

Data Management: Reducing the cost of storage.

Summary of Infrastructure perspective (Europe)

Alastair Dewhurst

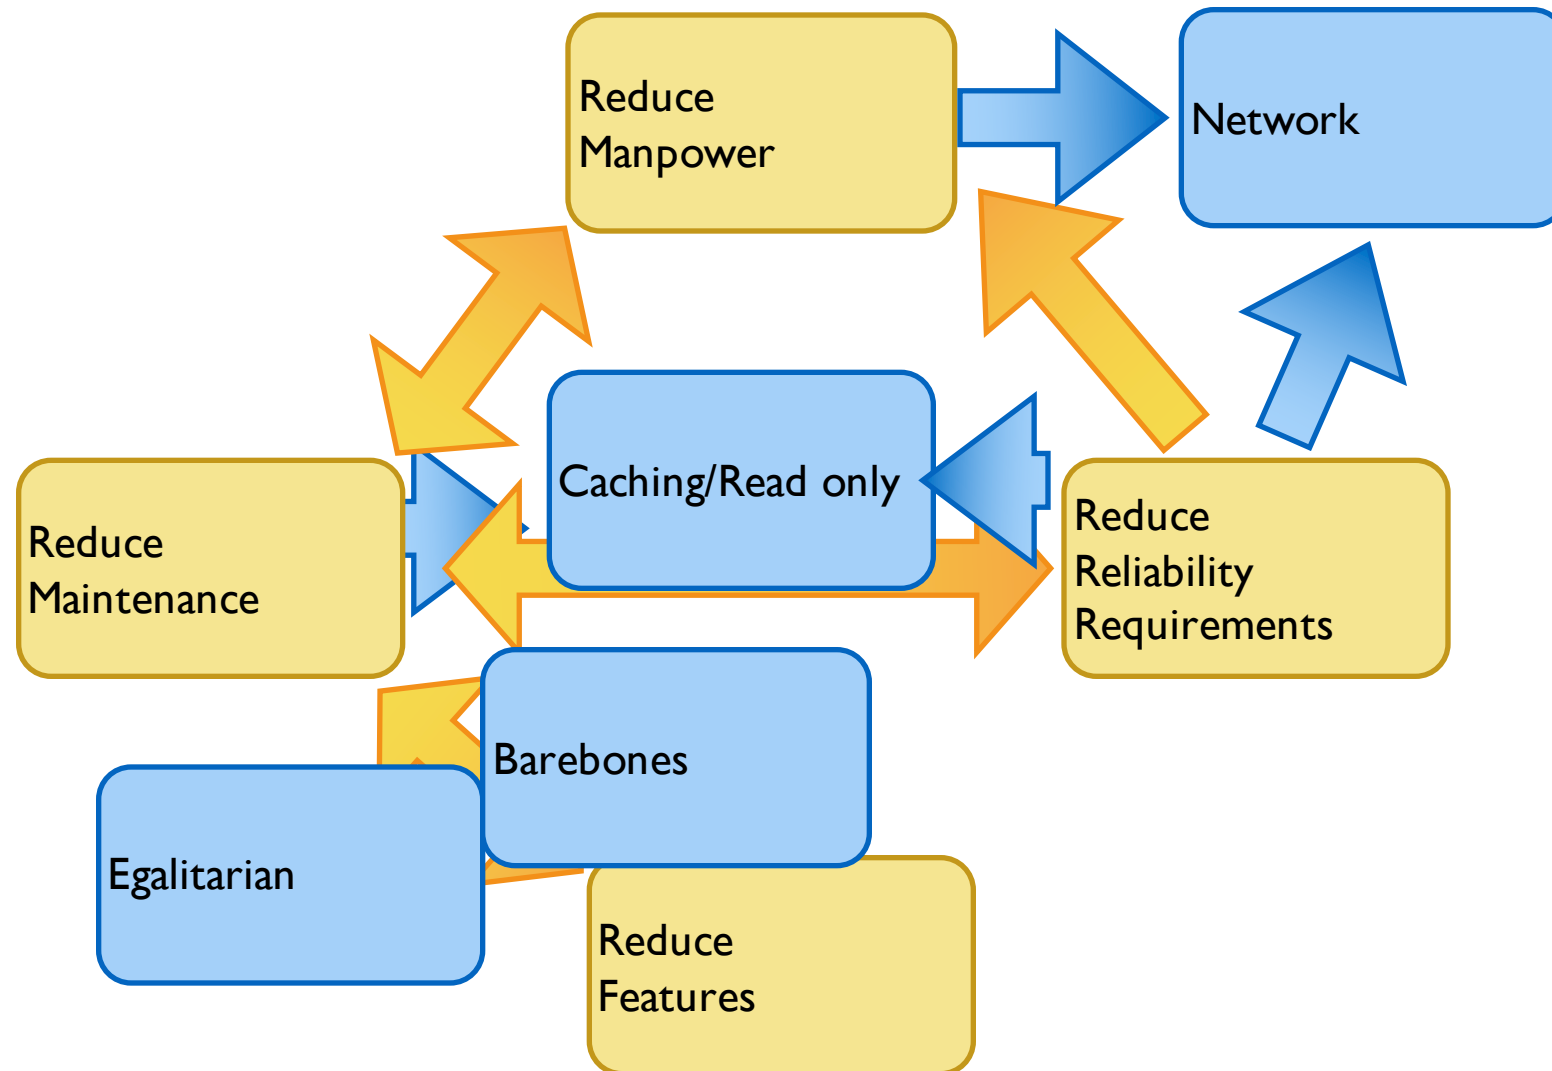


Context

- European sites often support many VOs.
- Storage is the most time consuming service for site admins.
 - This is our expertise compared to other sciences (HPC).
 - This is a way to attract other sciences and extra-income.
- Sites are making long term decisions regard Run 3. Need clear understanding on evolution of VO computing models:
 - Is the evolution because of a physics requirements? (e.g. Higher pile-up)
 - Is the evolution to save VO costs? (Fine as long as it doesn't just transfer it to sites)
 - Is the evolution to help Sites reduce cost? (Thank you - as long as it will help!)



How can we reduce costs?



Storage-less / Caching Tier 2s

- Sites see caching model as a good solution:
 - Allows a transition between normal site and storage-less site.
 - Reduces reliance on network compared to completely storage-less.
 - Allows hardware to be used longer and lowers maintenance requirements.
- How do storage-less sites affect ratio between CPU/disk?
- How do Federations fit in?
 - Significant effort from sites for no obvious benefit.
- What does this imply for network evolution?



RAL new storage experience

- RAL has been building a new Ceph based storage.
- Setting up S3 / Swift API was trivial.
- The VO 'requirements' made it hard:
 - Required VO expertise to understand what VO really need.
 - No solution for ALICE (yet) because we don't have an expert.
- Protocol Zoo implies choice:
 - Have to have GridFTP for WAN transfers.
 - Have to provide XrootD because direct I/O is necessary (fo
 - Want to provide https because easier to 'sell' to new users.



Questions?



Backup



Overview of Major EUDAT Services

Component	Role	Current impl. based on:
B2SAFE	Replication of data across infrastructure	iRODS
B2SHARE	Collaborative sharing of docs/data	Invenio
B2FIND	Metadata/discovery	CKAN
B2DROP	Dropbox-type service	OwnCloud
B2STAGE	Move data between infrastructures	GridFTP
B2ACCESS	Authentication and authorisation	Unity

EUDAT does not write software from scratch, so each service is implemented on top of *something* – however, services could be implemented independently



EUDAT & WLCG - Similarities and Differences

- Similarities
 - Distributed infrastructure, IGTF host certs.
 - Data focus, GridFTP/HTTP protocol
 - EUDAT Community == WLCG VO
 - Replicating data between sites
 - Operations - SIRTFI
- Differences
 - Data only – no user defined compute
 - Very different user communities (e.g. several in each of: arts & humanities, Earth sciences, bioinformatics, libraries)
 - Users generally use portals, multi-LoA authentication



Opportunities?

- Move data between infrastructures
 - Initial targets are EGI and PRACE
 - ... could be WLCG, too?
 - Technically – prob. few community reqs
 - Use of open standards & interoperable implementations...
- Common authentication between infrastructures
 - Move data with GridFTP (B2STAGE, B2SAFE)
 - Internal X.509 certificates available
 - Investigating RCauth with EGI
 - Support Globus?
- Sustainability

