



UiO : **Department of Physics**
University of Oslo

David Cameron
Andrej Filipcic



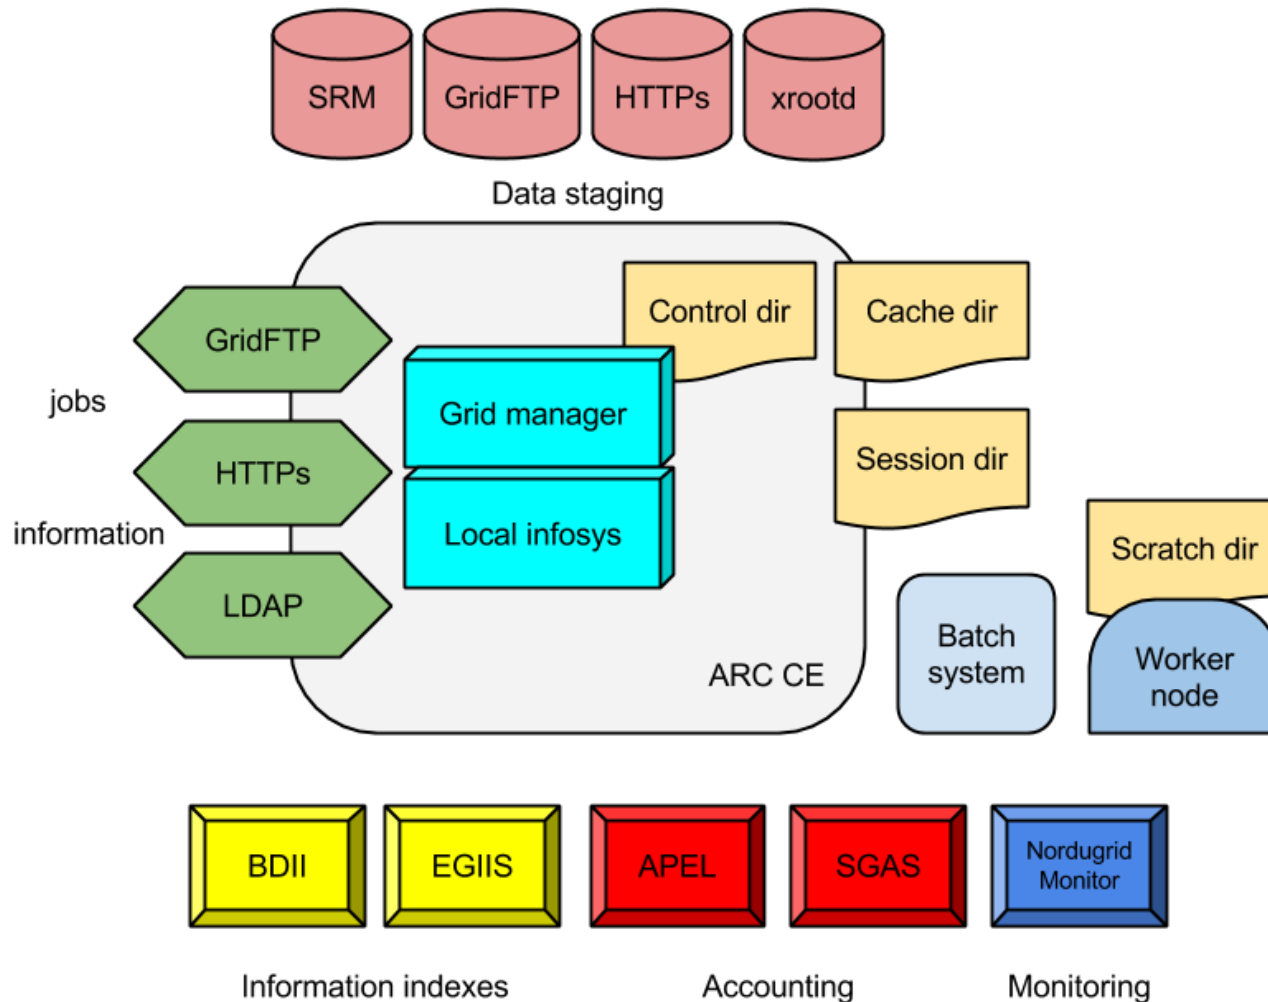
ARC Control Tower and ARC CE

ARC CE and NDGF

- Distributed dCache
 - Requires synchronous data transfer and local caching
- No middleware or connectivity from worker nodes
- ARC CE performs pre and post job data transfer
- Input files are downloaded to local shared filesystem cache

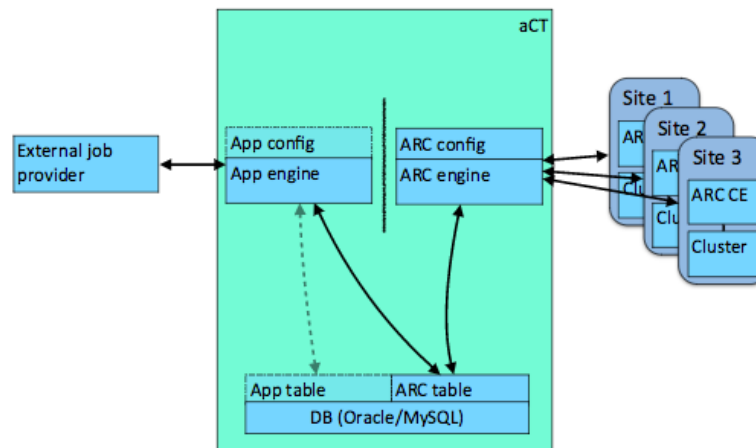


ARC CE Components

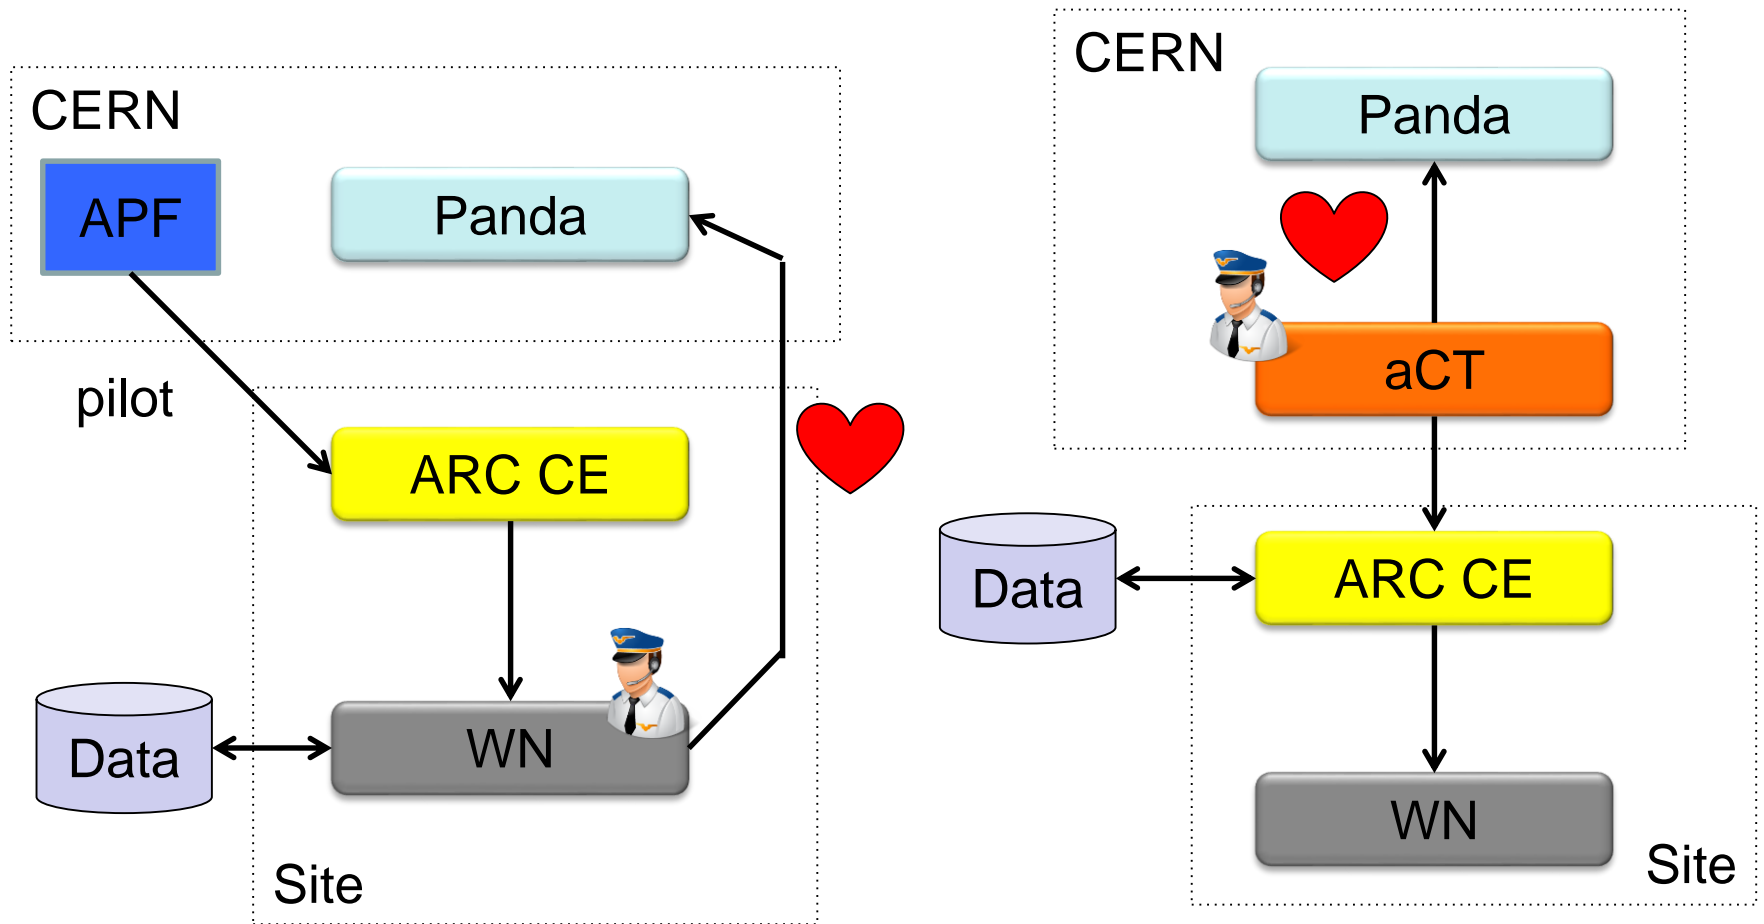


Need for ARC Control Tower

- The pilot model does not work well for NDGF
 - Requires middleware on the WN, outbound connectivity, synchronous data-staging
- Put a service in between Panda and ARC CE
 - A “fake fat pilot” pulling jobs, sending them to ARC and dealing with communication with Panda
 - Started as ATLAS-specific project, now generic



APF vs aCT job submission



Issues with pilot job framework

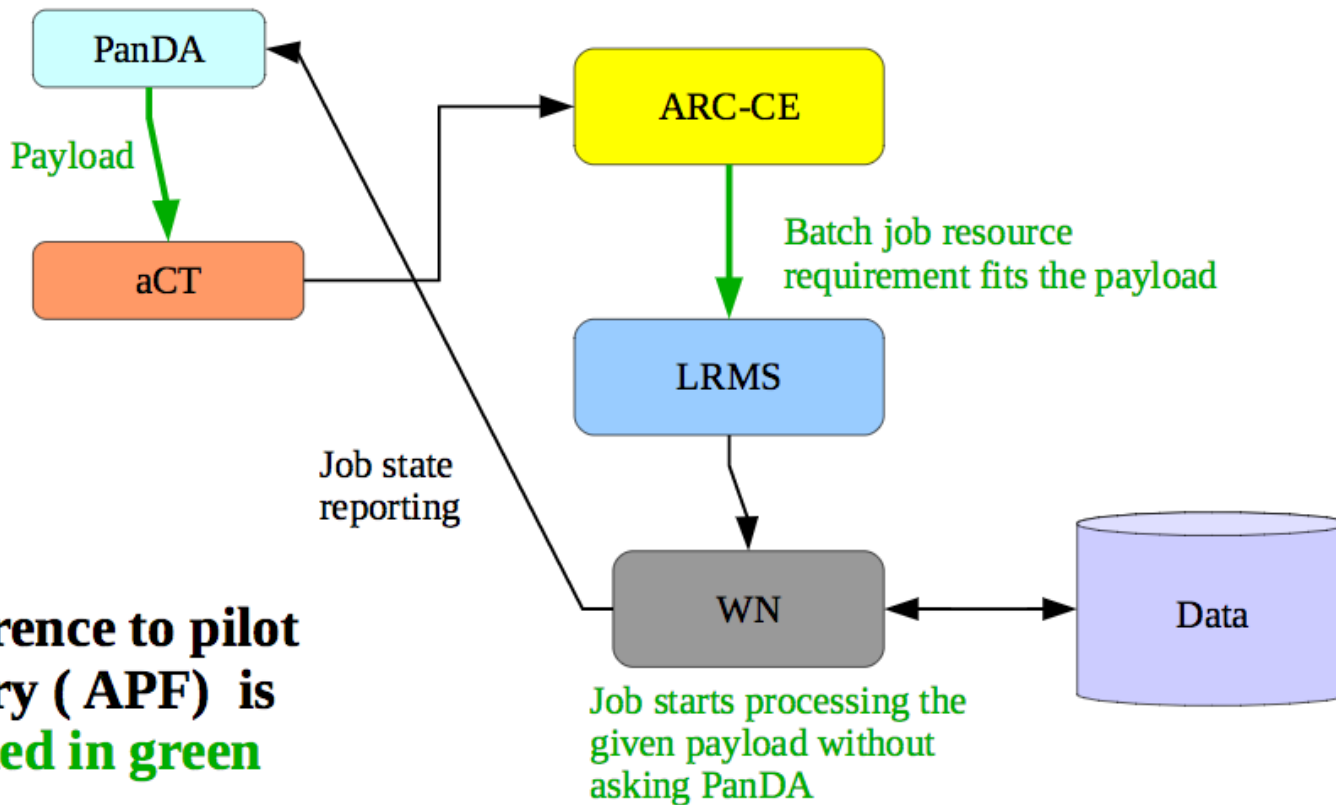
- ATLAS uses different workloads by memory, cputime, corecount requirements
 - But pilots are submitted with fixed requirements
- The only way to run different workloads is different Panda queues:
 - single-core
 - multi-core
 - analysis
 - hi-memory
 - very hi-memory
 - ...
- A task with radically new requirements requires a new Panda queue – heavy maintenance cost
- Fixed requirements lead to sub-optimal batch system scheduling

aCT “Truepilot” mode

- For sites running ARC CE who do not need the full “native” mode with staging etc
- aCT fetches the payload and submits it to the ARC-CE
- ARC-CE submits the batch job with predefined payload and requirements
- Pilot on the worker node does the same as on the conventional pilot sites, but skips the fetching of payload from PanDA
- aCT sends heartbeats to Panda up until job is running, then leaves it to pilot

aCT Truepilot

A. Filipcic, CHEP 2015



Difference to pilot factory (APF) is marked in green

ATLAS Collaboration

Slide: 9

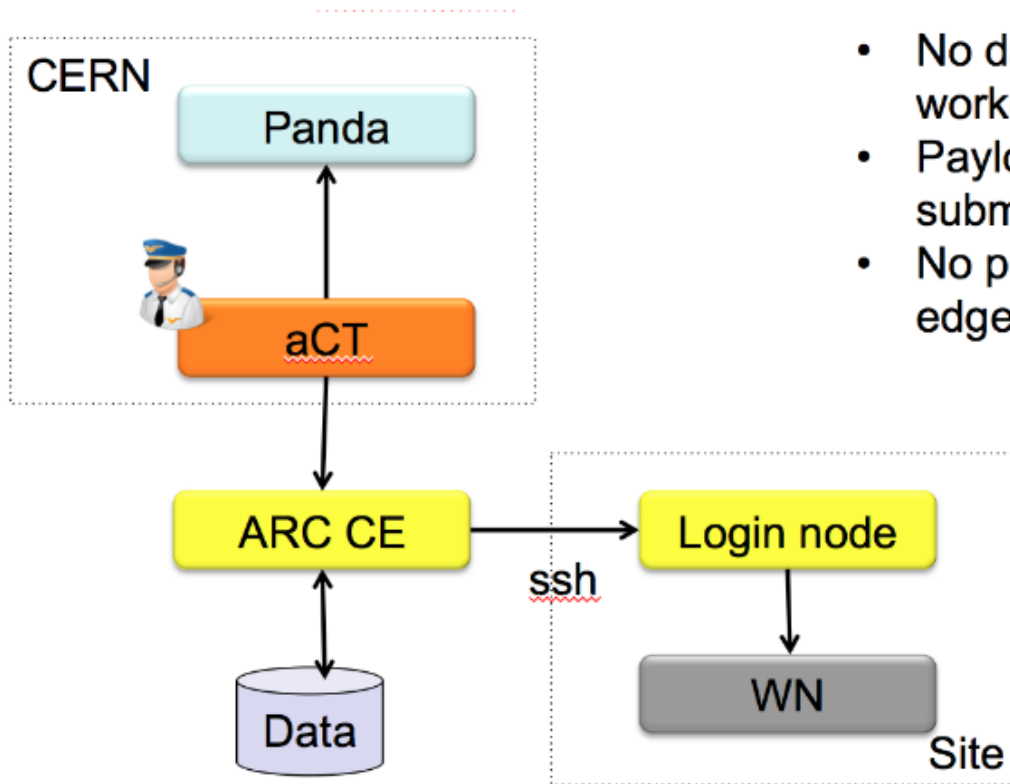
Pilot factory vs aCT Truepilot

- Pilot factory:
 - Highest priority jobs start running first
 - But the batch jobs have all the same resources
- aCT truepilot:
 - Payload known in advance – the batch job has the resource requirements fit to the job
 - Payload can request any memory, cputime, corecount, of course in agreement with site capabilities
 - But the late-binding is partially lost – highest priority jobs need to wait some time in the batch
 - Bad worker nodes can cause black holes – fast resubmission cycle

A. Filipcic, CHEP 2015

aCT and Supercomputers - HPCs

A. Filipcic, CHEP 2015



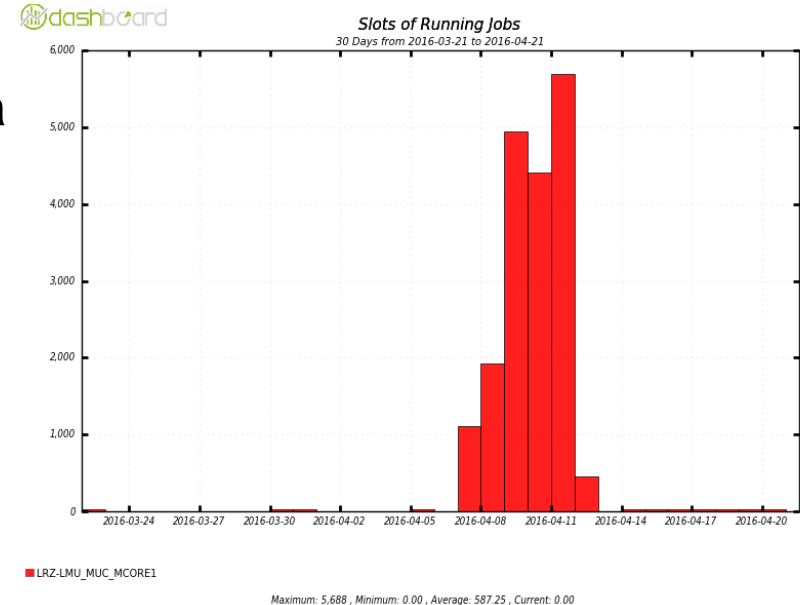
- No data access from worker node ✓
- Payload known at job submission ✓
- No persistent service on edge node or open ports ✓

Using aCT native node

Used at HPCs in Switzerland, Germany, China

aCT, ARC CE and Event Service

- Solution for HPC backfilling where there is no WN connectivity
- Instead of uploads to OS from WN, ARC CE uploads all events at the end of the job
 - This was shown to be a bottleneck – zipping by pilot can address this issue
- aCT reports completed events to Panda
- For jobs preempted by the batch system:
 - An ARC CE plugin “unfails” the job and does pilot post-processing (metadata xml, log tarball etc)
 - Tested in SuperMUC HPC (300 x 24 cores)



aCT service and sites

- 4 machines at CERN
 - 2 prod: large physical machines, one master, one hot spare
 - 2 dev: VMs, for testing, eg Event Service workflow
- Databases:
 - Prod machines: MySQL database on demand
 - Dev machines: local MySQL
- Serving ~12% of all ATLAS jobs:
 - ND, DE T2s, RAL, TW, BOINC, ...
- Very basic monitoring:
 - <https://voatlas403.cern.ch/data/aCTReport.html>
- Kibana service status probe is running (but reporting 100%)
- Source code (5k lines of python): <https://gitlab.cern.ch/arc/aCT>

Monitoring

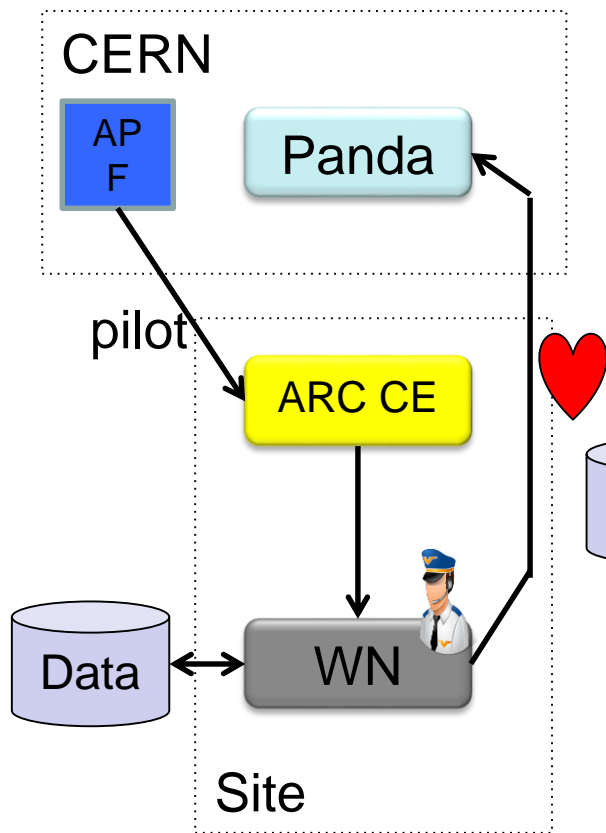
All ARC jobs: 26834

	Undefined	Accepted	Preparing	Submitting	Queuing	Running
WaitingSubmission:	1085	-	-	-	-	-
abisko-ce.hpc2n.umu.se:	1	-	-	-	-	31
alarik-grid.lunarc.lu.se:	-	-	2	-	152	22
arc-boinc-01.cern.ch:	5	-	-	-	1033	5502
arc-boinc-02.cern.ch:	-	-	-	-	28	-
arc-boinc-03.cern.ch:	20	-	14	12	971	5648
arc-ce.pleiades.uni-wupperta:	179	-	33	-	62	695
arc-ce01.gridpp.rl.ac.uk:	9	-	-	-	231	234
arc-ce02.gridpp.rl.ac.uk:	4	-	-	-	389	247
arc-ce03.gridpp.rl.ac.uk:	26	-	1	1	271	172
arc01.grid.sinica.edu.tw:	27	-	-	-	177	550
arc01.hpc.ku.dk:	-	-	109	-	49	183
arc01.lcg.cscs.ch:	-	-	-	3	-	-
arc02.grid.sinica.edu.tw:	13	-	-	-	168	534
arc02.ihep.ac.cn:	-	-	-	-	-	3
arc02.lcg.cscs.ch:	-	-	24	-	-	-
arc03.grid.sinica.edu.tw:	23	-	1	-	231	533
arc03.lcg.cscs.ch:	-	-	41	-	-	5
arcbrisi.cscs.ch:	2	-	-	-	-	-
atlas.triolith.nsc.liu.se:	58	4	116	-	14	259
c2papdata2.lrz.de:	-	-	-	-	154	280
ce01.grid.uio.no:	20	-	-	-	125	95
ce01.lhep.unibe.ch:	-	-	-	-	6	493
ce02.lhep.unibe.ch:	-	-	-	-	62	559
ce04.lhep.unibe.ch:	7	-	29	-	-	46
f-arc01.grid.sinica.edu.tw:	7	-	-	-	182	303
f-arc02.grid.sinica.edu.tw:	-	-	-	-	144	293
gar-ex-etpgrid2.garching.phy:	20	-	120	13	-	279
grid-arcce0.desy.de:	32	-	-	-	-	468
gridvm3.unige.ch:	-	-	-	-	-	24
jost.arnes.si:	74	-	-	-	118	222
lcg-lrz-ce0.grid.lrz.de:	-	-	-	-	161	223
lcg-lrz-ce3.grid.lrz.de:	40	-	-	-	3	36
nordugrid.unibe.ch:	2	-	1	-	110	18
pikolit.ljs.si:	37	-	69	39	90	740
Totals:	1691	4	560	68	4931	18697

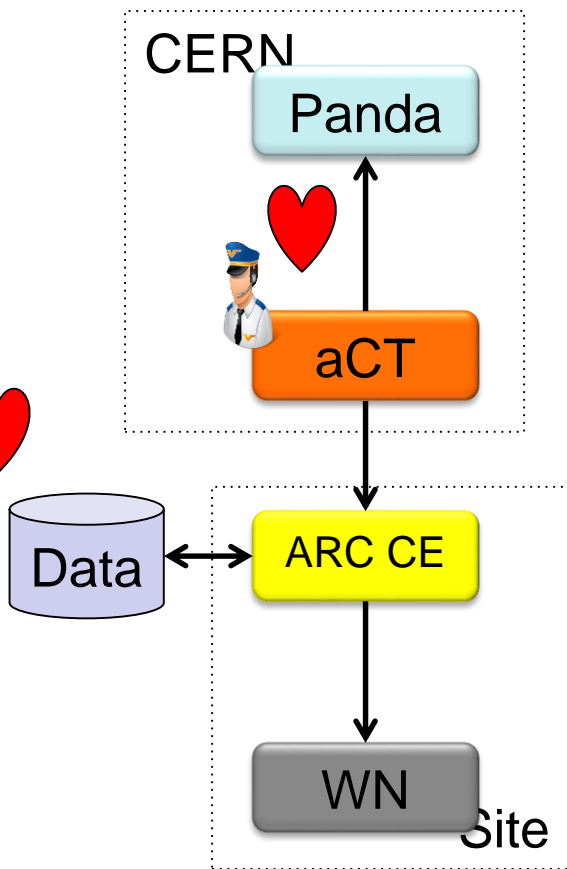
	sent	starting	running	toV
ANALY_ARNES:	21	72	130	
ANALY_CSCS:	2	49	50	
ANALY_CSCS-HPC:	2	1	-	
ANALY_DCSC:	21	74	13	
ANALY_HPC2N:	-	-	2	
ANALY_LRZ:	-	26	8	
ANALY_LRZ_TEST:	13	-	-	
ANALY_LUNARC:	23	98	19	
ANALY_NSC:	-	47	91	
ANALY_RAL_SL6:	-	42	418	
ANALY_SiGNET:	21	84	321	
ANALY_SiGNET_DIRECT:	7	16	27	
ANALY_TAIWAN_PNFS_SL6:	-	-	10	
ANALY_TAIWAN_SL6:	1	2	1	
ANALY_UIO:	21	68	24	
ANALY_UNIBE-LHEP:	-	14	-	
ANALY_UNIBE-LHEP-UBELIX:	-	-	2	
ANALY_wuppertalprod:	-	26	292	
ARC-TEST:	-	1	5	
ARNES:	26	100	-	
ARNES_MCORE:	29	19	73	
BEIJING-ERAII_MCORE:	-	-	-	
BEIJING-ERA_MCORE:	-	-	3	
BOINC:	30	2106	11146	
BOINC-TEST:	-	28	-	
CSCS-LCG2:	-	8	45	
CSCS-LCG2-HPC:	-	-	-	
CSCS-LCG2-HPC_MCORE:	-	-	-	
CSCS-LCG2_MCORE:	-	-	6	
DCSC:	26	84	174	
DCSC_MCORE:	-	2	-	
DESY-HH_Condor:	-	52	223	
DESY-HH_MCORE_Condor:	-	24	225	
HPC2N:	21	-	29	
HPC2N_MCORE:	21	-	1	
LRZ-LMU:	22	38	13	
LRZ-LMU_C2PAP:	28	-	-	
LRZ-LMU_C2PAP_MCORE:	21	149	296	
LRZ-LMU_MCORE:	24	102	199	
LRZ-LMU_MUC1_MCORE:	16	152	309	
LRZ-LMU_MUC_MCORE:	26	-	-	
LUNARC:	24	85	3	
LUNARC_MCORE:	-	51	-	
NSC:	7	74	49	
NSC_MCORE:	20	51	123	
RAL-LCG2_CEPH:	22	-	-	
RAL-LCG2_HIMEM_SL6:	29	-	11	
RAL-LCG2_MCORE:	30	880	191	
RAL-LCG2_SL6:	22	-	3	
RAL-LCG2_VHIMEM:	2	-	4	
SiGNET:	27	18	164	
SiGNET_MCORE:	25	115	373	

Summary of aCT/ARC modes

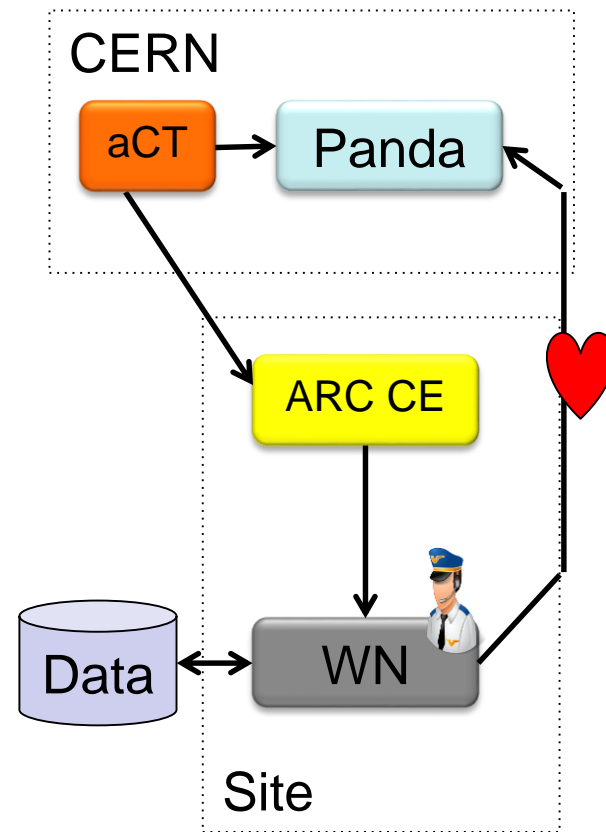
1. APF



2. NorduGrid



3. aCT Trueepilot



Current Issues and Future Work

- Lost heartbeat problem:
 - In true pilot mode, bad nodes can cause job to fail before any heartbeat is sent
 - aCT doesn't send any heartbeats or resubmit jobs so info/logs are lost (in fact aCT has the logs but can't tell panda where they are)
- Top priority jobs
 - Inherent latency in the model but not really an issue for high-throughput sites where batch system priorities work
- Missing software release/badly matched job requirements can block a site
 - Improve monitoring and reporting back to panda
- Improve service resilience
 - Remove afs dependency
 - Automatic failover to hot spare
- Monitoring improvements
 - Jobs and monitoring of service itself
- Move more sites to aCT – reduce panda queues!

References

- CHEP papers:
 - <http://iopscience.iop.org/article/10.1088/1742-6596/664/6/062042/pdf>
 - <http://iopscience.iop.org/article/10.1088/1742-6596/664/6/062042/pdf>
- CHEP talks:
 - <https://indico.cern.ch/event/304944/contributions/1672319/attachments/578552/796696/chep2015act.pdf>
 - https://indico.cern.ch/event/304944/contributions/1672645/attachments/578833/797024/ARC_Control_Tower_-_CHEP2015v5.pdf
- Operations twiki:
 - <https://twiki.cern.ch/twiki/bin/view/AtlasComputing/ARCControlTower>