


Tier-2 Evolution @ Glasgow

Gareth Roy
GridPP36

Outline

- Local Evolution
 - Stimulus
 - Mutation
 - Crossover
- Technological Evolution
 - Inheritance
 - Selection



```
def getSolutionCosts (navigationCode):  
    fuelStopCost = 15  
    extraComputationCost = 8  
    thisAlgorithmBecomingSkynetCost = 999999999  
    waterCrossingCost = 45
```



GENETIC ALGORITHMS TIP:
ALWAYS INCLUDE THIS IN YOUR FITNESS FUNCTION

Local Evolution

Stimulus: University of Glasgow - Master Plan

I am convinced that this is as significant a development in the on-going story of the University as the relocation to Gilmorehill from the city centre was in 1870. It presents us with a unique opportunity to provide modern, fit for purpose facilities that are in keeping with Glasgow's status as a world leading research-intensive university. As ever, the student experience will be at the heart of all that we do, and the campus redevelopment will allow us to focus on what our students want and need, as well as providing an environment that enables our staff to flourish and is open and accessible to the wider community.

*Professor Anton Muscatelli
Principal and Vice Chancellor
University of Glasgow*



Stimulus: University of Glasgow - Master Plan

I am convinced that this is as significant a development in the ongoing story of the University as the relocation to Gilmorehill from the city centre was in 1870. It presents us with a unique opportunity to provide modern, fit for purpose facilities that are in keeping with Glasgow's status as a world leading research-intensive university. As ever, the student experience will be at the heart of all that we do, and the campus redevelopment will allow us to focus on what our students want and need, as well as providing an environment that enables our staff to flourish and is open and accessible to the wider community.

*Professor Anton Muscatelli
Principal and Vice Chancellor
University of Glasgow*

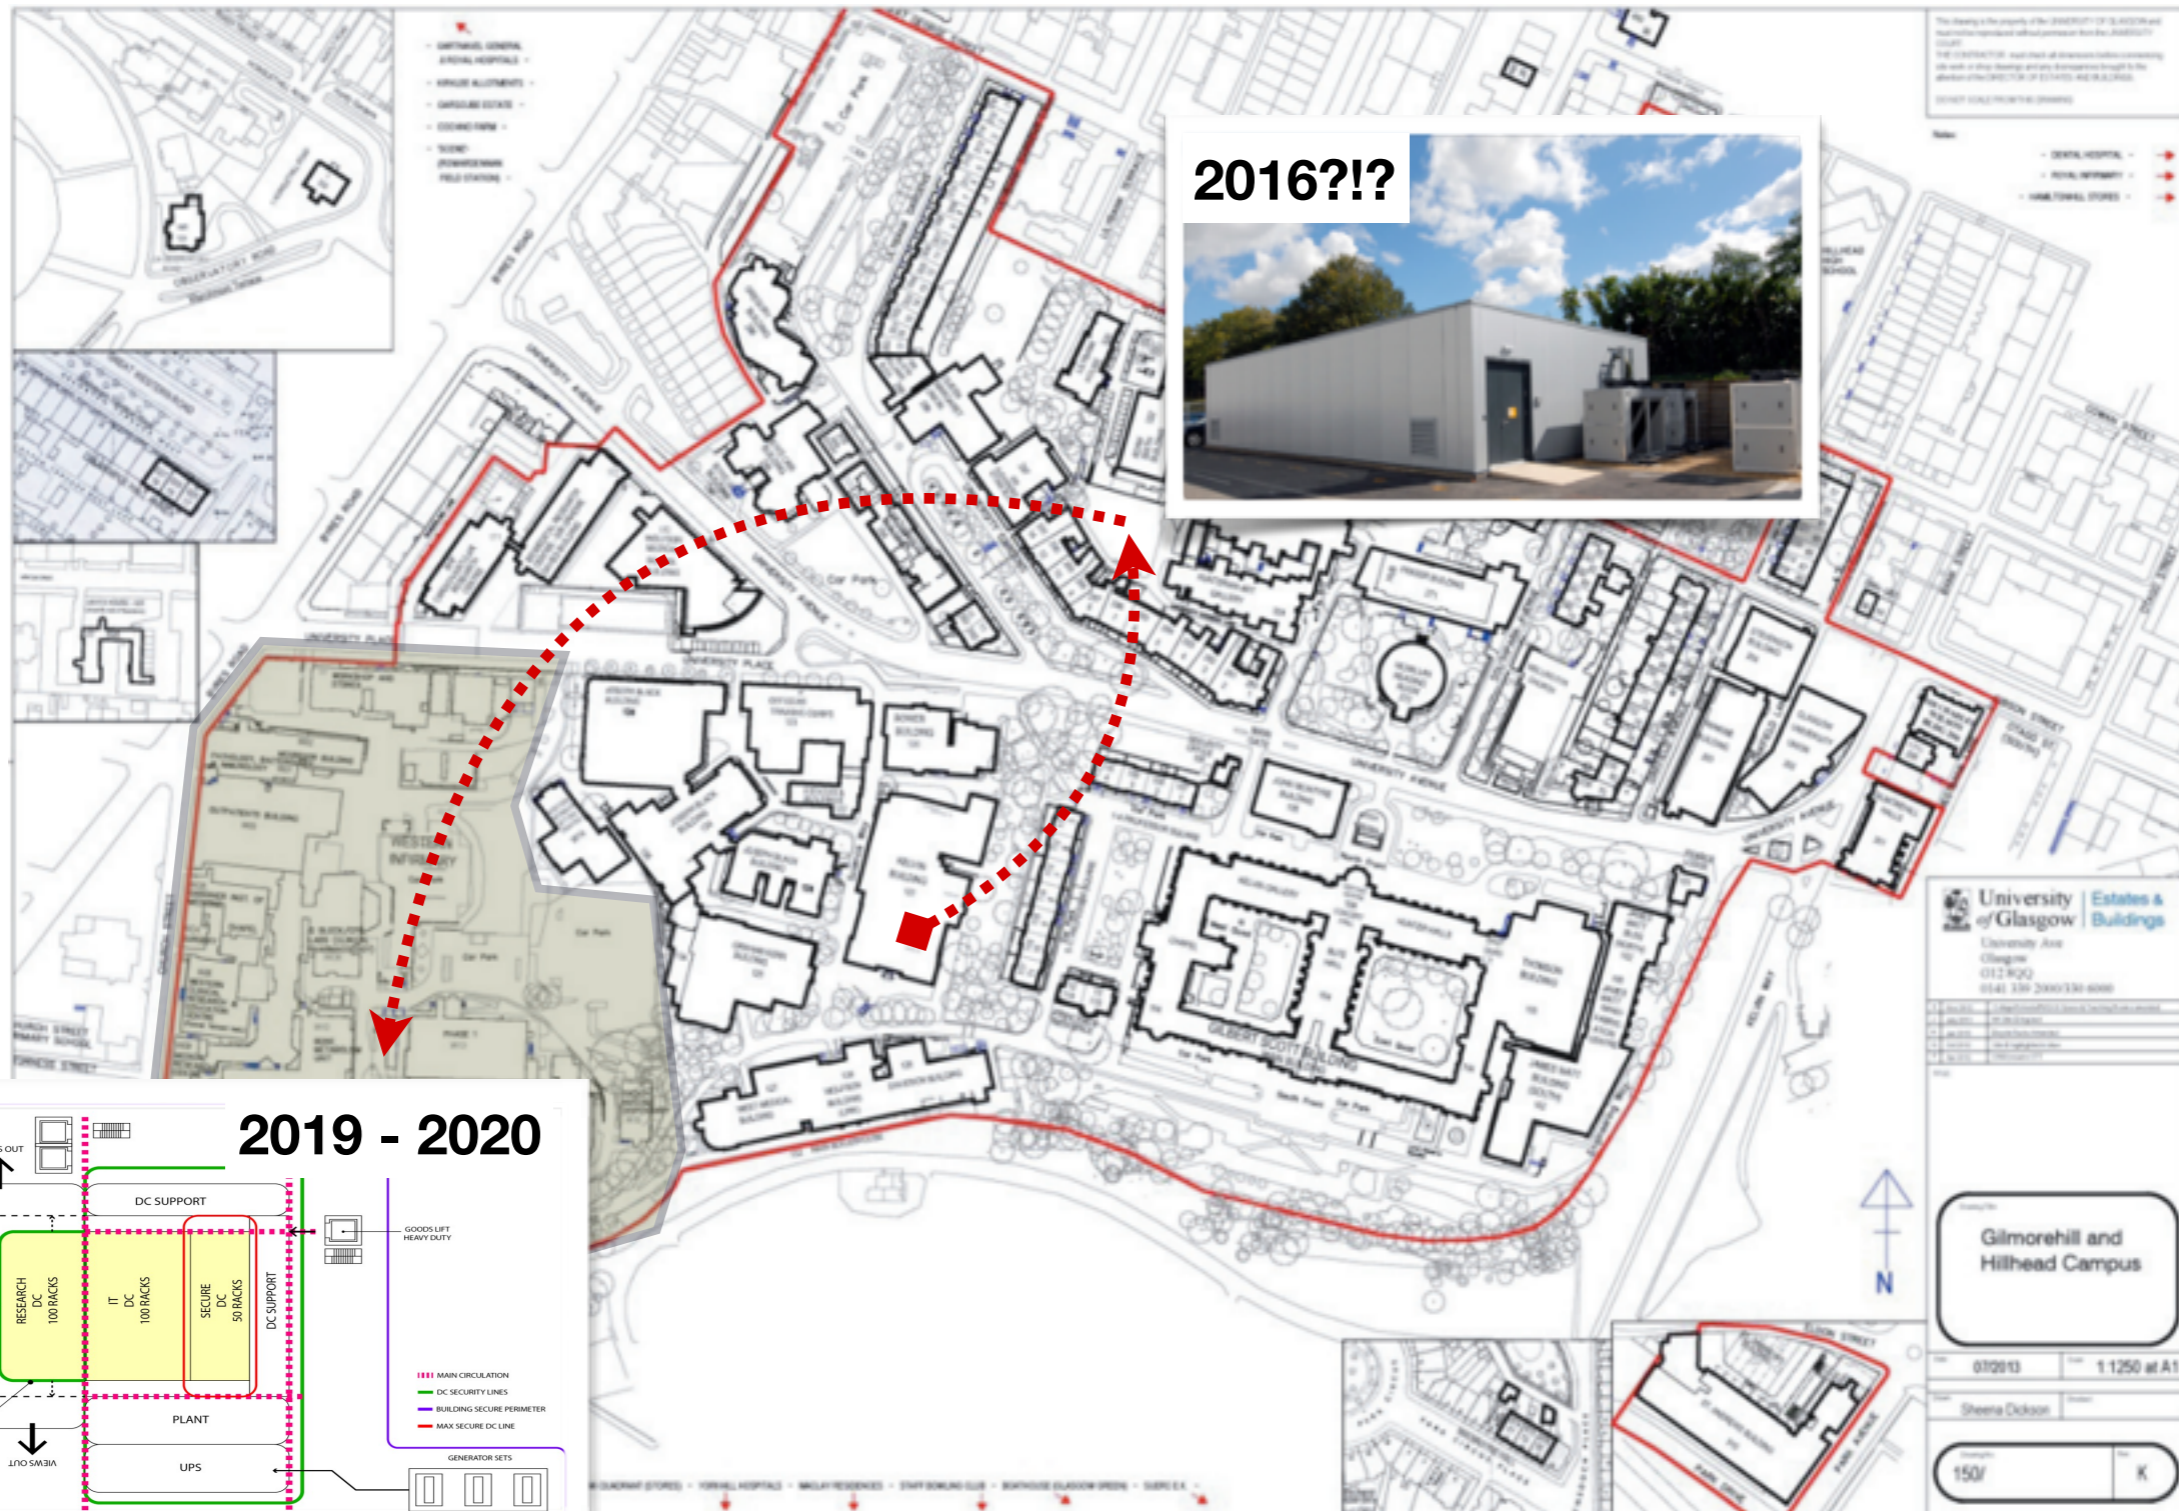




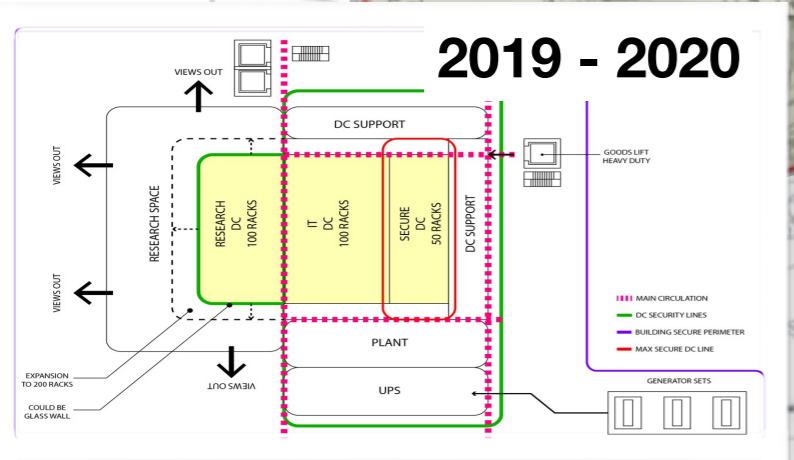
Mutation: Seed of a Shared Cluster


- As part of the Campus redevelopment Glasgow is planning to construct a “state-of-the-art” Data Centre.
- For ScotGrid it’s slightly more complicated as we need to move to an intermediate solution due to renovations to our current location.
- It’s hoped that we can leverage our 10% “other” share (as well as additional match funding) to become the seed of a Shared University Cluster.
 - Importantly, this gives us the opportunity to shape its requirements and guide its development.
 - The ultimate goal is to leverage local funding to increase overall resources available to the WLCG.
 - e.g. Lancaster, ECDF, others...
 - ScotGrid Glasgow is already working with other communities within the University (Arts, UBDC, NHS), potentially creating new user communities for the Grid.

ScotGrid Glasgow - Master Plan



2019 - 2020





Cloud

Tier-2.5

Tier-3

Tier-2 (pt1)


Crossover: Tier3 shared resources

- Previously Tier 2.5, allowing local ATLAS users to use Grid resources and access local storage systems:
 - Many issues encountered with NFS stalls/timeouts & local AFS requirements.
 - Planning to resurrect with a more “up-to-date” approach (xrootd interfaces, preferential network paths).
- Tier 3 resources are often unused, vast majority only required in earnest during conference season.
 - Plan on merging resources, simplifying management and allow local research groups to “burst” into the grid resources, and grid jobs to “burst” into Tier-3 resources (cf. RAL cloud strategy).
 - Also take advantage of shared staffing and development work.
 - In particular looking to move services to local cloud development which will likely be shared with the Grid (more later).

Technological Evolution

MARGINAL GAINS

HOW THE PROFESSIONALS MAKE SMALL CHANGES TO IMPROVE THEIR PERFORMANCE



Inheritance: Where are we now?

- What makes it difficult (at Glasgow) to quickly support new VO's?
- A number of tools and systems still in place that should be phased out:
 - cfengine/puppet duplication
 - ypf/cobbler duplication
 - yaim/puppet duplication
 - WMS/planet/etc...
- Initial goal to declutter and removing unused service (e.g. WMS)
- Centralise services, one source of truth:
 - FreeIPA (cf. Argus)
 - Ansible/Git
- New lighter weight services:
 - S3
 - Submission via arctsub/arctproxy/arctstat

Whenever you make a change to your codebase, there's always going to be a risk that you're about to break something.

No one likes downtime, no one likes cranky users, and no one enjoys angry managers. So the act of deploying new code to production tends to be a pretty stressful process.

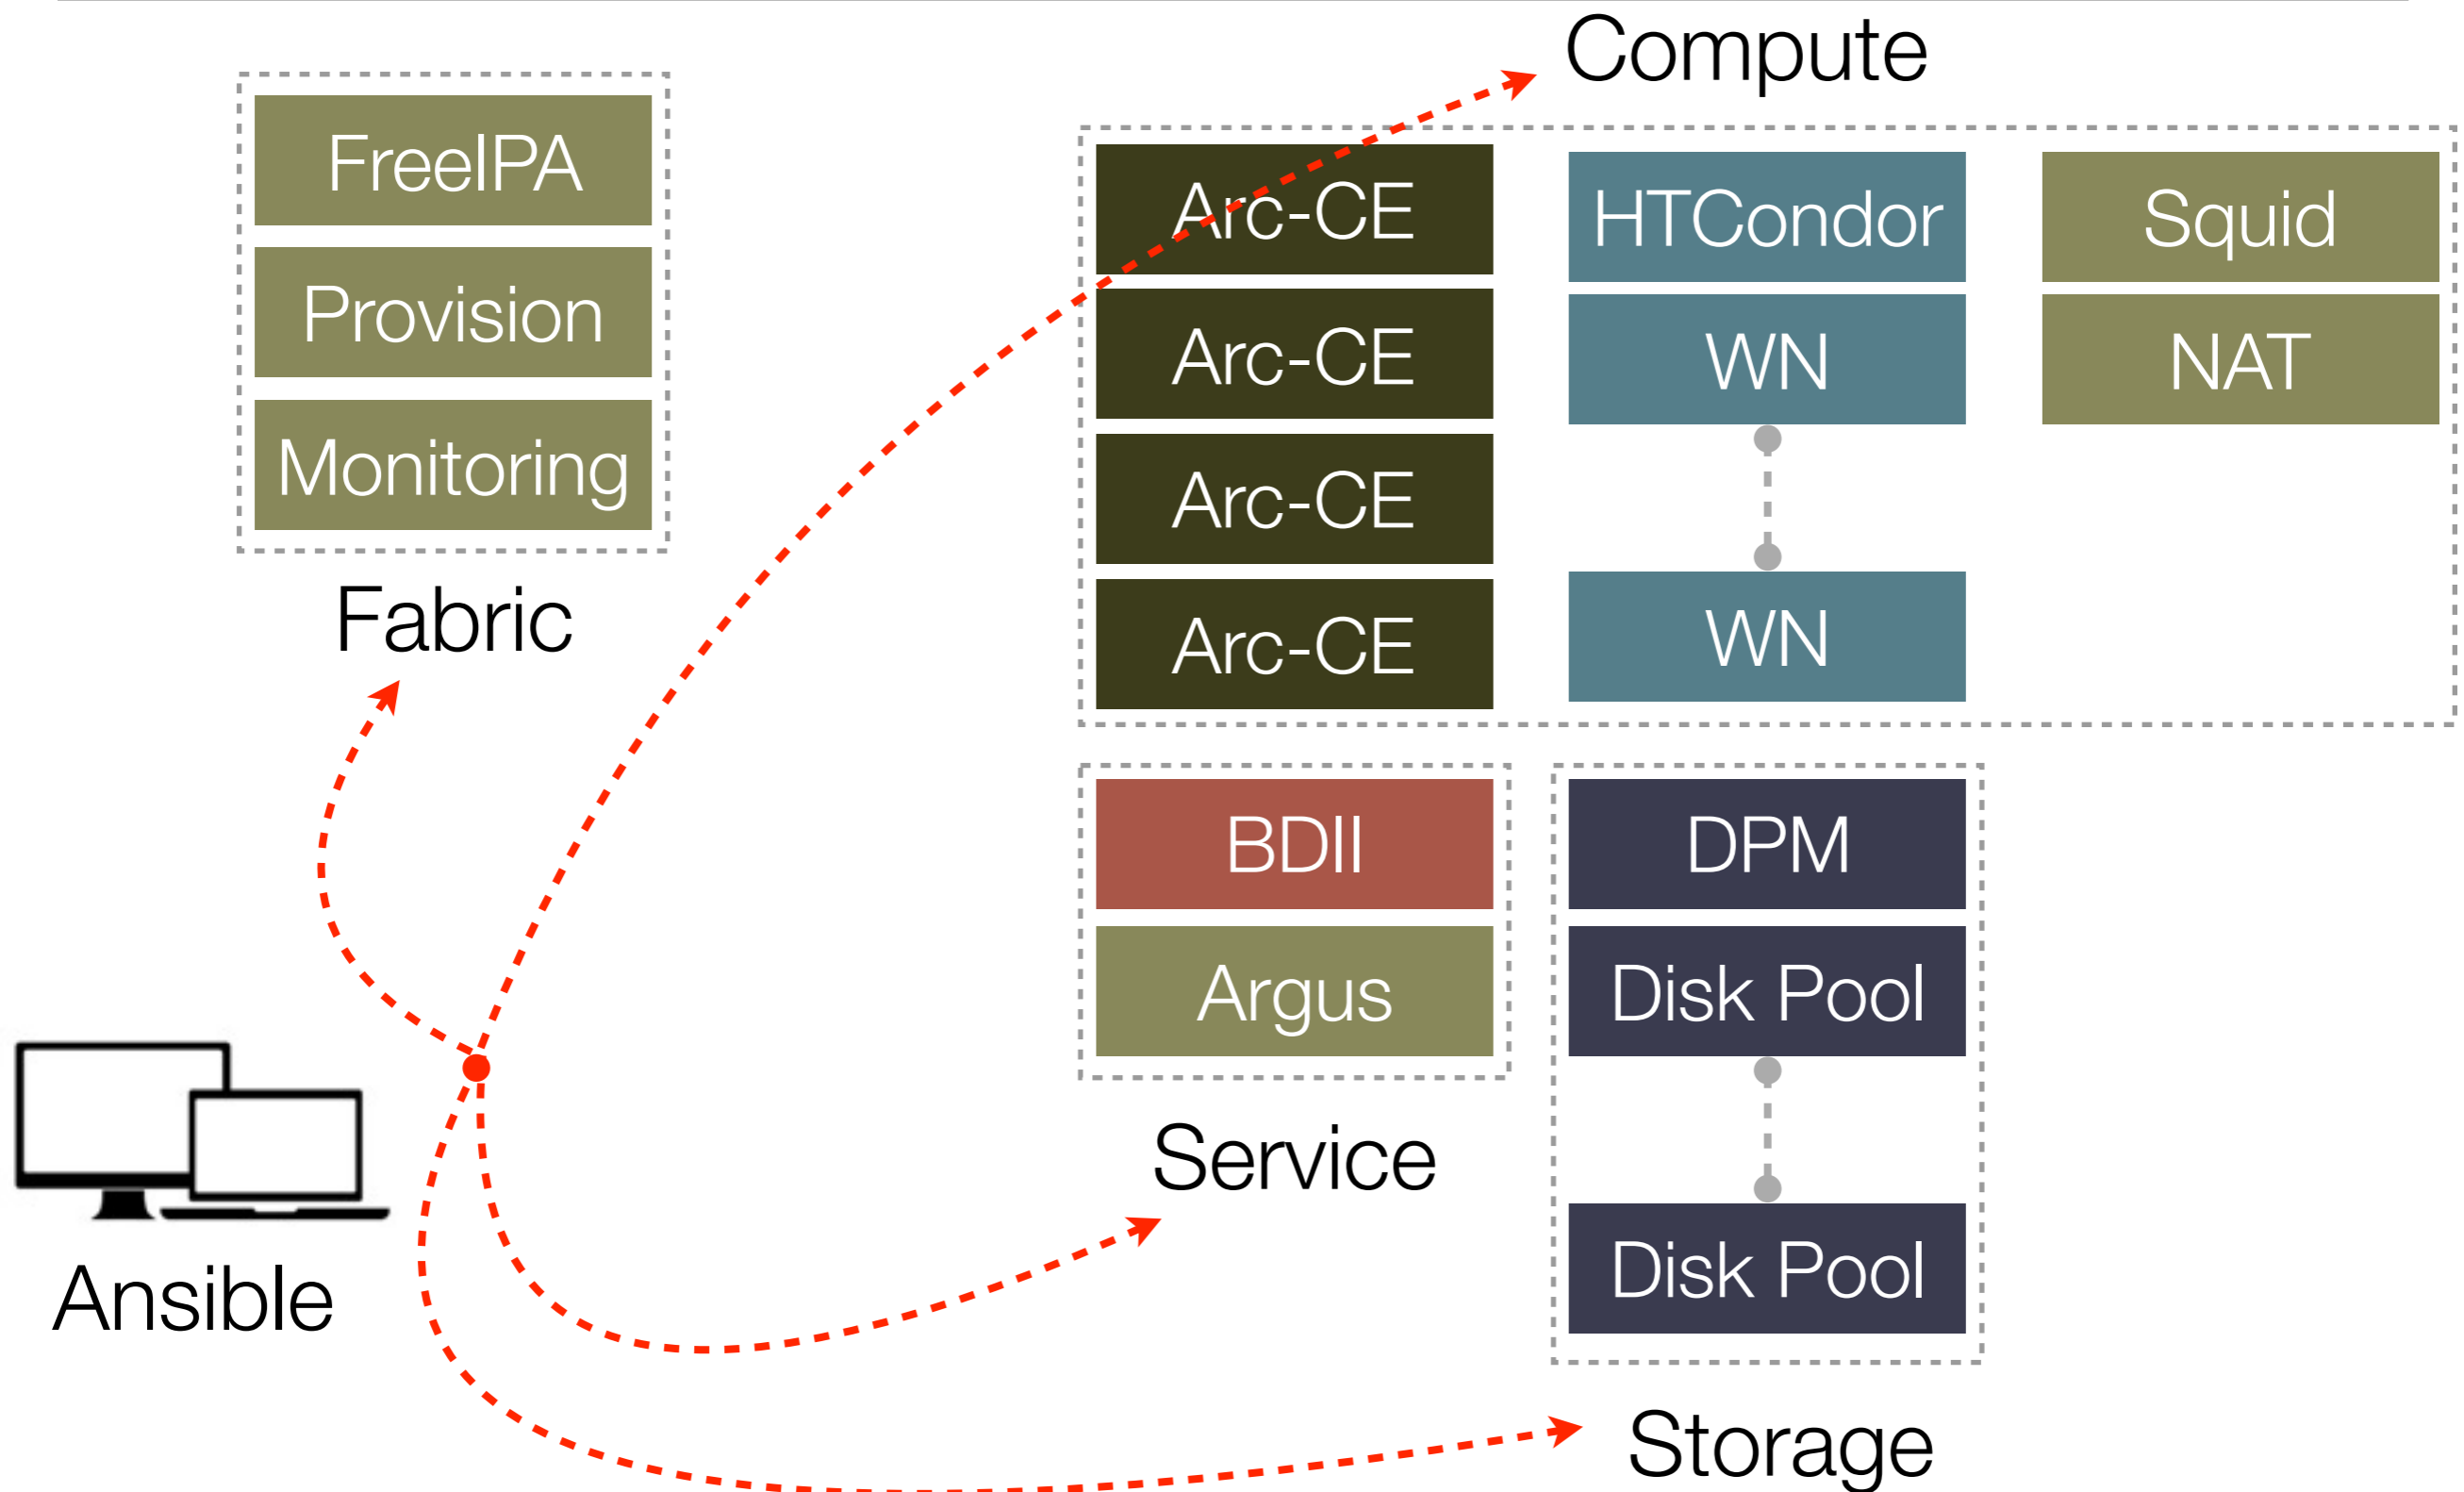
It doesn't have to be as stressful, though. There's one phrase I'm going to be reiterating over and over throughout this whole piece:

“**Your deploys should be as boring, straightforward, and stress-free as possible.**”

Deploying major new features to production should be as easy as starting a flamewar on Hacker News about spaces versus tabs. They should be easy for new employees to understand, they should be defensive towards errors, and they should be well-tested far before the first end-user ever sees a line of new code.

<https://zachholman.com/posts/deploying-software>

A simplified Glasgow



Selection: Ansible

- Evaluated Chef/Puppet/Salt/Cfengine/Quattor/Ansible/Fabric/pdsh/etc...
- Pros:
 - Lightweight push model.
 - No single point of failure.
 - No agents or CA's required. SSH keys only.
 - Allows a "Task" based approach rather than a "State" based approach.
 - Closer to how we work, doesn't feel like your fighting the CM Tool.
- Issues:
 - Initial scaling tests gave worrying performance.
 - Old versions of ssh (CentOS6) causes Ansible to use paramiko (very slow).
 - Enforce ssh and turn off ControlMaster and ControlPersist, performance similar to **pdsh**
 - Should not be a problem with CentOS 7

```
svr000:~/gdr/Ansible# cat update-anchors.yml
- hosts: demo
  tasks:
  - name: Install/Upgrade lcg-CA packages.
    yum: name=lcg-CA state=latest
svr000:~/gdr/Ansible# ansible-playbook -i inventory update-anchors.yml

PLAY [demo] *****

GATHERING FACTS *****
ok: [node001]

TASK: [Install/Upgrade lcg-CA packages.] *****
ok: [node001]

PLAY RECAP *****
node001                : ok=2    changed=0    unreachable=0    failed=0

svr000:~/gdr/Ansible# _
```

Selection: FreeIPA


- Upstream project for the Redhat Identity Management Platform
- Centralised account and group management, via LDAPv3.
- Uses know projects such as 389 Directory services, Kerberos, SSSD etc...
- Web and CLI interfaces for ease of use and scripting.
- Allow creation and deployment of users/groups in a much simpler fashion than at present at Glasgow (no more static /etc/passwd files)
- Successfully deployed at Durham for local and pool accounts.
- Test system at Glasgow, simple to deploy and setup, waiting for Hardware to bring online a production instance.

A screenshot of the FreeIPA web interface. The top navigation bar includes 'IDENTITIES', 'POLICIES', and 'IPA CONFIGURATIONS'. Below this, there are tabs for 'Users', 'Groups', 'Hosts', 'Host Groups', 'Services', and 'Netgroups'. The main content area is titled 'Managing User: Pat D. Bunny' and includes links for 'Details', 'Enrollment in Groups', 'Enrollment in Netgroups', and 'Roles'. A table lists groups with columns for Group, GID, and Description. The table contains seven rows of data. Below the table, there are instructions for viewing current enrollment, enrolling in new groups, and withdrawing from groups.

Group	GID	Description
Group10	2344545	Group description1
Group132	8659751	Group description2
Group5548	4213659	Group description3
Group7895	8546958	Group description4
Group5432	8742154	Group description5
Group754	7542154	Group description6
Group7845	8456923	Group description7

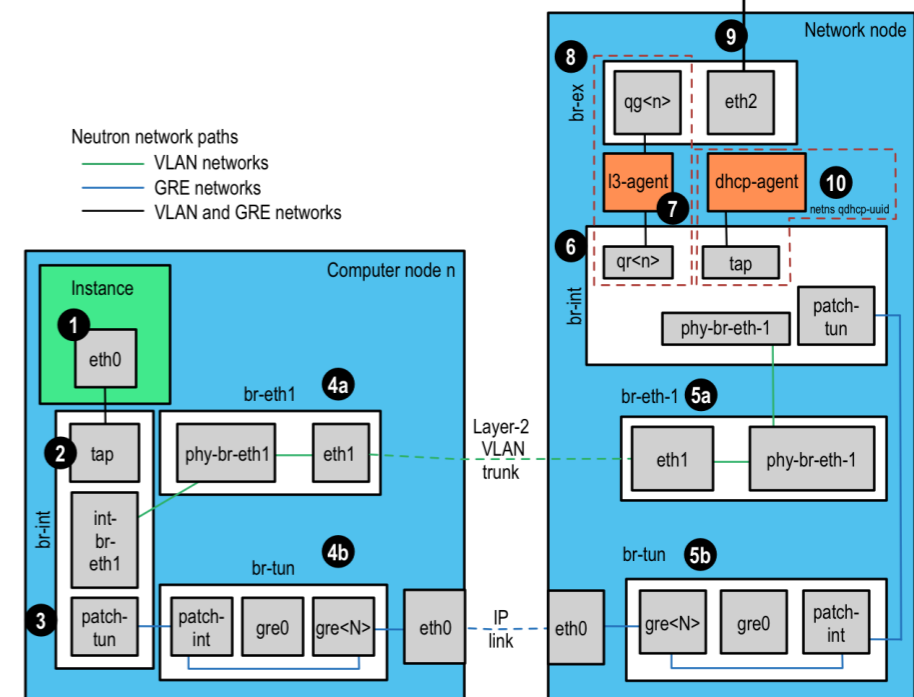

Selection: OpenStack

- Proof of concept:
 - Deployed Juno/Kilo/Liberty
 - OpenStack Liberty
 - 12TB of Block Store
 - 128 CPU cores for compute
 - Neutron Networking
 - Looking at Swift



Selection: OpenStack

- Issues:
 - Difficulty in upgrading (OS and OpenStack) and following security updates.
 - Setup was brittle and easy to break (could be our inexperience).
 - Opaque configuration utilities, lots of Puppet.
 - Overly complex network configuration for **our** purposes.
- Way forward:
 - Provider Networking (Flat or VLAN) to simplify deployment.
 - Using Ansible, writing own deployment scripts to understand components and interactions (see Gordon).



Selection: Storage - S3

- How do we replace DPM (if we need to/future proofing)?
- Can we replace DPM?
- Can we do something in parallel, Atlas Event Service will use S3 for small files.
- Can we use S3 for Local VO's?
 - Building out services, easy to use Platform, etc.
- In planning, initial test systems to be set up with older storage, **not** a production system.
- Exploring technologies that have a compliant interface:
 - Ceph/RADOS
 - OpenStack Swift
 - Riak CS/HAProxy

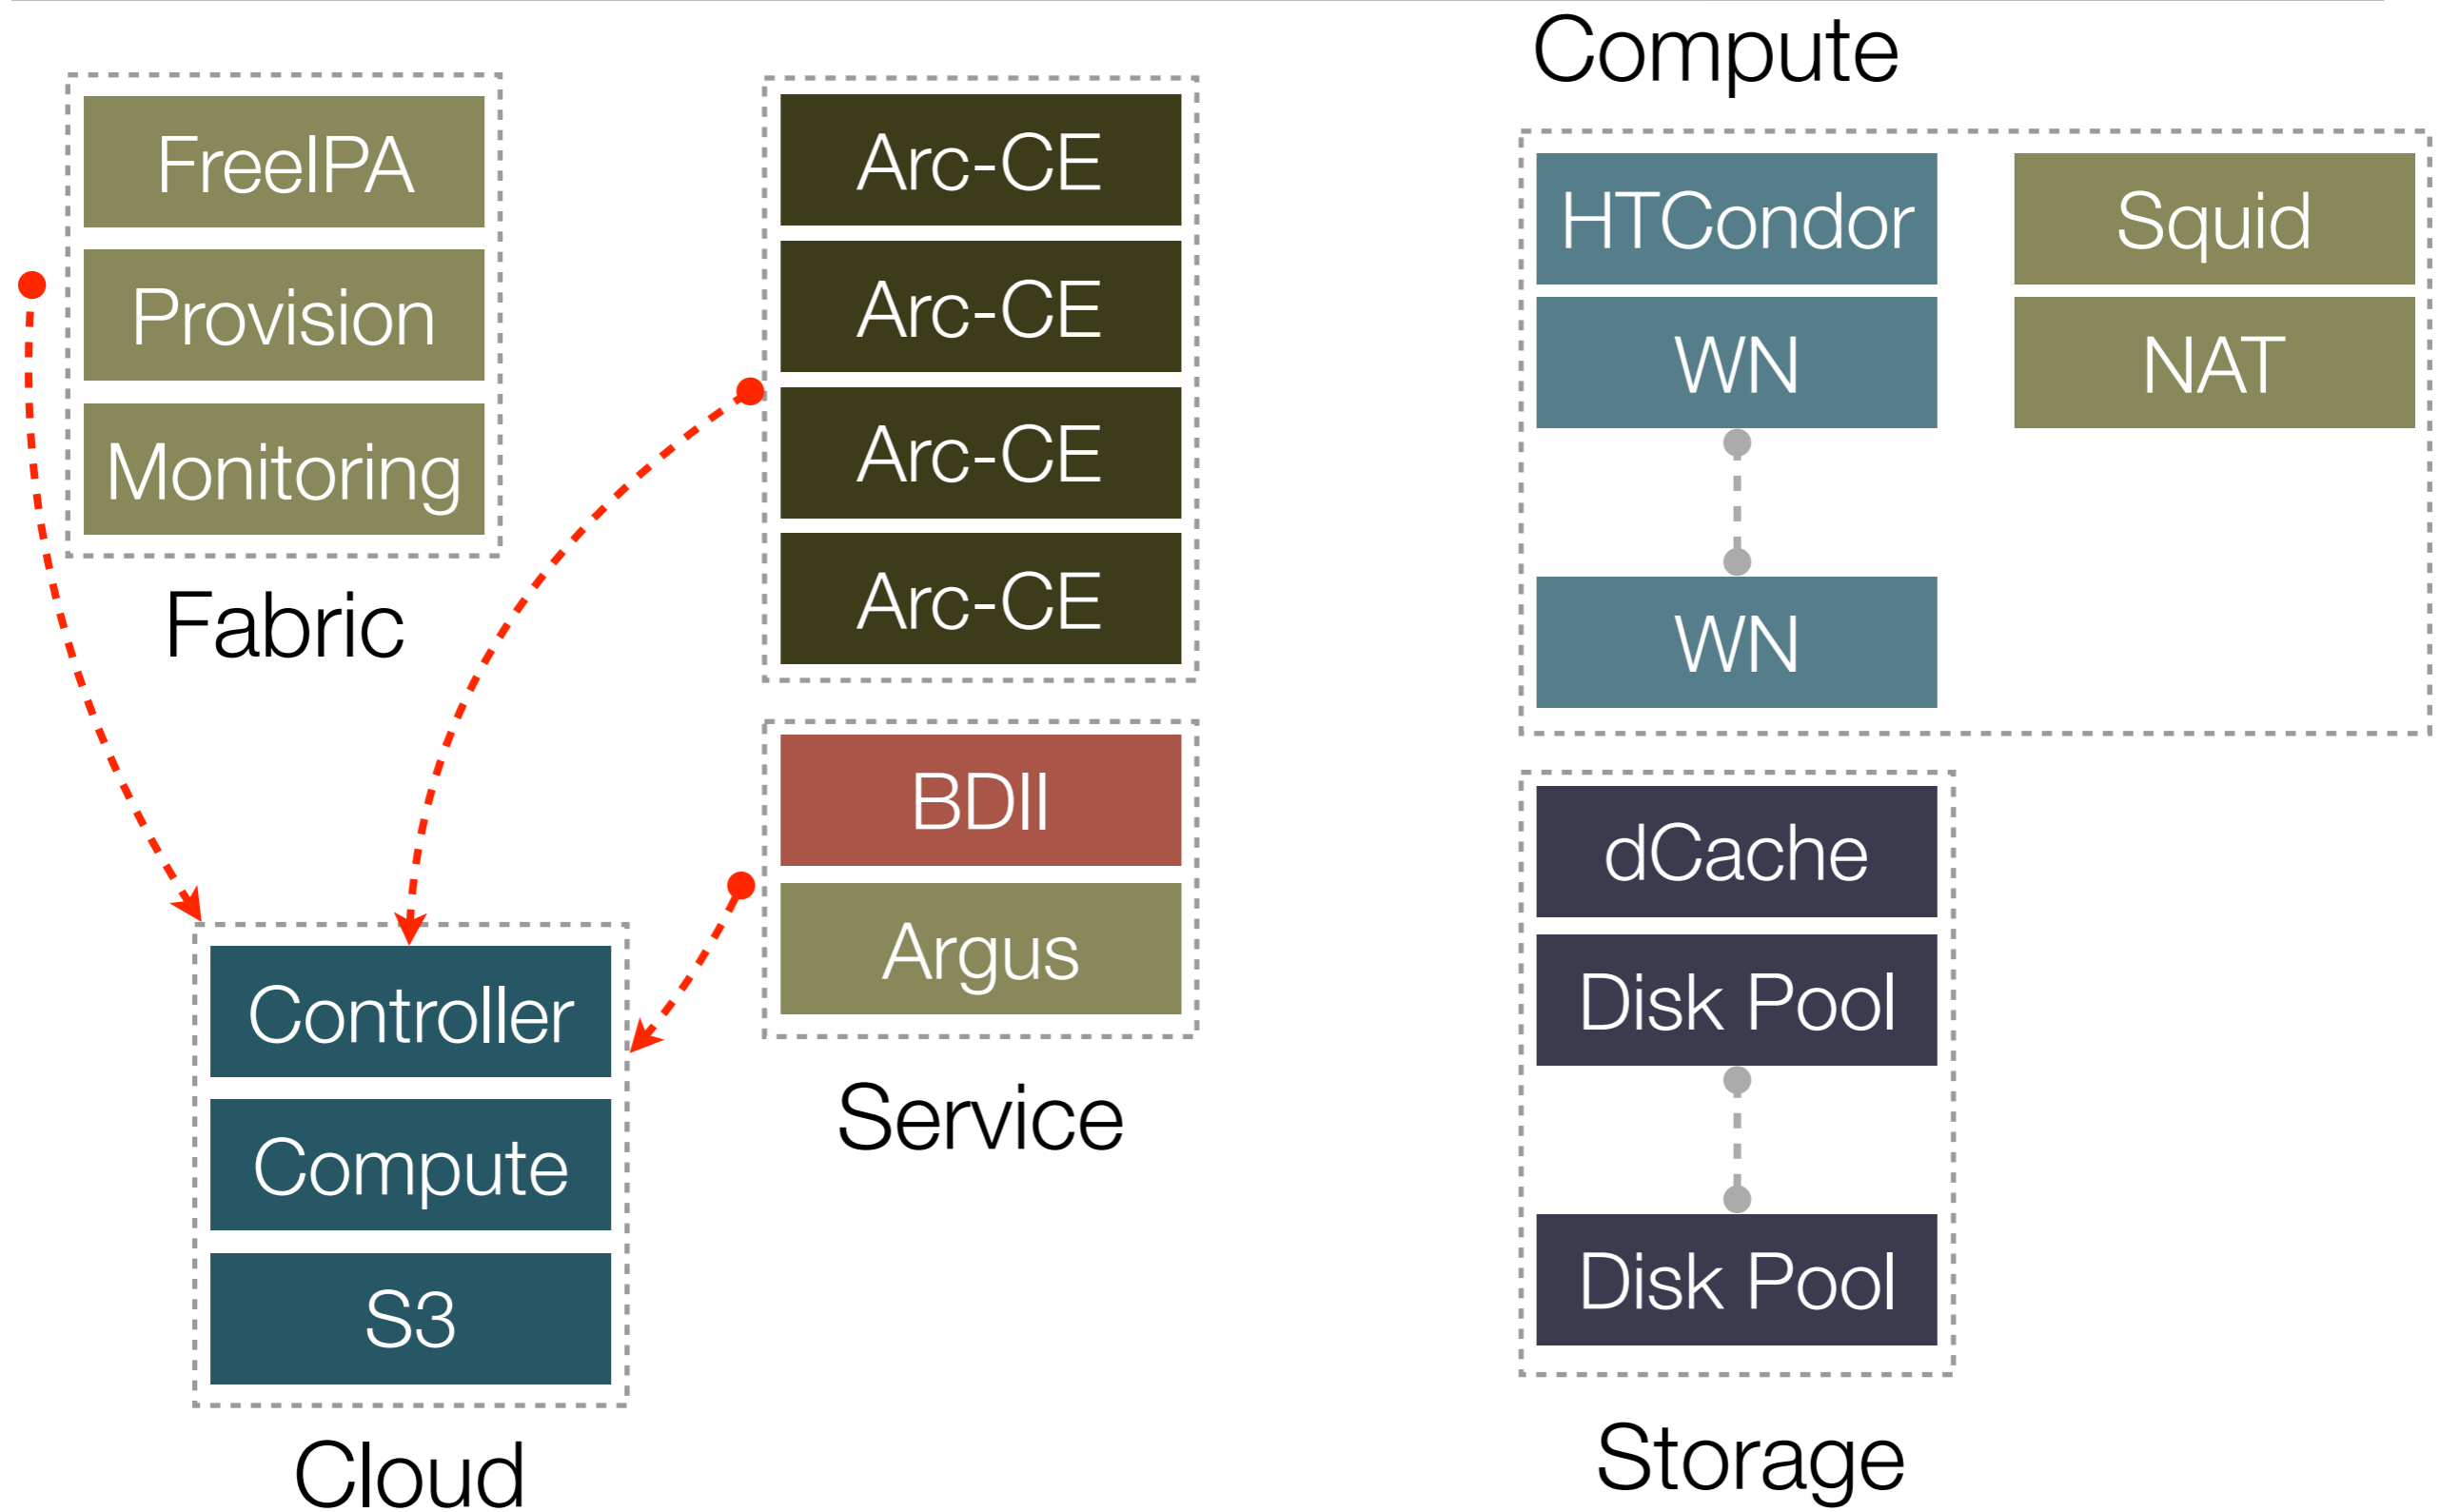


Selection: Storage - dCache

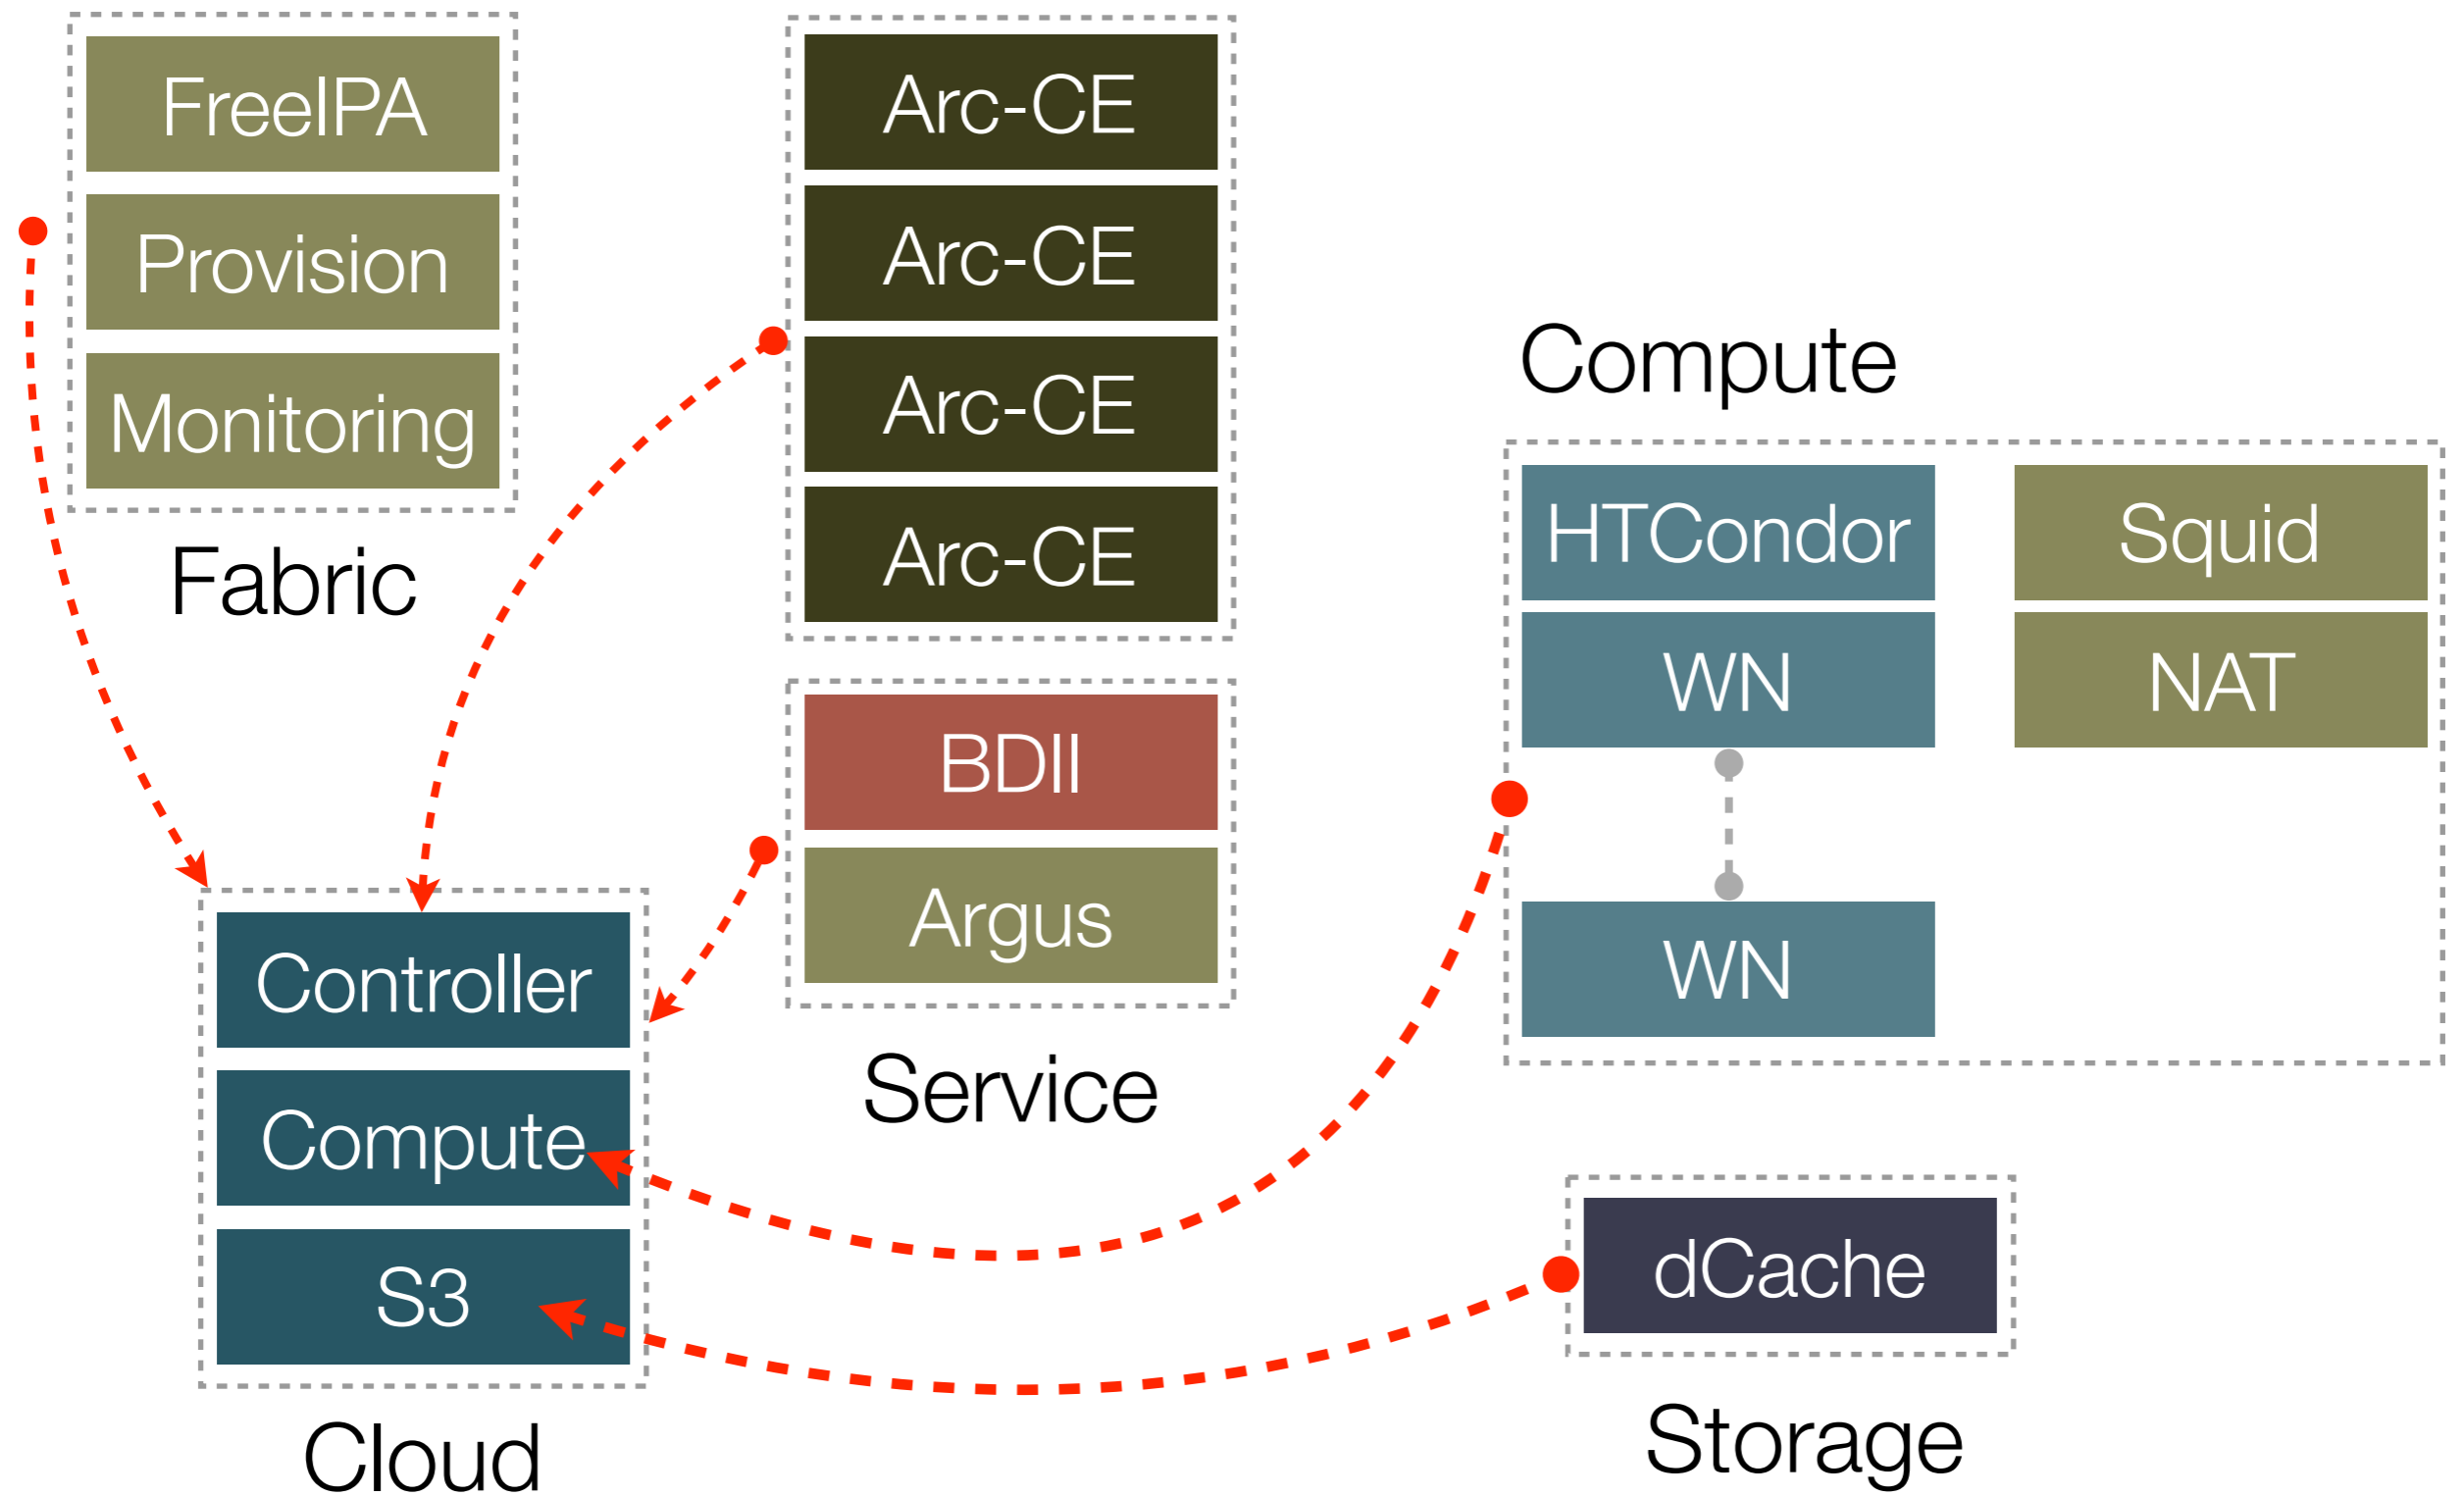
- If S3 isn't a solution then what is?
- Existing grid technologies:
 - Gaster/EOS/dCache
- Experience in the UK using dCache at T2's.
- Experience across the WLCG using dCache at T2's.
- Support for non-Grid technologies:
 - Ceph/S3/Http/Oauth2
- How do we migrate?
- Can we move a DPM site to dCache?
- Initial work by Sam (and Paul Millar) indicates this is "not impossible".

dCache.org 

A Future Glasgow?



A Future Glasgow? - Flexibility



An aside, what can be done with 0.25 FTE (with investment)

by SITE and VO																
SITE	None	atlas	cms	enmr.eu	gridpp	llc	lhcb	lsst	mice	na62.vo.gridpp.ac.uk	nordugrid.org	ops	pheno	vo.scotgrid.ac.uk	total	%
UKI-SCOTGRID-DURHAM	0	19,839,713	0	0	7,729	225,445	1,441,236	0	0	0	0	554	16,375,898	11	39,892,087	40.49%

Q116 - 39,892,087 HEPSPPEC/hrs

Q115 - 9,568,140 HEPSPPEC/hrs

~19x Atlas Improvement

by SITE and VO																		
SITE	None	atlas	biomed	cernschool.org	cms	dteam	enmr.eu	epic.vo.gridpp.ac.uk	gridpp	llc	lhcb	mice	na62.vo.gridpp.ac.uk	ops	pheno	vo.scotgrid.ac.uk	total	%
UKI-SCOTGRID-DURHAM	0	1,104,388	0	0	0	50	0	0	0	313,665	1,238,217	0	0	0	6,911,820	0	9,568,140	5.55%

Well done Oliver and the Team!!!

Conclusion

MARGINAL GAINS

HOW THE PROFESSIONALS MAKE SMALL CHANGES TO IMPROVE THEIR PERFORMANCE

