

ARC-CE Patch Policy by/for/in the WLCG

WLCG Ops Coord Mai 2017

Max Fischer

ALICE site contact and ARC-CE/HTCondor guy at FZK-LGC2

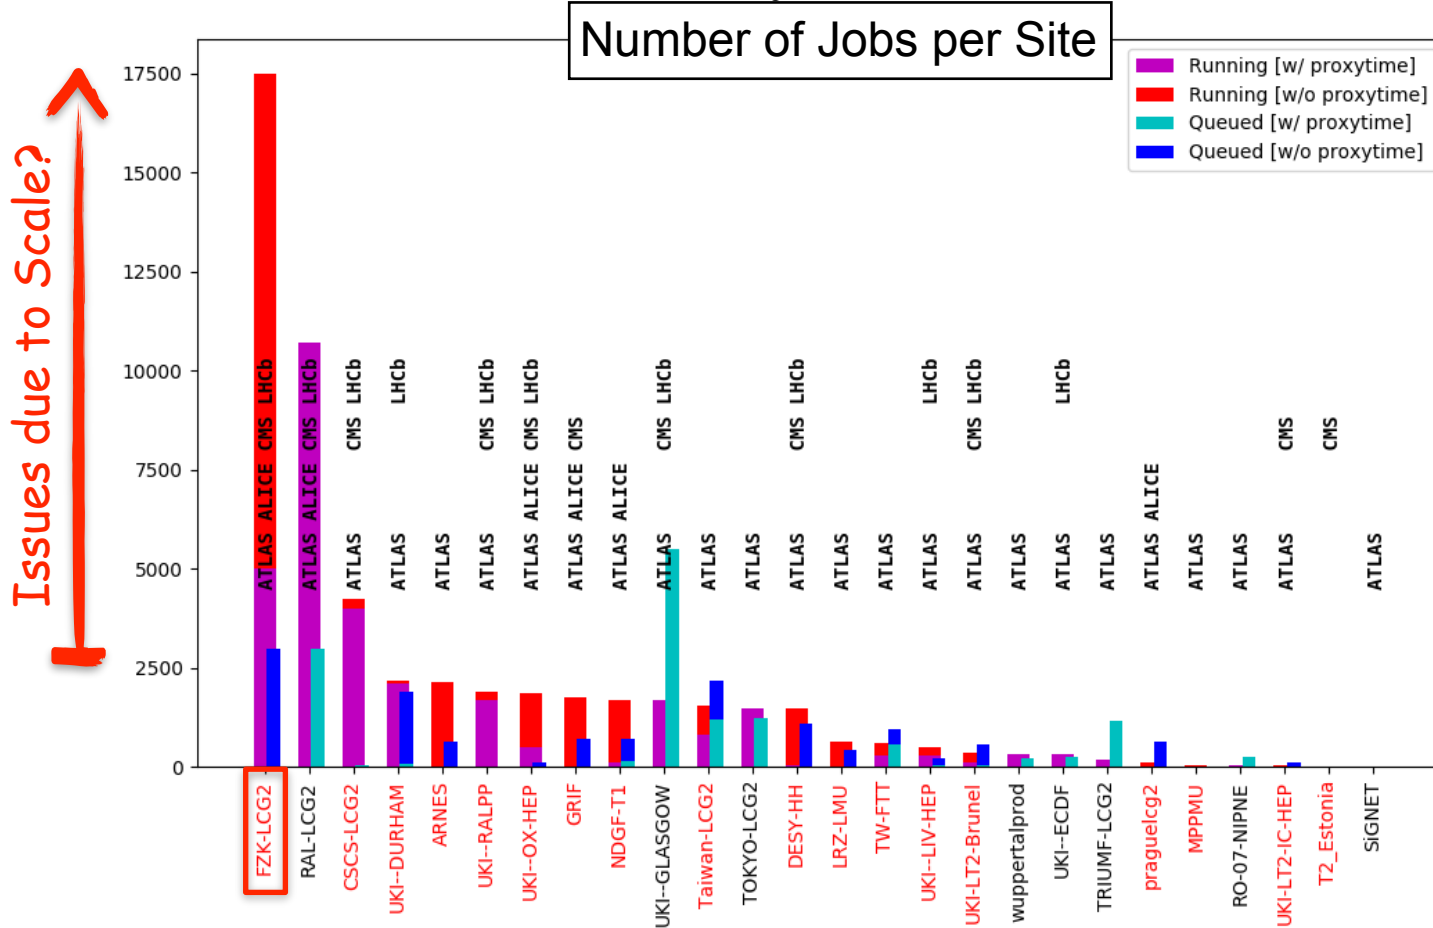
Steinbuch Centre for Computing



Image By Matt Howard [CC-BY-SA 2.0 (<http://creativecommons.org/licenses/by-sa/2.0/>)], via Wikimedia Commons

Occasion for this talk

- Sites with **red name** are technically broken...



- ...running an unmodified, official release of ARC-CE

ARC-CE at FZK (as relevant here)

- FZK semi-recently switched to ARC-CE + HTCondor
 - Based on good feedback from community
 - Switched off CREAM-CE end of March
 - *We consider ourselves late to the party...*
- Encountered lots of technical issues
 - Infosystem unsuitable for some VOs
 - Services and clients regularly die/hang/deadlock
 - *...and by default assume we messed up*
- Applied custom patches ^[1] requested by VOs
 - Site contacts mentioned/found out only after things broke down
 - Overhead for upgrades/debugging, works as temporary fix
 - *...fine for us after assurance that patches will go upstream*

[1] https://www.gridpp.ac.uk/wiki/Example_Build_of_an_ARC/Condor_Cluster

Case Study: ProxyExpirationTime

- FZK site contact for a VO reports broken Infosystem
 - *ProxyExpirationTime* of jobs not published by CE
 - Urgent: client constantly „updates“ proxy and frequently deadlocks
 - Site contact assures problem occurs only with FZK CE
- Roughly 2 weeks trying to locate the issue
 - Easy to reproduce - affects most jobs from all VOs
 - Not *reliably* reproducible - 10% of jobs just work
 - Check implemented/known custom hacks, docs, scripts, code, ...
- Preliminary results: Frustration
 - VO still has to babysit/restart/hack its client, issues only at FZK
 - Site staff doubts determinism of the universe...

Case Study: ProxyExpirationTime


- Question to the world: How does this even work? [answers paraphrased]
 - Site: „It actually does not. No one complained, so we never looked.“
 - Devs: „[Holiday pause] We made a patch just now in a few hours.“
 - FZK: „Wait, what?“
- Post-Mortem Debugging
 - **Problem exists for all ARC-CE versions in the last 6 months** *likely much longer*
 - Functionality was „used“ by at least one VO client tool
 - Functionality was patched on several WLCG sites
- Sites did not know others had the same problem or fixes
- Site contacts did not know others had the same problem or fixes
- Developers did not know there was a problem
- *We wasted precious time needed for debugging other problems...*

Personal [Site + Experiment] Perspective

- ARC-CE in WLCG has a problem due to technical debt
 - Manually maintained patches delay trouble shooting + updates
 - Site/VO specific hacks means duplicate work (of dubious quality)
 - *We probably spend most time chasing known problems*
- Custom patches not acceptable for either site or experiments
 - Sites need official fixes to avoid technical debt
 - VOs need official fixes to rely on features and services
 - *Temporary hacks must be semi-official, push for official fixes*
- Current problems likely hurt badly due to scale
 - FZK-LGC2 runs 2-10x jobs versus other sites (due to VO mix)
 - Issues will (probably do) affect other sites now or in the future
 - ***Need to debug/fix/address such issues at WLCG scale***

Conclusion for this talk

- Sites with **red name** are technically broken...

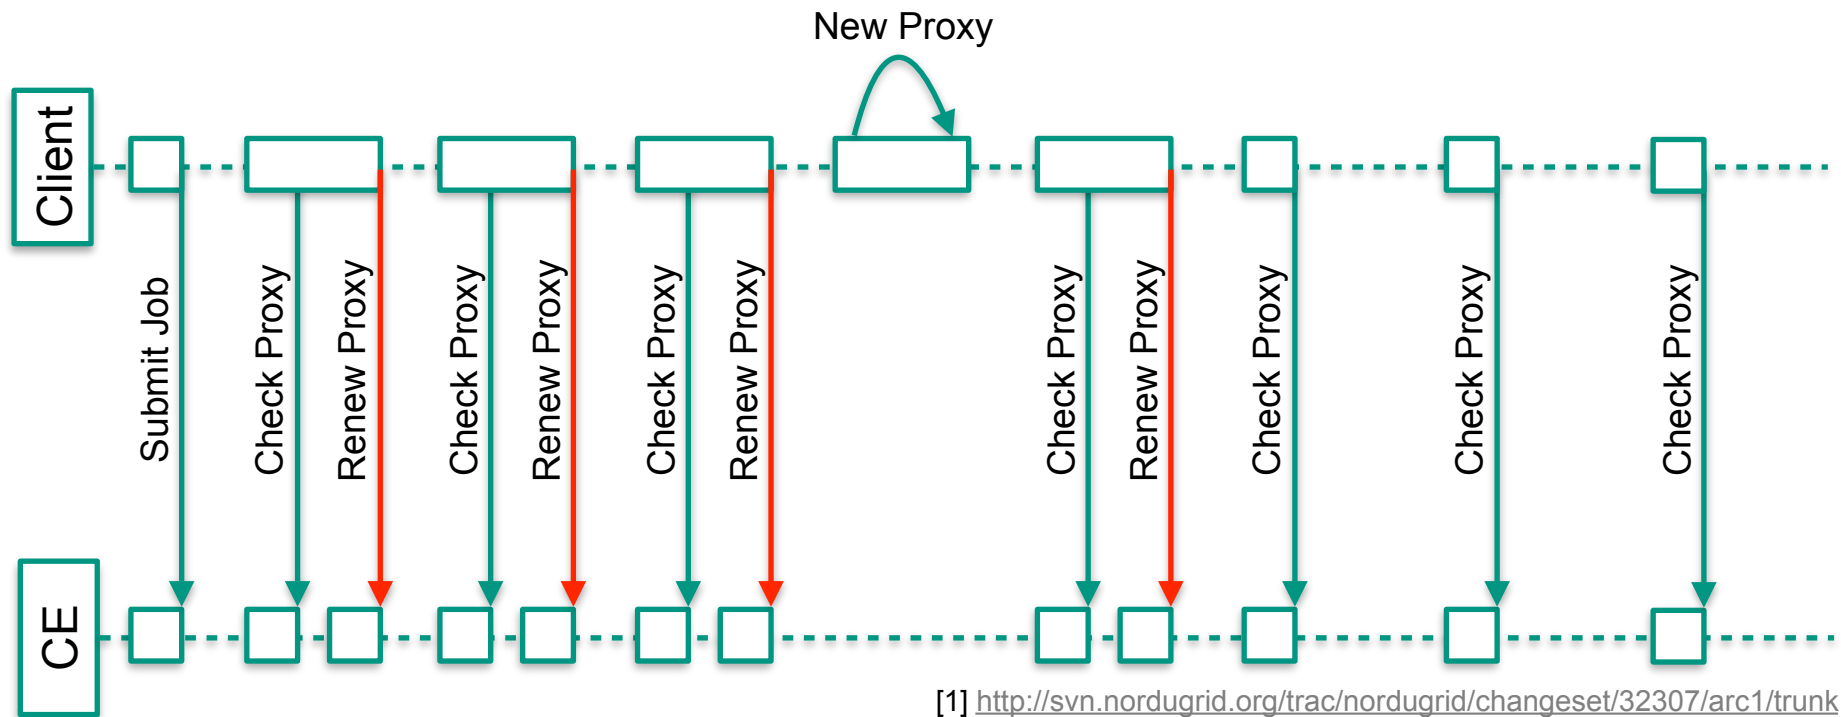


- ...what can we do about it in the long run?

Backup

Backup: Proxy Lifetime/Update Issue

- Job proxy lifetime not visible after submission^[1], only after proxy renewal
- ARC-CE only accepts proxy with longer lifetime for proxy renewal
- Client sees 0 proxy lifetime, constantly renews with same proxy
- Client can **deadlock** during proxy renewal^[2]



[1] <http://svn.nordugrid.org/trac/nordugrid/changeset/32307/arc1/trunk>

[2] http://bugzilla.nordugrid.org/show_bug.cgi?id=3534