



ASGC site report

Felix Lee

**Academia Sinica Grid Computing Centre (ASGC),
Taiwan**

HEPiX Spring 2014 in Annecy



ASGC site report

- Facility update
- Network(WAN) update
- Activities/events update for Grid/Cloud/HPC
- Software collaboration
- Hardware collaboration



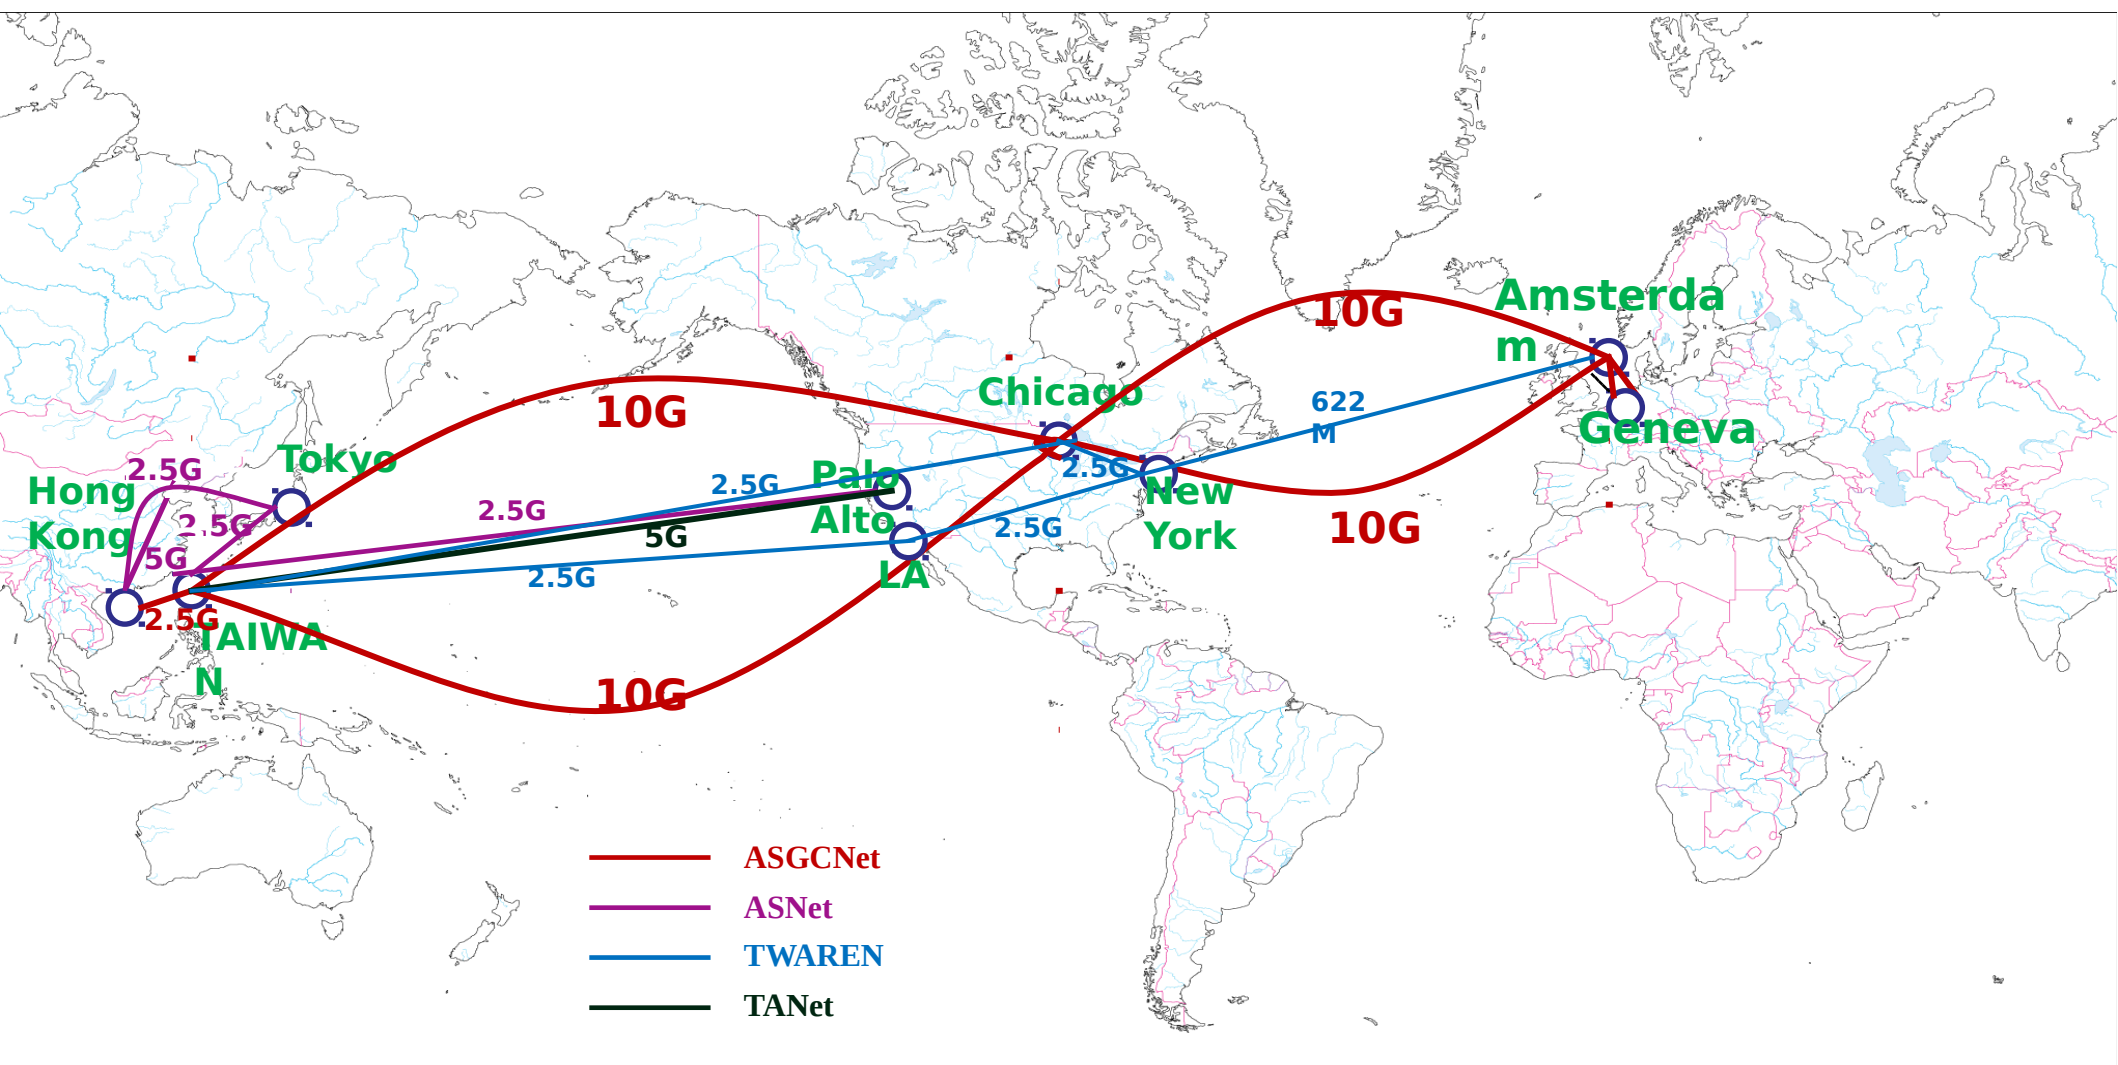
Facility update

- Still targeting high density solutions because the rack space is pretty tight in our DC.
- Another Dell C6220 * 54, 648 cores + Dell 8164F 10Gb edge switch – for Cloud as hypervisors.
 - E5-2630L, 48 GB mem, 10GbE.
- HP S6500, SL2300s Gen8 * 96, 1536 cores – for HPC resources.
 - E5-2650L, 64 GB mem, IB QDR.
- Dell MD3660f, 480TB + 640TB
- Now, we have 18k cores and ~11PB storage in total in DC.



Network update

- Upgrade 10Gb+2.5Gb -> 10Gb*2 to CERN.





Grid service update

- NTP DDoS vulnerability is pain...
 - Several stratum1 ntp servers are being effected in Taiwan.
 - Things are improved, but still getting instability sometimes...
- Migrating most of services to VM.
- Evaluating EOS.
 - Start with AMS, so far so good.
- Withdrawing from CMS T1.
 - There are still some fights for that., but no big hope..
 - Looking for another possibility to support CMS



Cloud service update

- Openstack Grizzly -> Havana
 - 108 hypervisors (1296 cores).
 - The upgrade to Icehouse is coming.
- OpenNebula 3.8.1(no more changes, waiting to be decommissioned)
 - There are 960 cores now, but would migrate some hypervisors to Openstack.
- Ceph is now in production.
 - Cinder + Ceph RBD is working well.

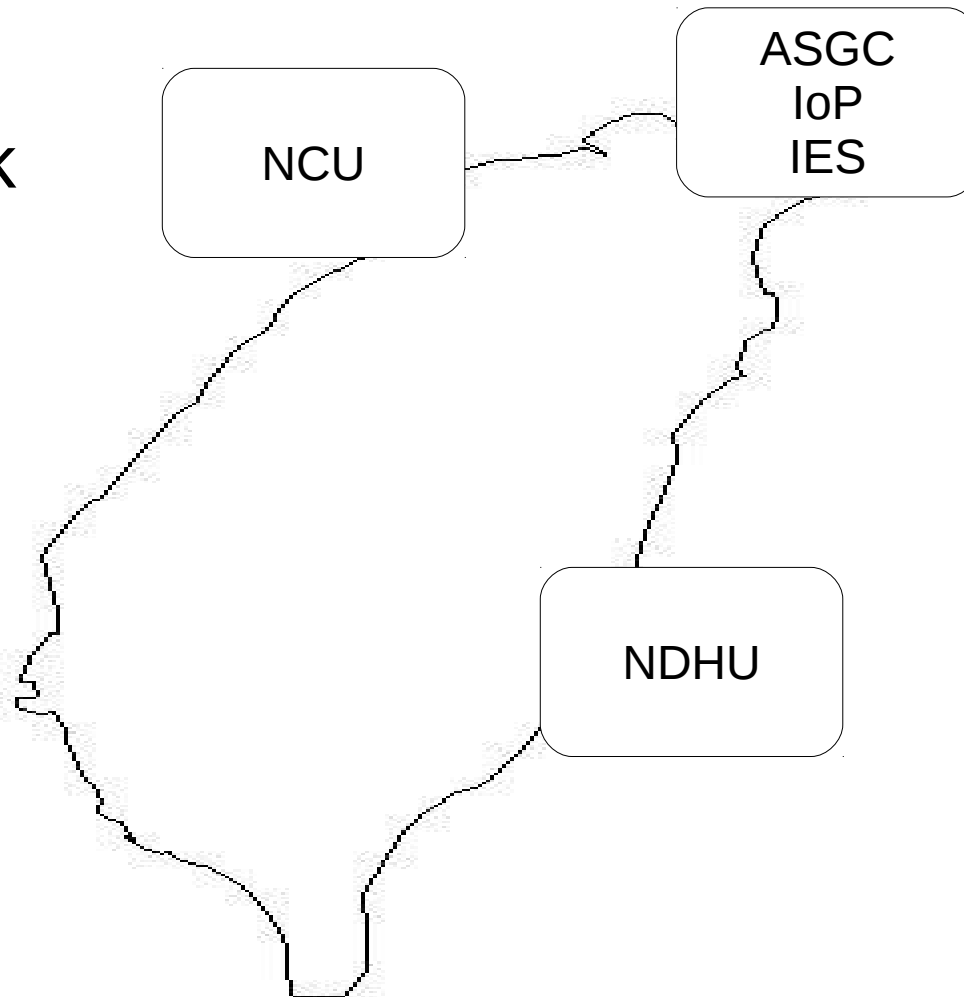


Cloud service update

- A problem with Dell 8154F 10GbE edge switch.
 - The 20Gb uplink to our core switch could only get 100kB throughput...
 - After some debugging, it turned out that was compatibility issue between Dell 8154F and our Cicso core switch Nexus 7000.
 - To upgrade Dell switch firmware to fix this issue.

Cloud service update

- Building cloud federation in Taiwan.
 - 5 sites so far
 - With Openstack





HPC service update

- The debate between 10GbE and IB.
 - We did a lot of comparisons but..
 - IB QDR won because of \$\$... if we don't buy big IB switch.
 - Now, we have
 - 1984 cores with IB DDR, 3712 cores with IB QDR.
 - 1980 cores for 10 GbE, but will buy more IB QDR in the future, until the price of 10GbE is lower...
- Turning IB cluster into Openstack cloud is a little bit challenging.
 - Always has some problems with nova network.
 - Our staffs are still working on it..
- Evaluating EOS for HPC as well.
 - Its low latency gains good potential for HPC.
 - The EOS fuse performance is quite competitive to Lustre.
 - Need to also work on its native API.



HPC service update

- An issue when we were purchasing new IB QDR cluster.
 - Vendor offered us: Voltaire Grid director 4036 QDR switch + connectX-3 IB card.
 - Things didn't work.
 - Then we figured out that Voltaire 4036 is totally incompatible with connectX-3...
 - Vendor admitted their mistake and replaced Voltaire 4036 by SX6036



Software collaboration

- Mainly, collaborating with Atlas and the BigPanDA project.
 - Helping on PanDA - RUCIO integration.
 - Also working on couple RUCIO components:
 - File consistency check
 - Object storage support, e.g. CEPH
 - Sending three people to CERN in order to work closely with Panda and RUCIO developers.
- A little work with CERN IT-PES for EGI cloud SSM2 accounting.



Hardware collaboration

- Fanless & UPSless computer facility design
 - Using static heat sinking.
 - Using battery as second PSU.
- Now we have more industry partners:
 - Supermicro, Trendmicro and some local manufacturers.
- 1st prototype was reported in HEPiX spring 2013
- 2nd prototypes are almost done:
 - Fanless rack
 - Fanless 50 bay storage
 - Fanless server.

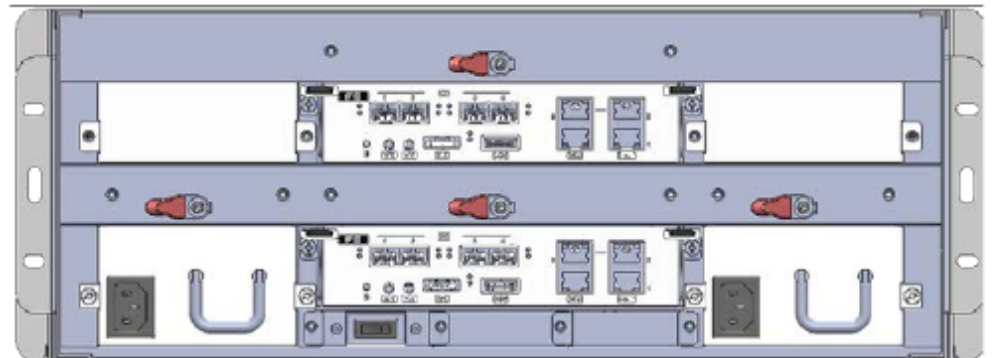
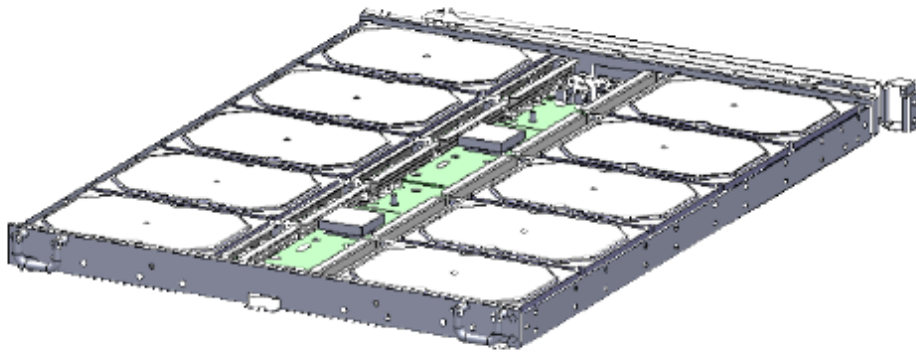
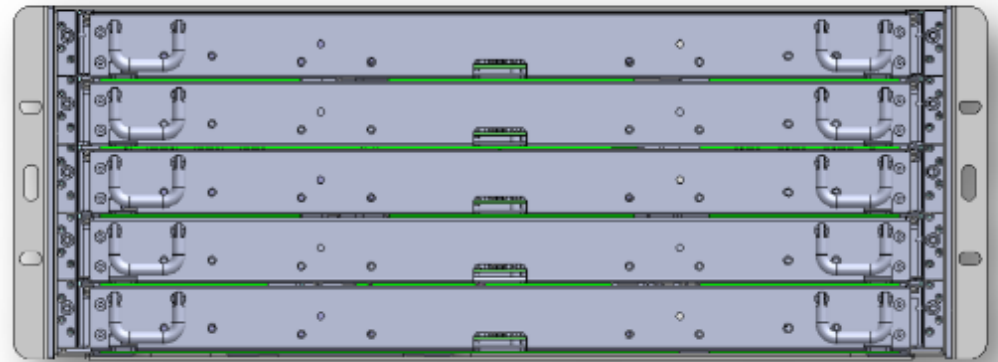
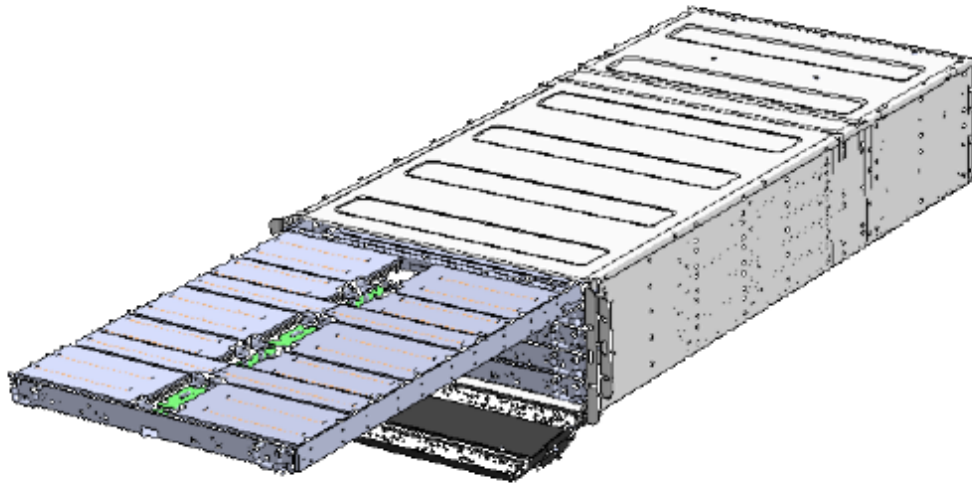
Hardware collaboration

- Fanless rack:



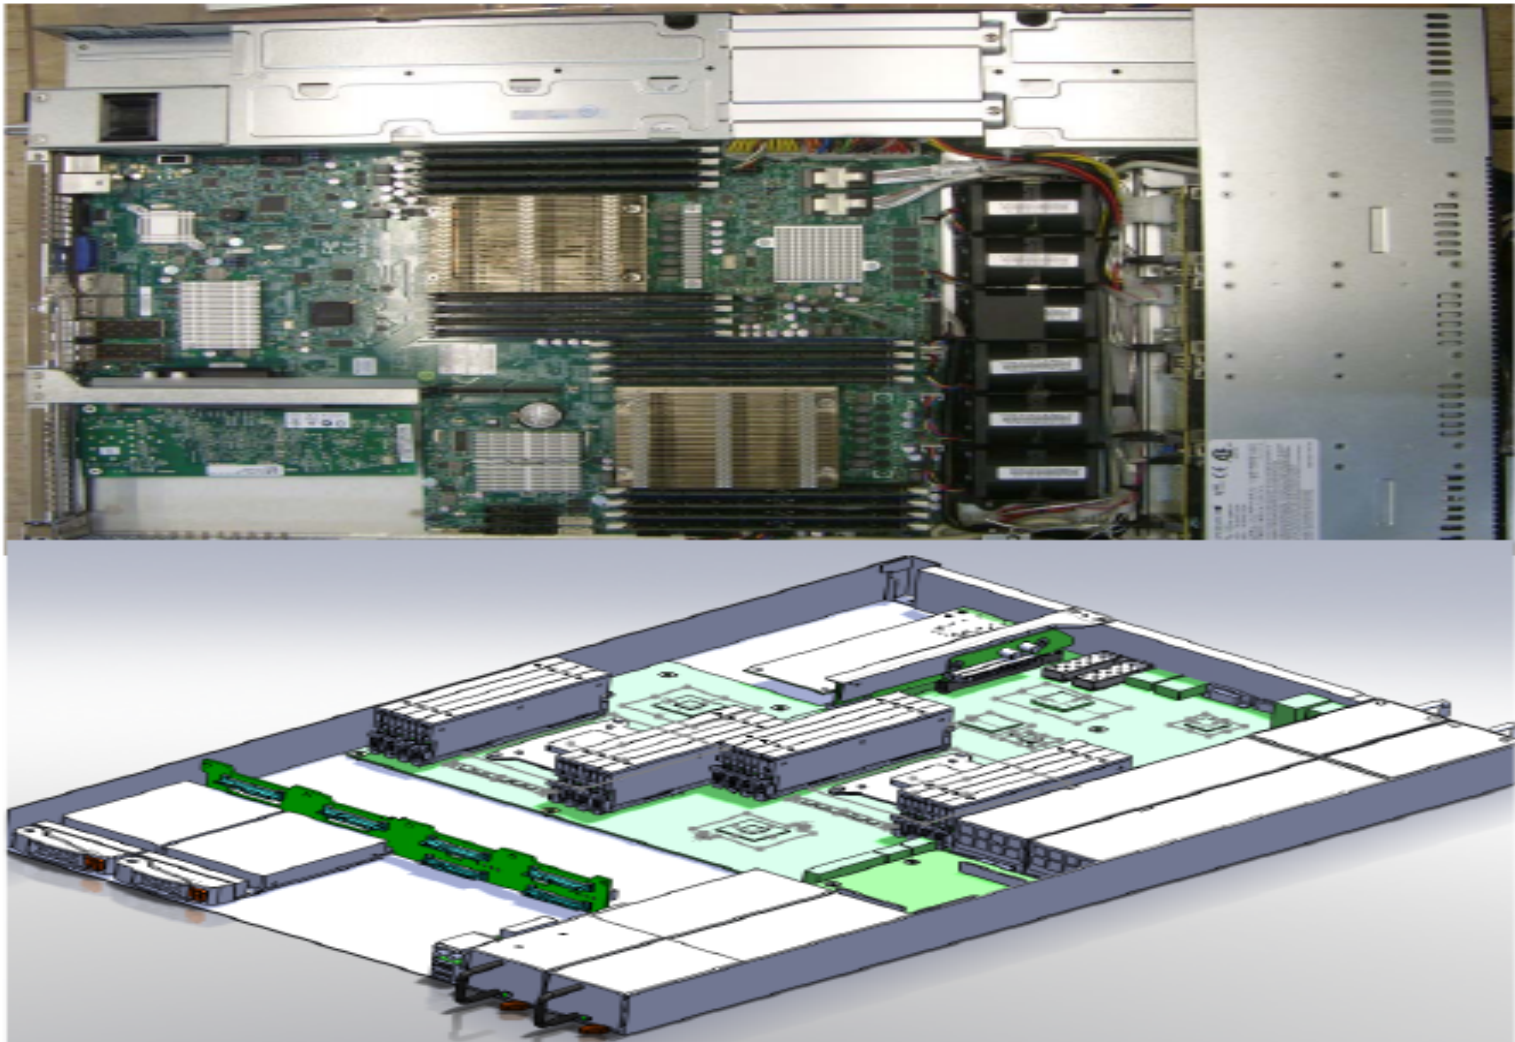
Hardware collaboration

- Fanless storage, 4U 50bays:



Hardware collaboration

- Fanless server, 1.5U:





Hardware collaboration

- Having a demonstration of UPSless servers along with Supermicro in Computex 2014.

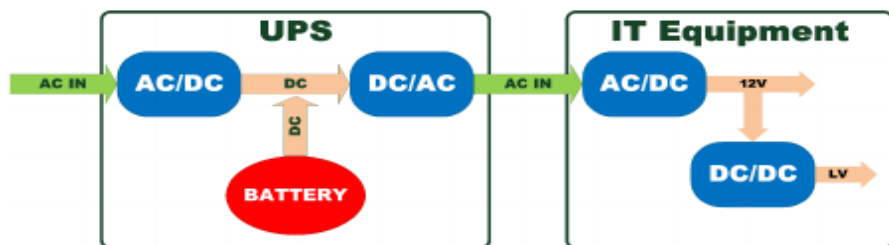
COMPUWARE
SuperMicro

ACADEMIA
SINICA

GCCA

Save energy by getting rid of your UPS

● *Traditional Power Distribution for IT equipment*



- 10%-20% of input power wasted.
- Heavy and dangerous.
- Only last for few hours normally.

○ *Solution – Battery Backup PowerSupply*



Thanks a lot!