



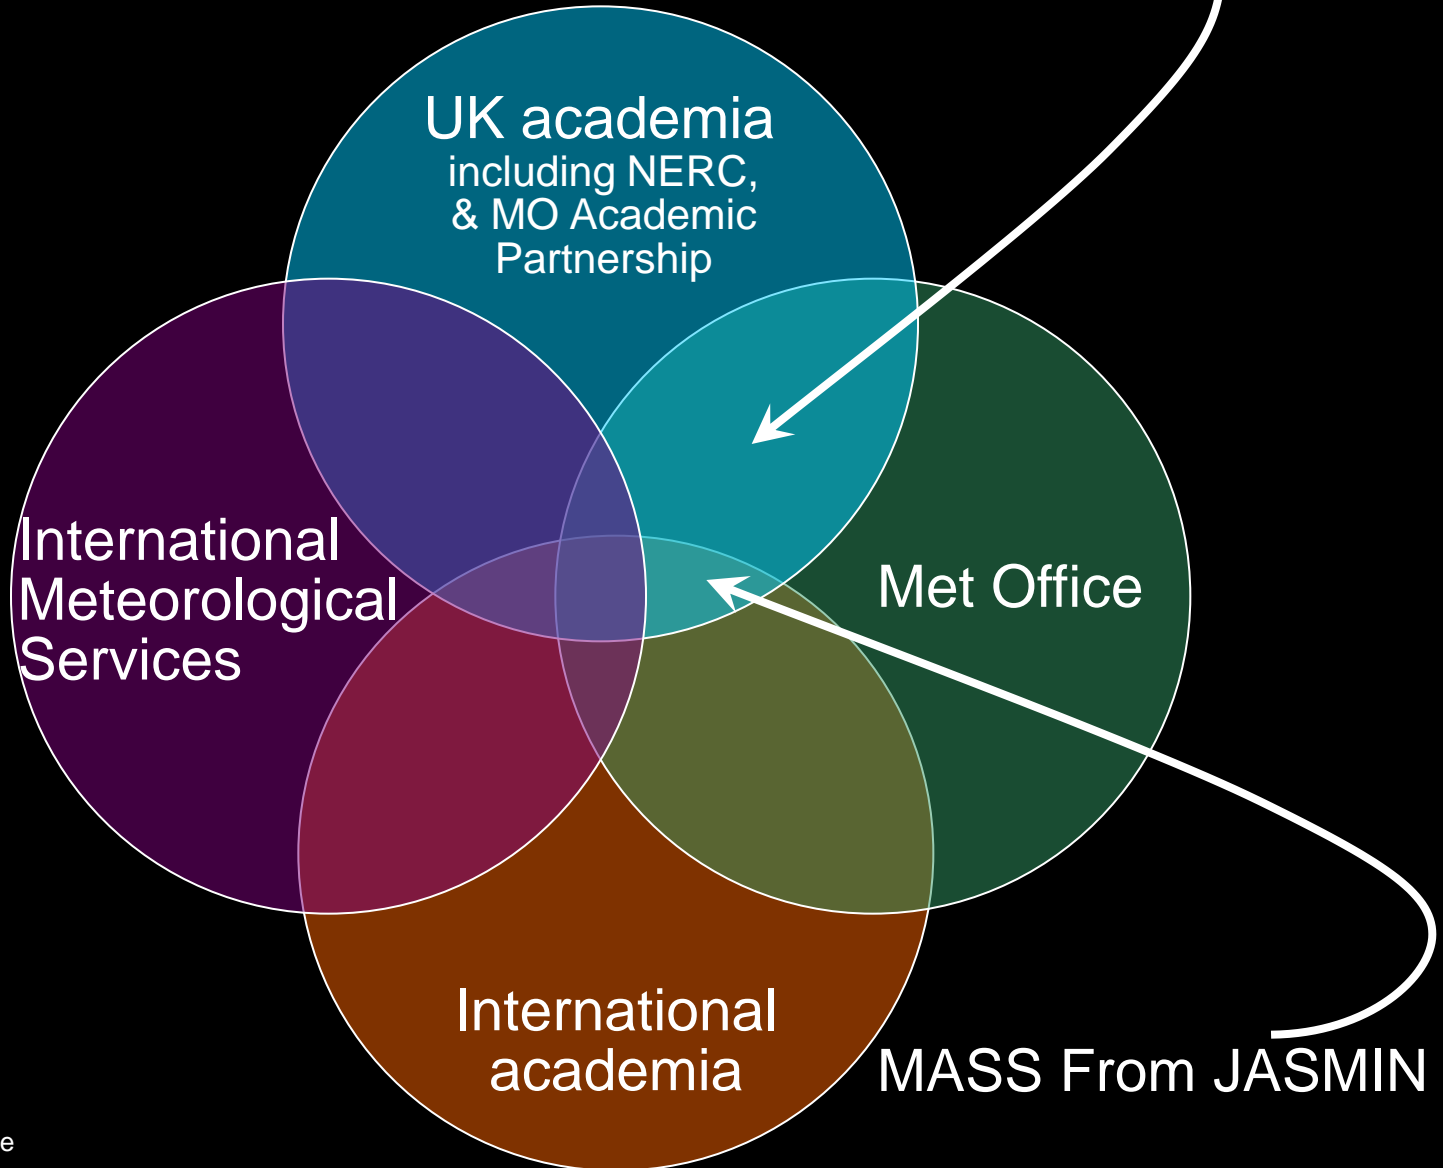
MONSooN, MASS, MOSRS

AJ Watling – MONSooN Technical Lead

Roger Milton – Scientific Collaboration Technical Lead



Collaboration on MONSooN





Context of MONSooN

- Joint Weather and Climate Research Programme (JWCRP) objectives:
 - Enable closer collaboration between Met Office and NERC scientists
 - Develop activities to address gaps in national portfolio of weather and climate research
 - Promote effective pull through of research to improved forecasts
- **MONSooN = Met Office / NERC Superco(o)mputer Nodes**
 - “To provide a shared supercomputing service in an environment within which collaborative projects between NERC and Met Office can be performed. The joint service is being provided to enhance collaboration between NERC and Met Office researchers and improve their ability to pull-through environmental science more rapidly.”*
- **All projects** must be a collaboration between NERC funded and Met Office scientists; applications are assessed for Technical & Scientific ‘Fit’



Why use MONSooN?

- Common access to Unified Model for scientists
 - at Met Office, NERC sites or UK universities
- Sharing of configurations, data and jobs
- Ease of entry for new users
 - UM is already set up and verified on MONSooN
- Access to dedicated UM support
 - from NCAS-CMS and Met Office
- Same HPC Architecture as Operational Met Office HPC



Technologies:

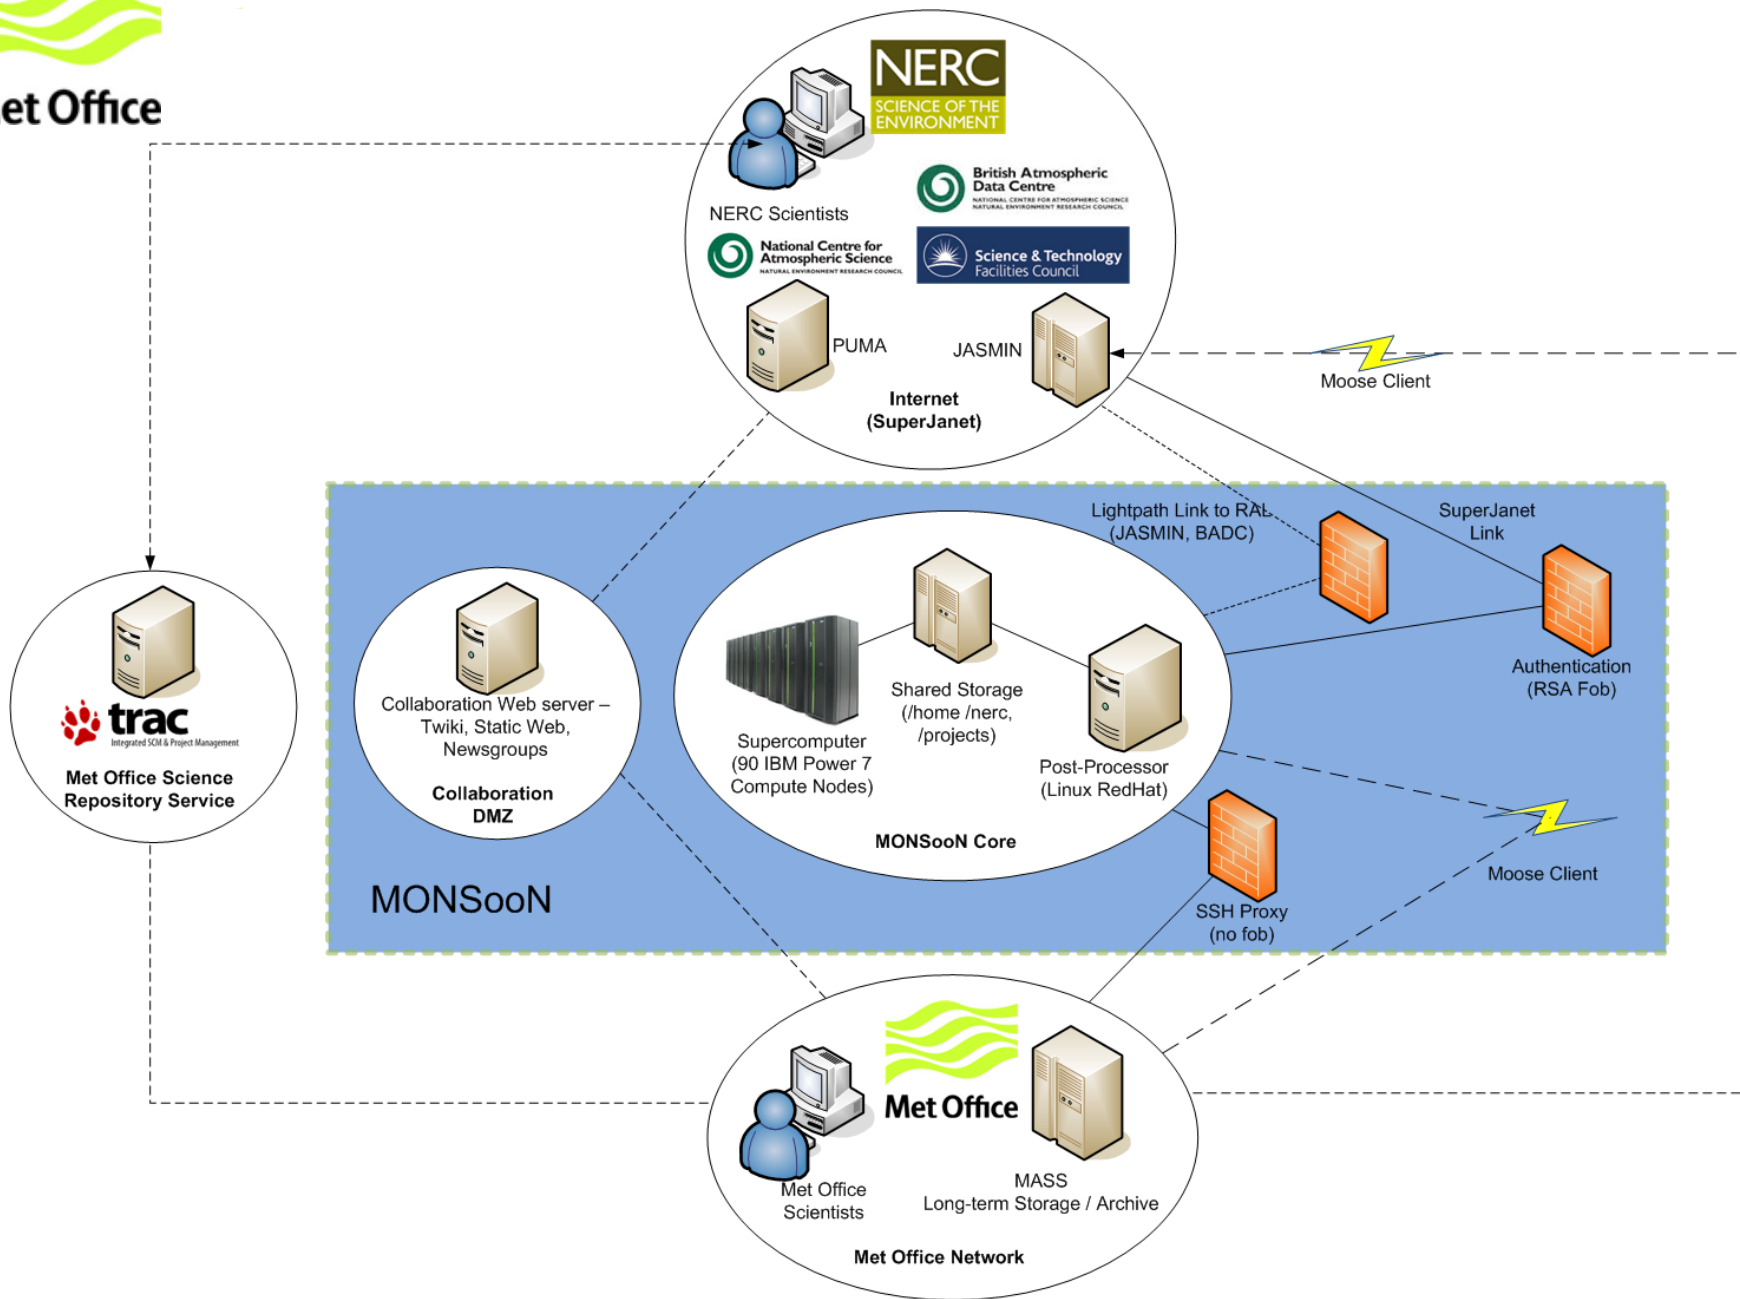
- *Not just a Supercomputer*



- 'Small' Supercomputer – IBM Power 7, migrating to Cray XC-40
- Post-Processor - Linux RedHat, like Met Office's Scientific Desktop
- Shared filesystem: `/home`, `/projects`
- Data transfers - SuperJanet and BADC (each 1Gbps)
- Read/Write Access to Data (e.g. MASS)
- Facilities for code development – e.g. Code repositories, configuration management.
- TWiki for collaborative documentation & discussion



Met Office Collaboration Architecture





Change of Supercomputers



2015:

June/July – Migration to Cray from IBM

(~same capacity as at present)

August – IBMs turned off & removed

November – Cray XCE expanded

2016:

February – Cray XCF expanded

2017:

February – New Cray XCS system



Exeter Science Park



- Third Phase of HPC ~April 2017
- Power limitations of HQ necessitates a purpose-built facility
- Earmarked for increased collaboration
- Configuration / access regime to be decided



And now... Roger



MASS Overview

- Main Met Office resilient tape-based archiving system
- 2 automated tape-libraries, ~50 PB each
- Duplexed copies
- 1 PB cache
- HPSS library-management system
- Bespoke user-interface 'MOOSE'
- Current stats:
 - > 25 PB archived
 - ~40 TB written per day, ~60 TB accessed per day
 - ~50% of retrieves served from cache



MASS upgrade for new HPC

- Will store up to 600PB of primary data by 2020
- Multiple automated tape libraries in 2 IT Halls
 - At least 240 tape drives
 - At least 55,000 tapes
- At least 6.5PB cache by 2017
- 345TB-445TB per day expected to be archived by 2017
- 150TB per day expected to be restored by 2017
- First phase of the upgrade in production by October 2015

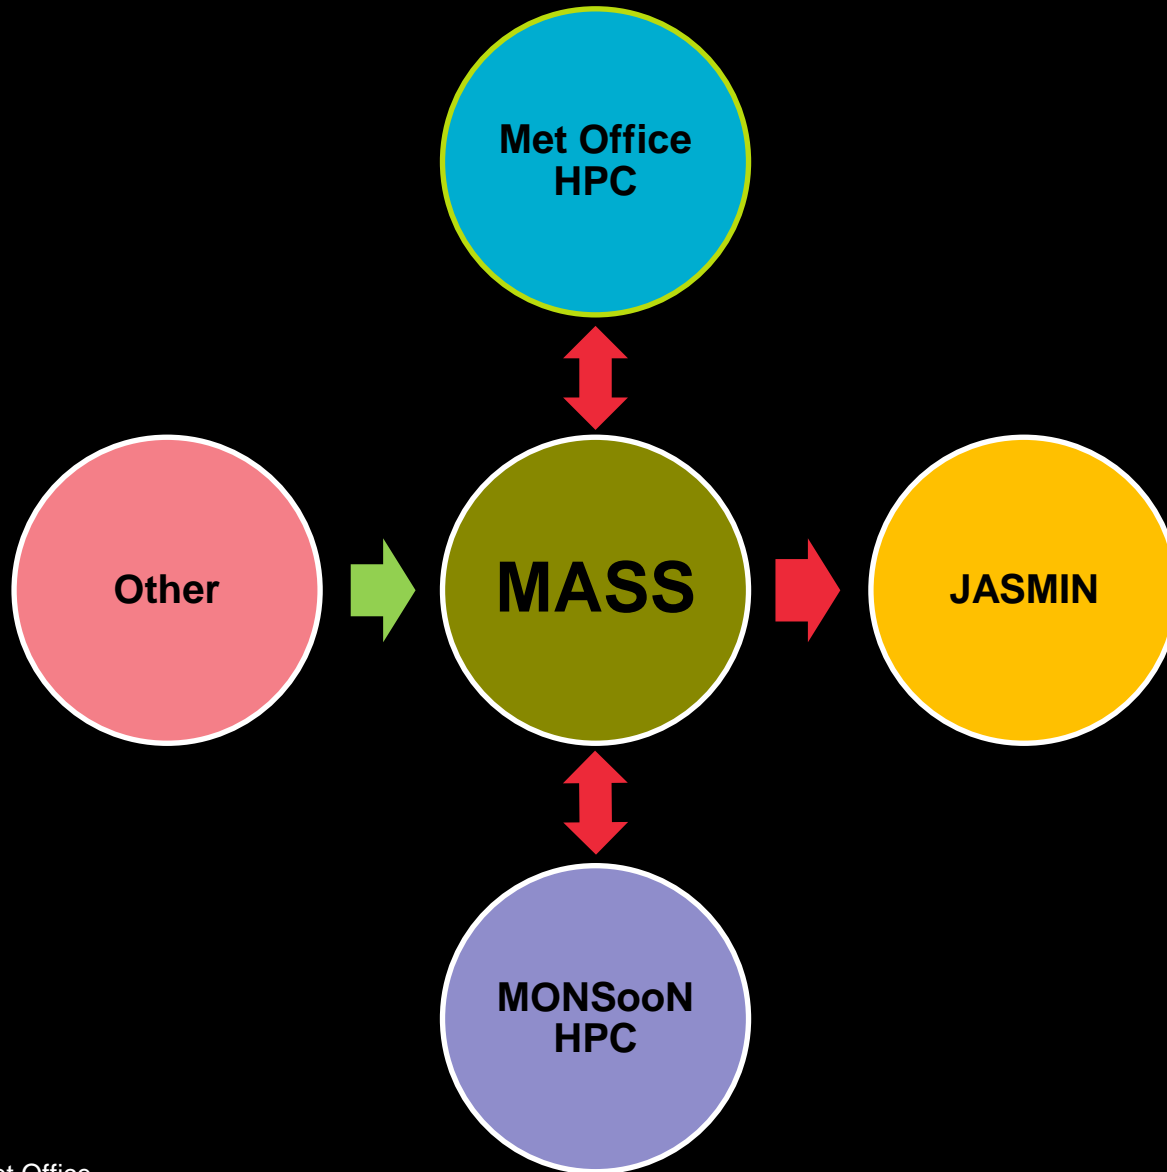


Rôle of MOOSE

- Common interface for all users
- Logical separation from vendor-specific interface
- Physical buffering, enhancing availability
- User account-management
- Access-control
- (Some) cataloguing
- Filtered retrieval for PP (UM Post Processing) and netCDF file-types
- Co-location of related data on physical media



Data Flows





MONSooN & JASMIN MASS Access allows you to...

- Have access based on a MASS project basis
- List contents of data-sets
- Retrieve files
- Filter files ('atomic access')
- Get information about data-sets:
 - ownership
 - quality-assessments, comments etc.



MONSooN & JASMIN MASS Access allows you to...

- MONSooN MASS Access allows you to:
 - Store and/or overwrite files
- JASMIN MASS Access does **not** allow you to:
 - Store or overwrite files
- Only Met Office dataset owners can:
 - Move, rename or delete files
 - Change metadata about files or data-sets



MOOSE data organisation

- Uses a directory hierarchy
- No fully searchable index exists
- Data is stored in Classes, Sets and Collections
- Data is accessed via MOOSE URIs (Uniform Resource Identifier)
- MOOSE provides ownership and access permissions to the data

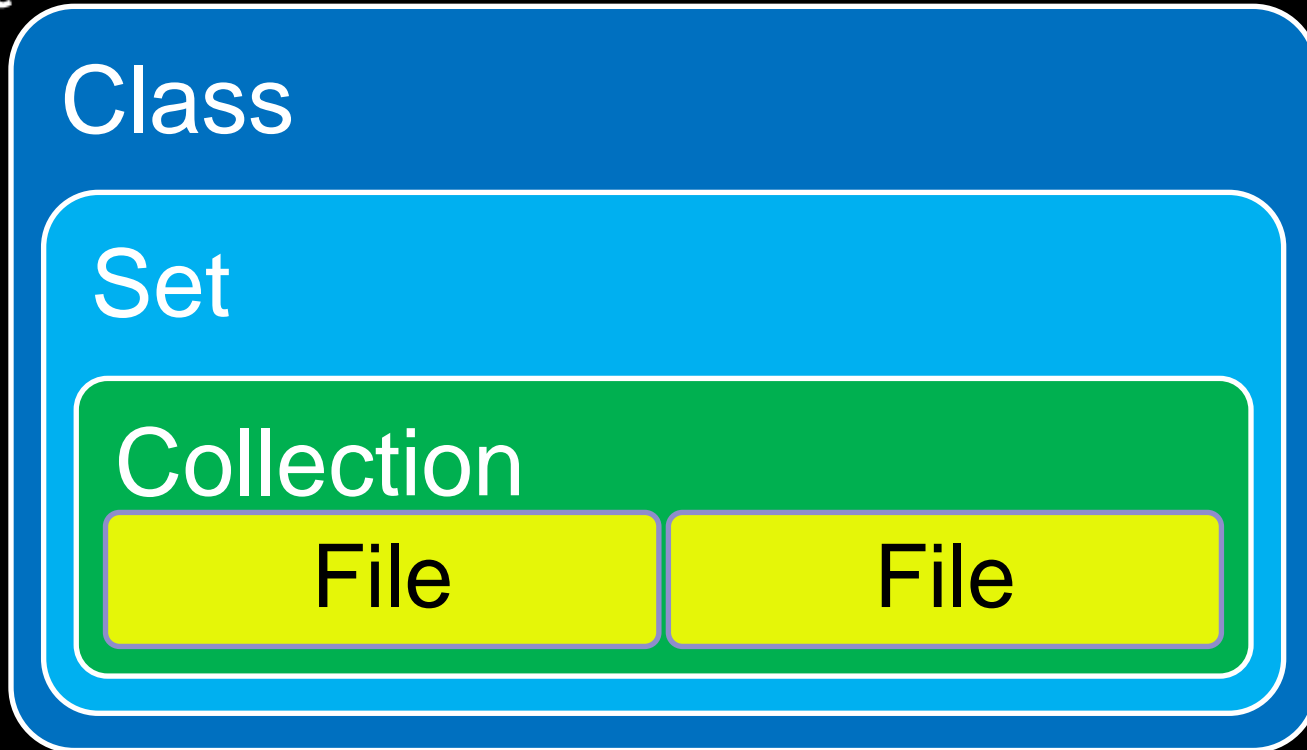


An example MOOSE URI

- `moose:/crum/antie/apa.pp`
- `moose:` - indicates a MOOSE URI
 - can be abbreviated to `mo:` or just `:`
- `crum` – is the data *Class*
- `antie` - is the data *Set*
- `apa.pp` - is the data *Collection*
 - files are stored in the Collection



MOOSE data organisation



To get access, see

<http://collab.metoffice.gov.uk/twiki/bin/view/Support/ExternalAccessToMASS>

or e-mail monsoon@metoffice.gov.uk to get started.



MO Science Repository Service

- Created a common shared environment
 - makes collaboration even easier
 - improved pull-through of science changes
 - developers work from common code base
 - share issue tracking, development
 - improves access to new releases
 - more up-to-date models at collaboration sites



What's provided?

- Subversion repositories
 - including source for UM 10.0+
- Trac environments
 - inc. documentation & commentary for UM development
- Static web pages
 - e.g. documentation for UM, JULES
 - more limited use
- Rosie suite repository & database



Implementation details – 1

- Service hosted by cloud provider
 - Open a TRAC ticket to get an account
 - <https://code.metoffice.gov.uk/trac/home/wiki/FAQ>
- Publically available repository
 - protected trunk
 - branch commits restricted to owner by default
 - strong password
 - uses keyrings (no plain text password storage)
 - Eclipse doesn't currently work with GNOME keyring



Implementation details – 2

- Repository will be mirrored at local sites
 - full mirror – not like puma / CDN
 - updated every five minutes
 - provides read-only access
 - in case service can't be reached
 - required for
 - fcm extract (performance)
 - access from batch (no authentication)
- Easy to set up and maintain



Questions and Answers...