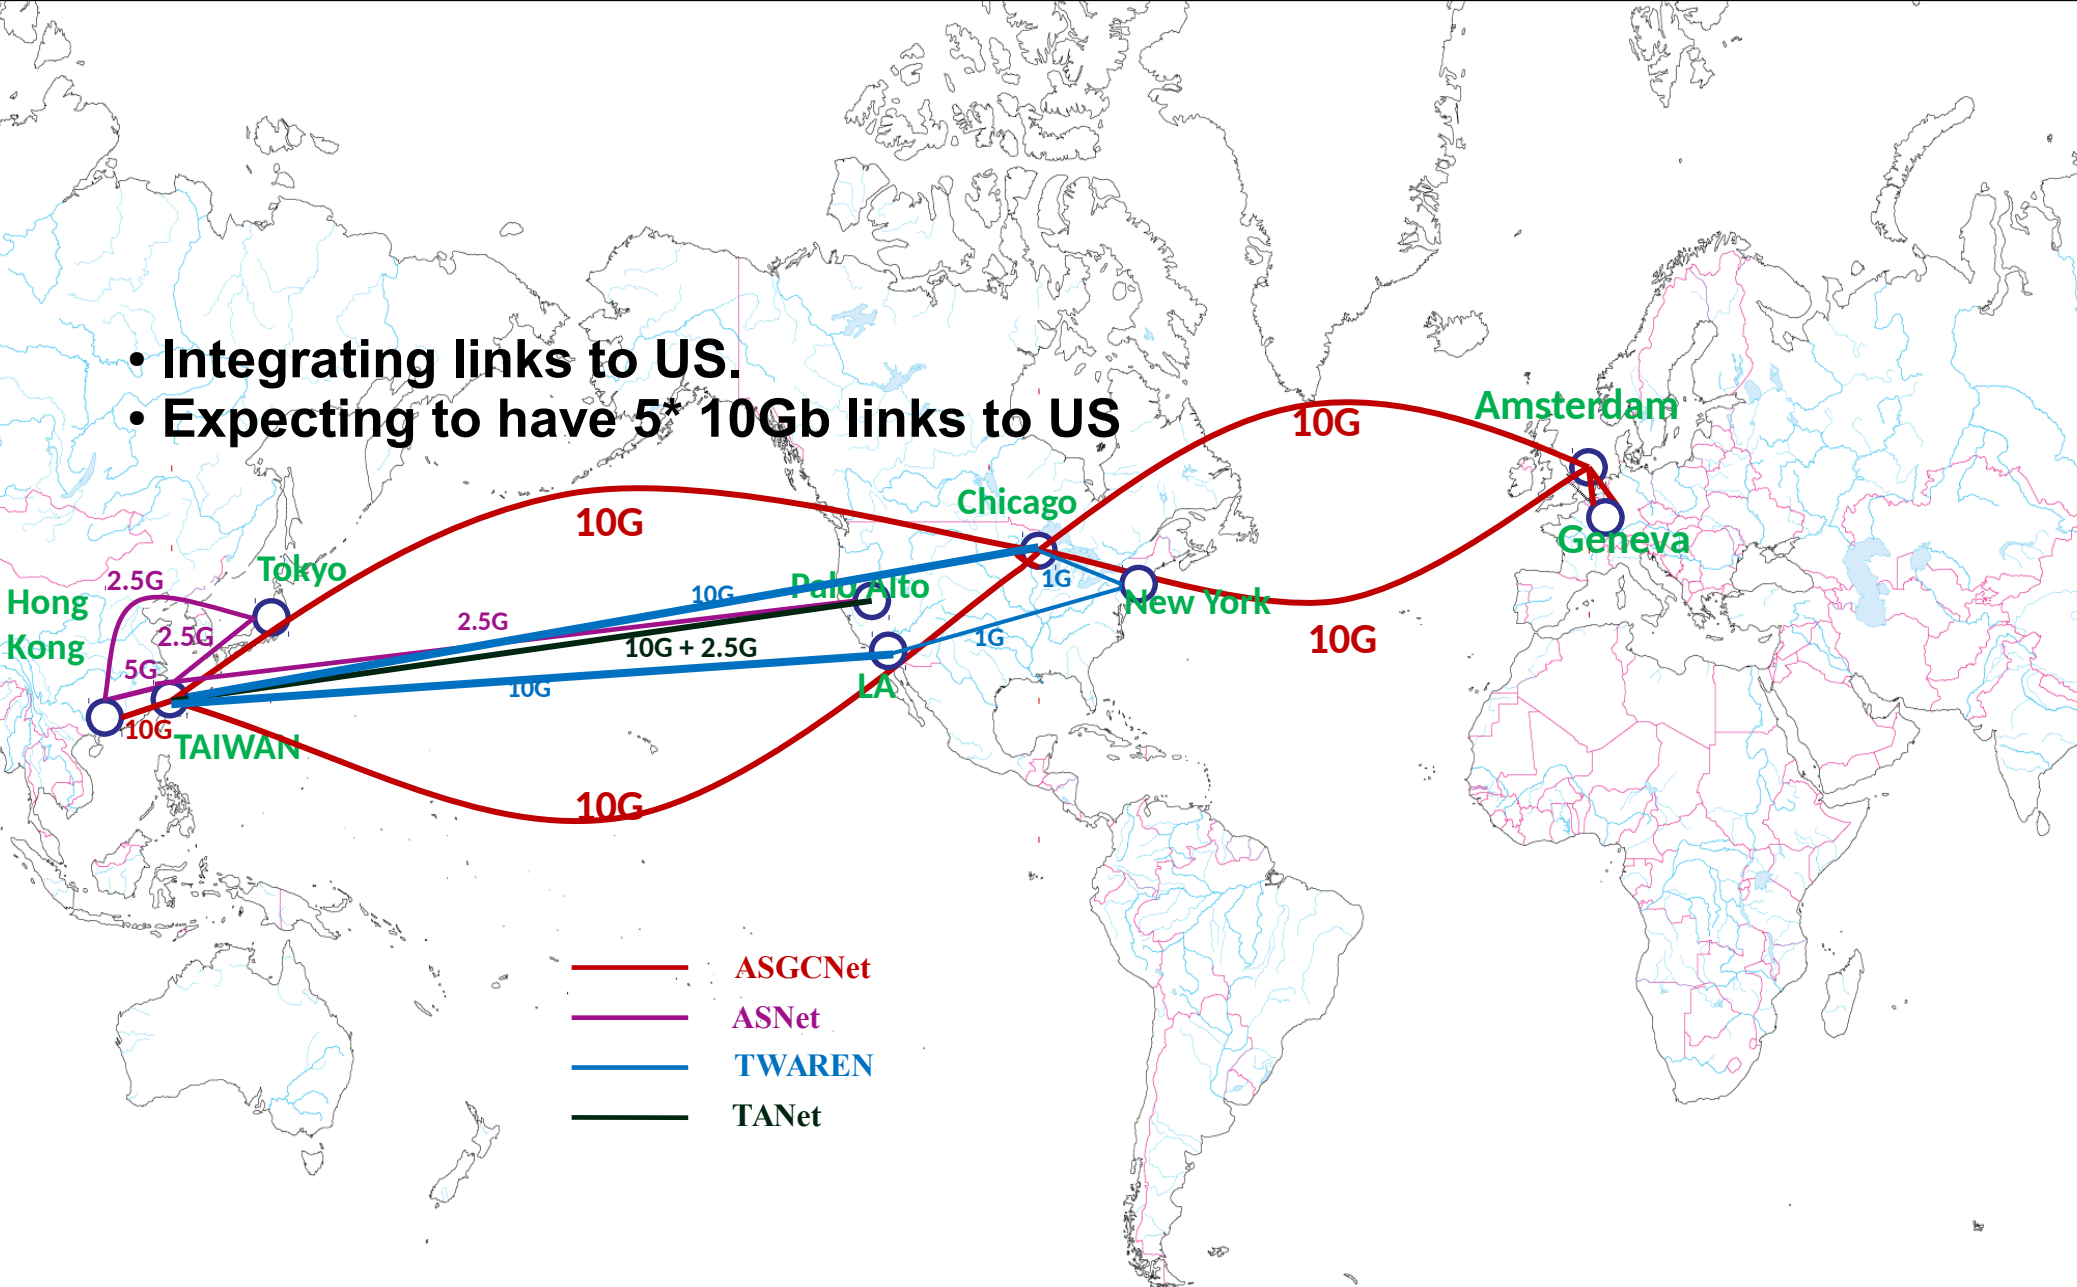
Monitoring the power consumption and temperature of every piece of equipment every 10 seconds.





Taiwan Global R&E network

- Integrating links to US.
- Expecting to have 5* 10Gb links to US





IPv6 Deployment

- We made it
 - After a long fight with our new Brocade router...
- Have established the BGP peering sessions
 - CERN, ESNET, GEANT, APAN-JP





CPU Resources for WLCG

- T1: 310 worker nodes, 2496 cores
 - Offers 38,002.20 HEP-SPEC06
 - Will soon reach the pledge for 2017.
 - 40,524 HEP-SPEC06
- T2: 94 worker nodes, 752 cores.
 - Offers 10,492.28 HEP-SPEC06
- Opportunistic resources:
 - Openstack cloud: 2256 cores
 - HPC clusters: 5000+ cores.



Storage Resources for WLCG

- DPM:
 - 3.2 PB, all are allocated to Atlas.
 - Direct links to router in order to get better performance
- EOS:
 - ~1PB.
 - Still under verification by Atlas.
 - Struggling at SRM...
 - Maybe will be all moved to DPM, if there is no more progress made.



Monitoring



Taiwan T1T2 - DPM Cluster Report for Fri, 13 Dec 2013 12:51:05 +0000

Get Fresh Data

Metric Last Sorted

[Physical View](#)

ASGC Grid > Taiwan T1T2 - DPM >

Overview of Taiwan T1T2 - DPM

CPU's Total: **236**
 Hosts up: **18**
 Hosts down: **0**





Avg Load (15, 5, 1m):
34%, 39%, 44%

Localtime:
2013-12-13 12:51



Cluster Load Percentages



- 100+ (5.56%)
- 75-100 (11.11%)
- 50-75 (22.22%)
- 25-50 (11.11%)
- 0-25 (50.00%)



Show Hosts: yes no | Taiwan T1T2 - DPM network_report last hour sorted descending Columns Size





Monitoring

- Nagios + SMS alarm notification to maintain 24/7 operation capability.

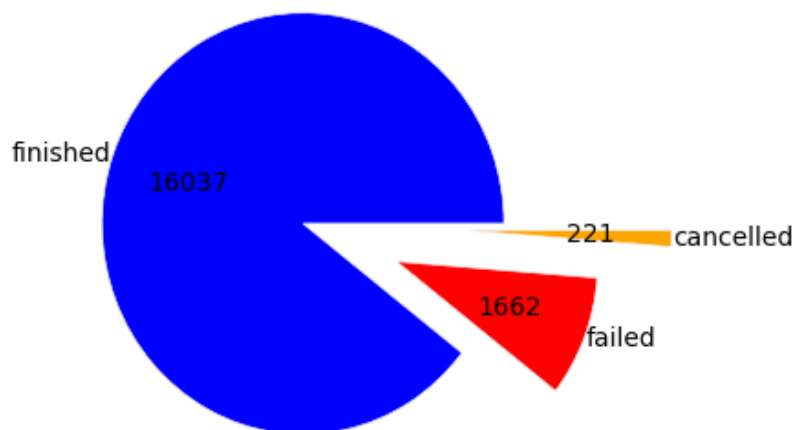
The screenshot displays the Nagios web interface. On the left is a navigation sidebar with sections: General (Home, Documentation), Current Status (Tactical Overview, Map, Hosts, Services, Host Groups, Service Groups, Problems), Quick Search, Reports (Availability, Trends, Alerts, Notifications, Event Log), and System (Comments, Downtime, Process Info). The main content area shows a list of hosts and their services. The top section shows a summary for 'FTS (fts)' with hosts as-sbdii, as-wn002, and nfs02. Below this, a detailed view for 'TWGrid FTT Site Service (ftt)' is shown, listing various hosts (apf2-tw01, apf2-tw02, apfmon-tw, db-backup, f-arc01, f-arc02, f-dpm000, f-dpmp22, f-dpmp23, f-dpmp24, f-dpmp25) and their associated services. Each service is represented by a colored box: green for OK, yellow for WARNING, and red for CRITICAL. Services like 'CRL' for apf2-tw02 and 'CA Packages' for f-dpmp22 are highlighted in yellow. The interface also includes 'Actions' icons for each host and service.



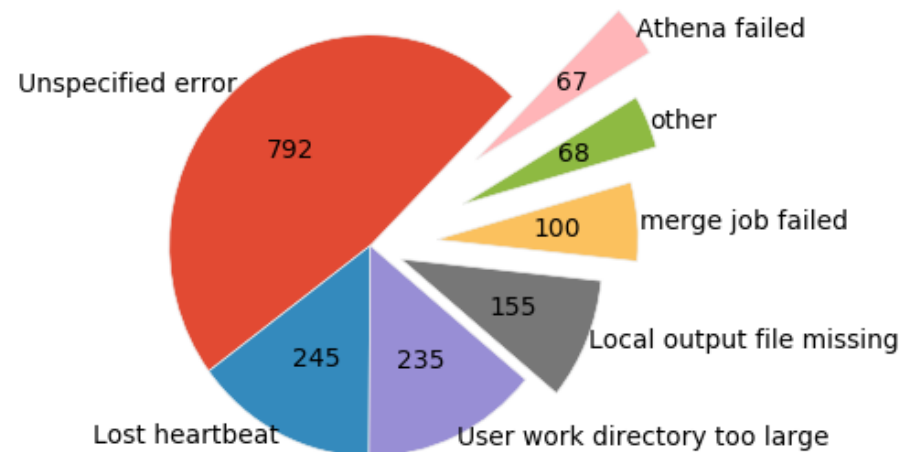
Monitoring

- Check job status, analyze the root cause and improve the condition every day.

jobstatus



errorlog

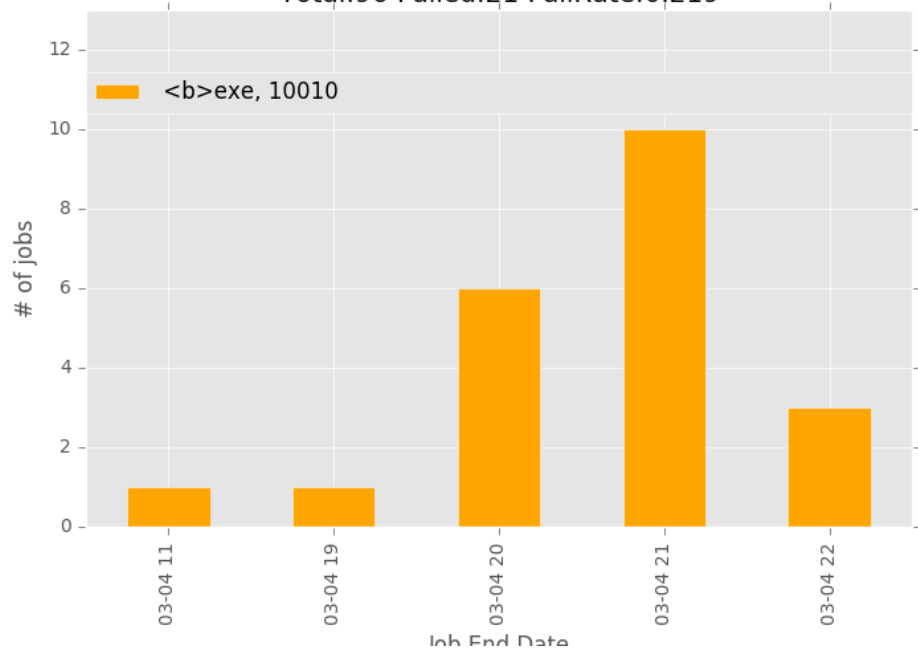





Monitoring

- Check job status, analyze the root cause and improve the condition every day.

Jobstatus of T1_SingleCore from 20170304 to 20170304 of user:jtanak
Total:96 Failed:21 FailRate:0.219



Fail job Top 10 Worknodes of T1_SingleCore
From 20170304 to 20170304 of user:jtanaka





Challenges

- Network is still our major problem for world wide grid.
 - Long latency issues..
 - Cisco N7K core switch has potential issue which would cause network performance drop to everywhere but local data center.
 - Its design is grouping 4 ports to share one buffer, which is not enough in certain condition.
 - e.g. long distance requires big tcp buffer, but big tcp buffer will make port buffer full quickly.
 - It's a bit hard to identify this issue, because everything is just fine but long distance transfer.
 - Thanks ESnet, OPN and WLCG-perfsonar people for helping us on this issue.
- Trying to reconnect disk server directly to router
 - Brocade router line card offers dedicated port buffer.



Challenges

- Tape system:
 - It's a bit expensive to maintain a decent tape system in Taiwan.
 - The cost of MA(Maintenance Agreement) per year allows us to buy additional 600+ TB disk per year.
 - Most of components (even technical support..) must be imported.
- Different timezone.
 - It wasn't big deal at beginning...



Thanks a lot!