

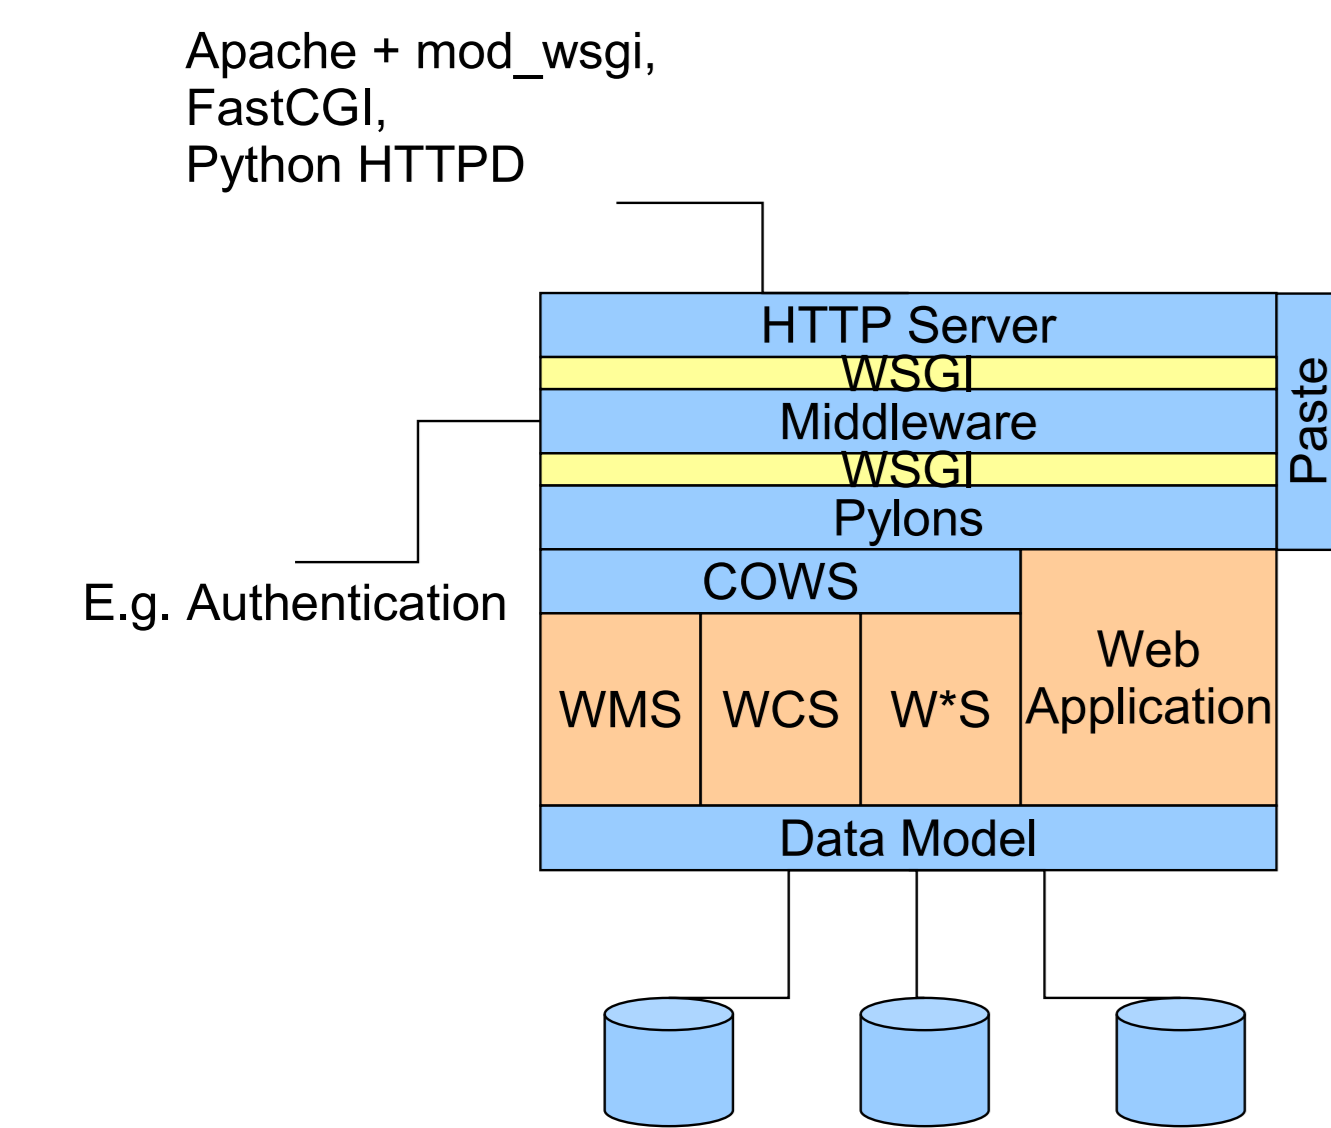
Pragmatic service development and customisation with the CEDA OGC Web Services framework (COWS)

Stephen Pascoe, Ag Stephens, and Dominic Lowe
British Atmospheric Data Centre, Rutherford Appleton Laboratory, UK (Stephen.Pascoe@stfc.ac.uk)

COWS Architecture

The CEDA OGC Web Services framework (COWS) emphasises rapid service development by providing a lightweight layer of OGC web service logic on top of Pylons, a mature web application framework for the Python language.

This approach gives developers a flexible web service development environment without compromising access to the full range of web application tools and patterns: Model-View-Controller paradigm, XML templating, Object-Relational-Mapper integration and authentication/authorization.



The Pylons Web Application Framework decouples application components from the server using the WSGI interface standard.

Application: QESDI Data Portal

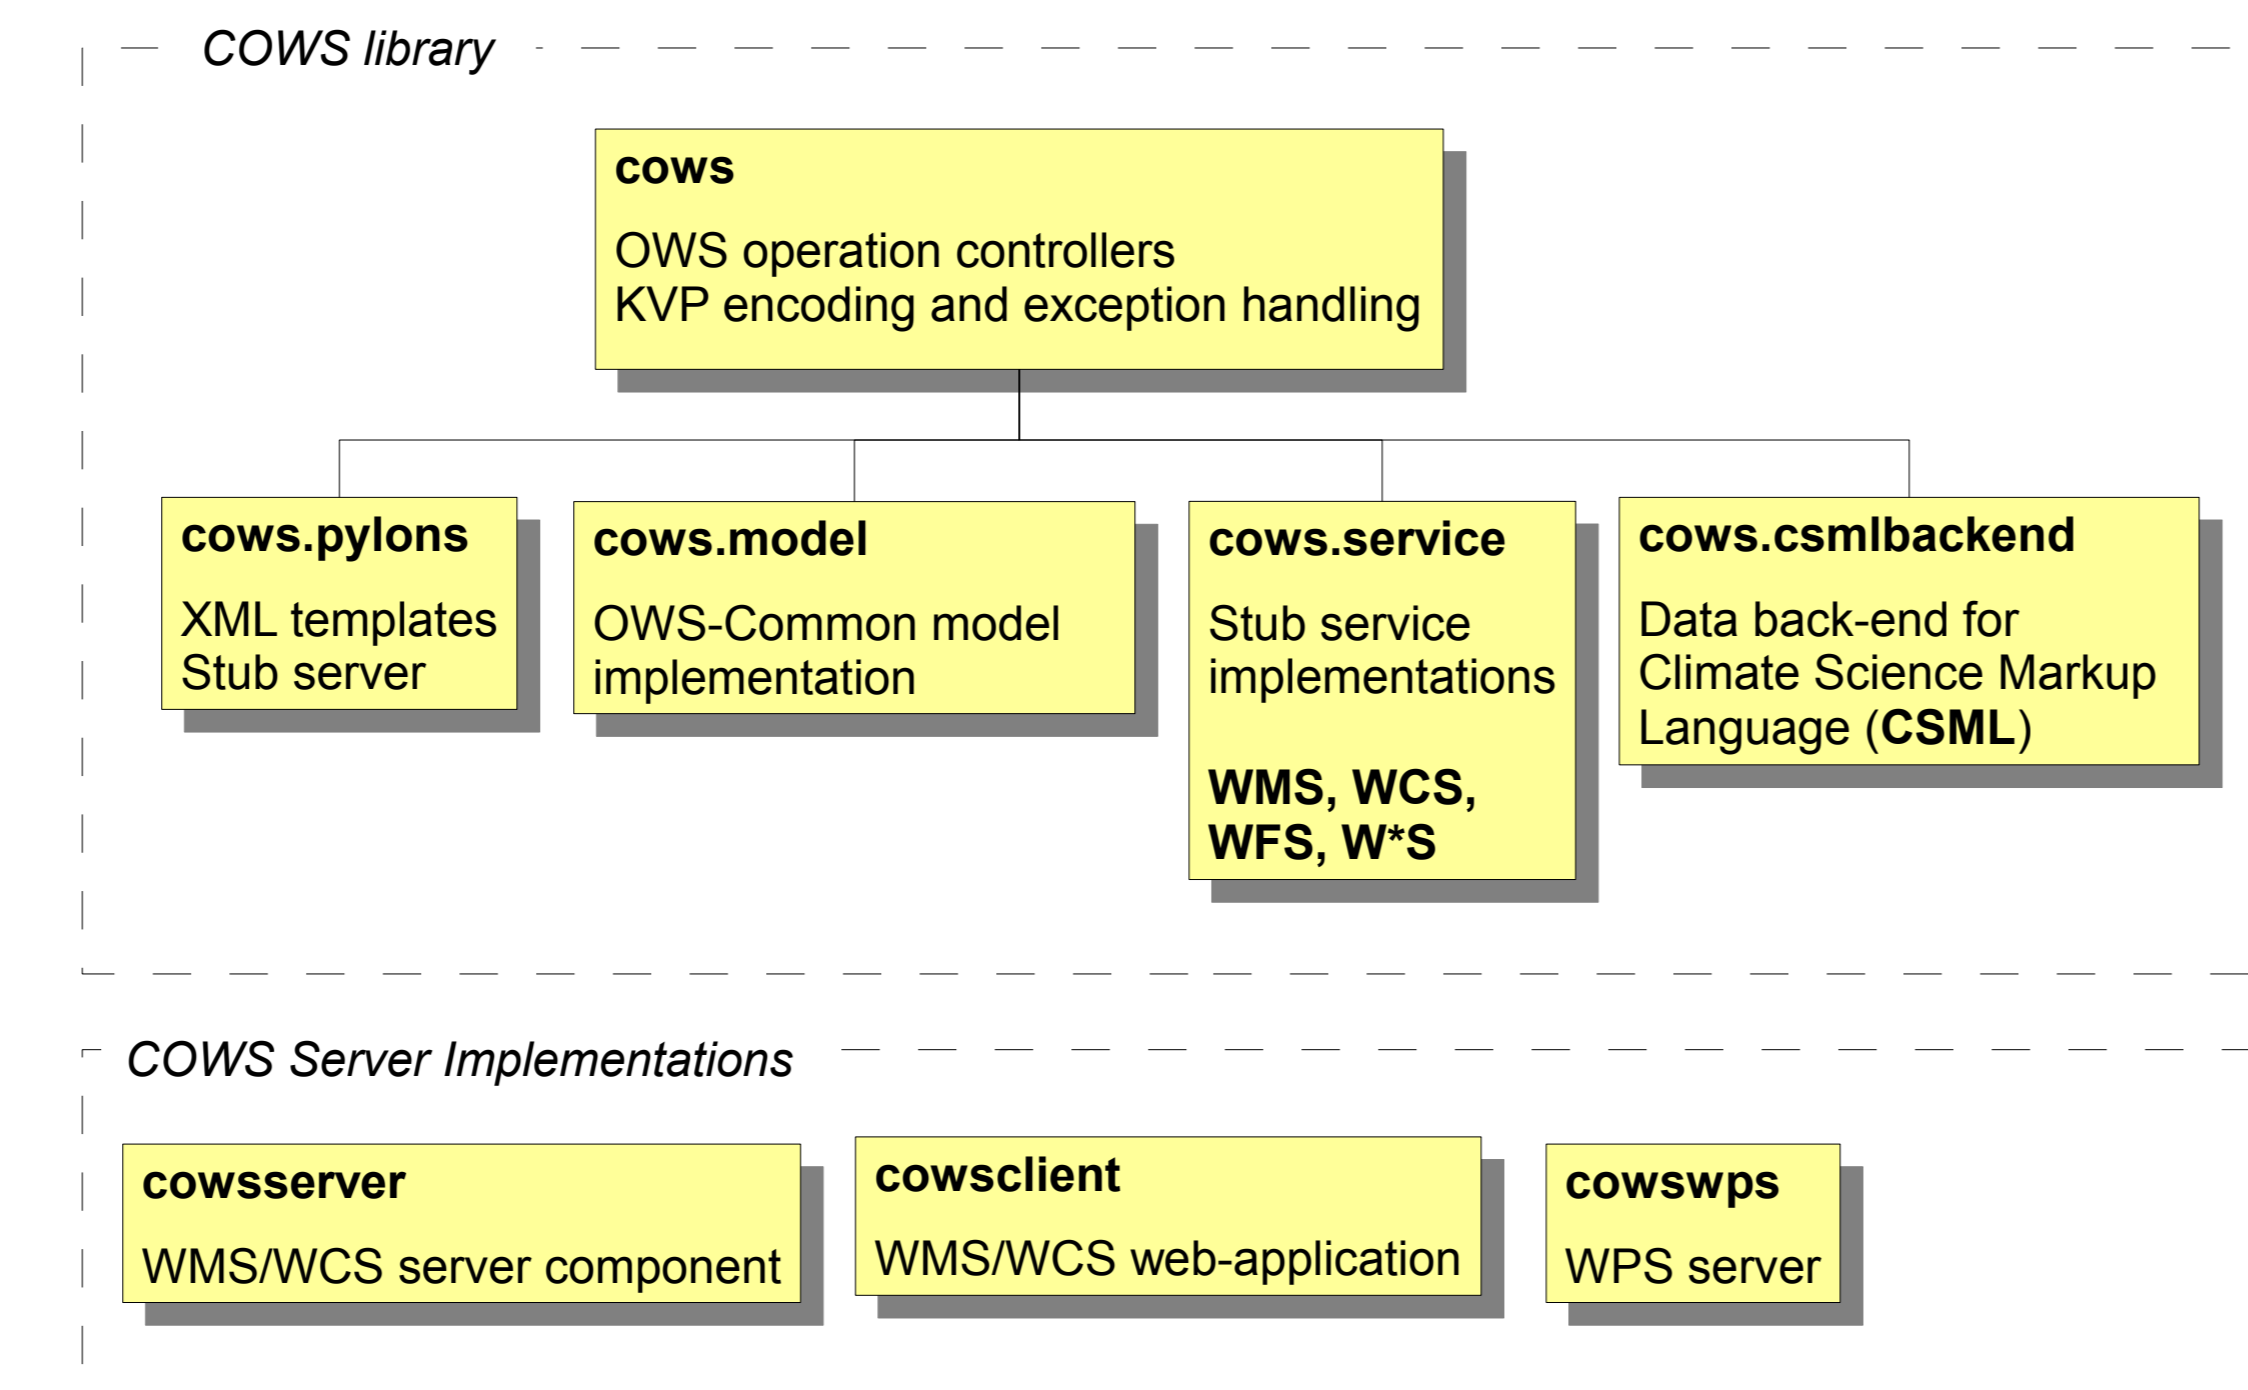
<http://phobos.badc.rl.ac.uk/qesdi/>

The QUEST Earth System Data Initiative (QESDI) facilitates exchange of data within the QUEST thematic programme. It provides access to a diverse range of data products produced within QUEST and complementary external datasets.

The QESDI portal's data visualisation interface is driven by COWS WMS server and client components. Data download is available through WCS.

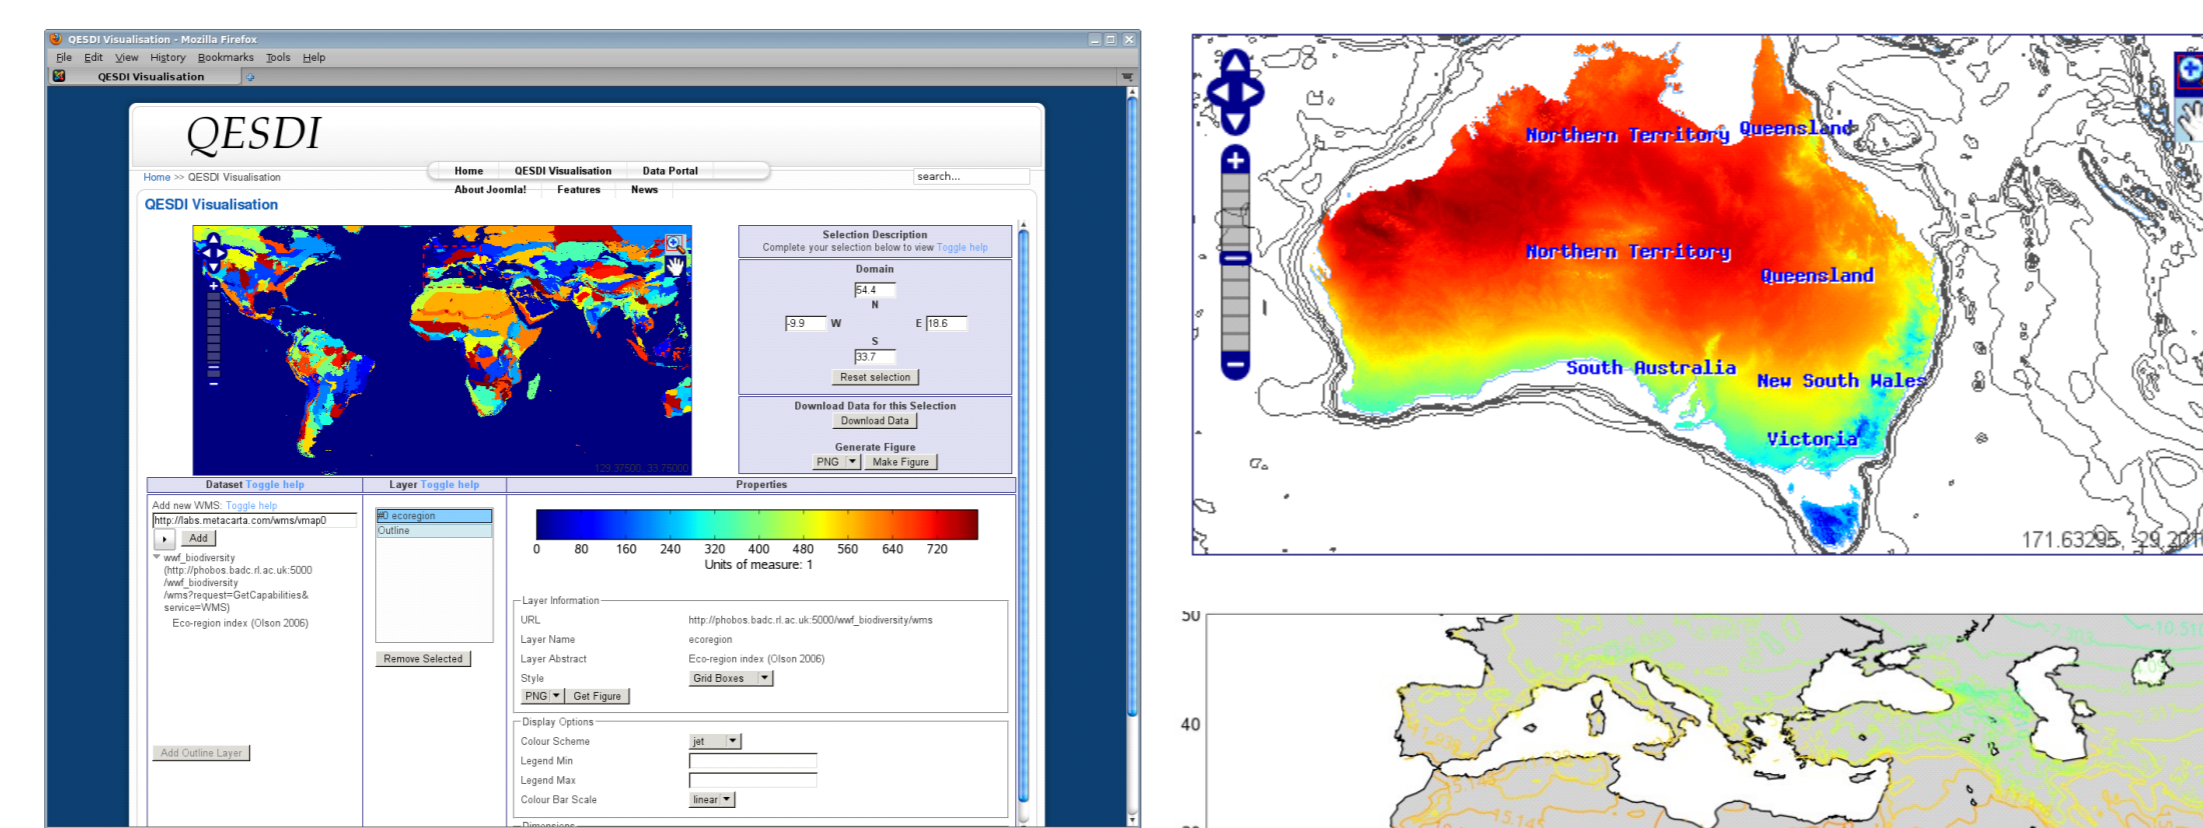
Features:

- Overlay WMS layers from internal and external WMS servers
- Browse arbitrary WMS dimensions
- Integrates with COWS Server to provide extended capabilities
- Styling Options: contouring, colourmap configuration
- Data download via WCS KVP request
- Publication quality plot generation
- Intuitive display of Climatology dimensions

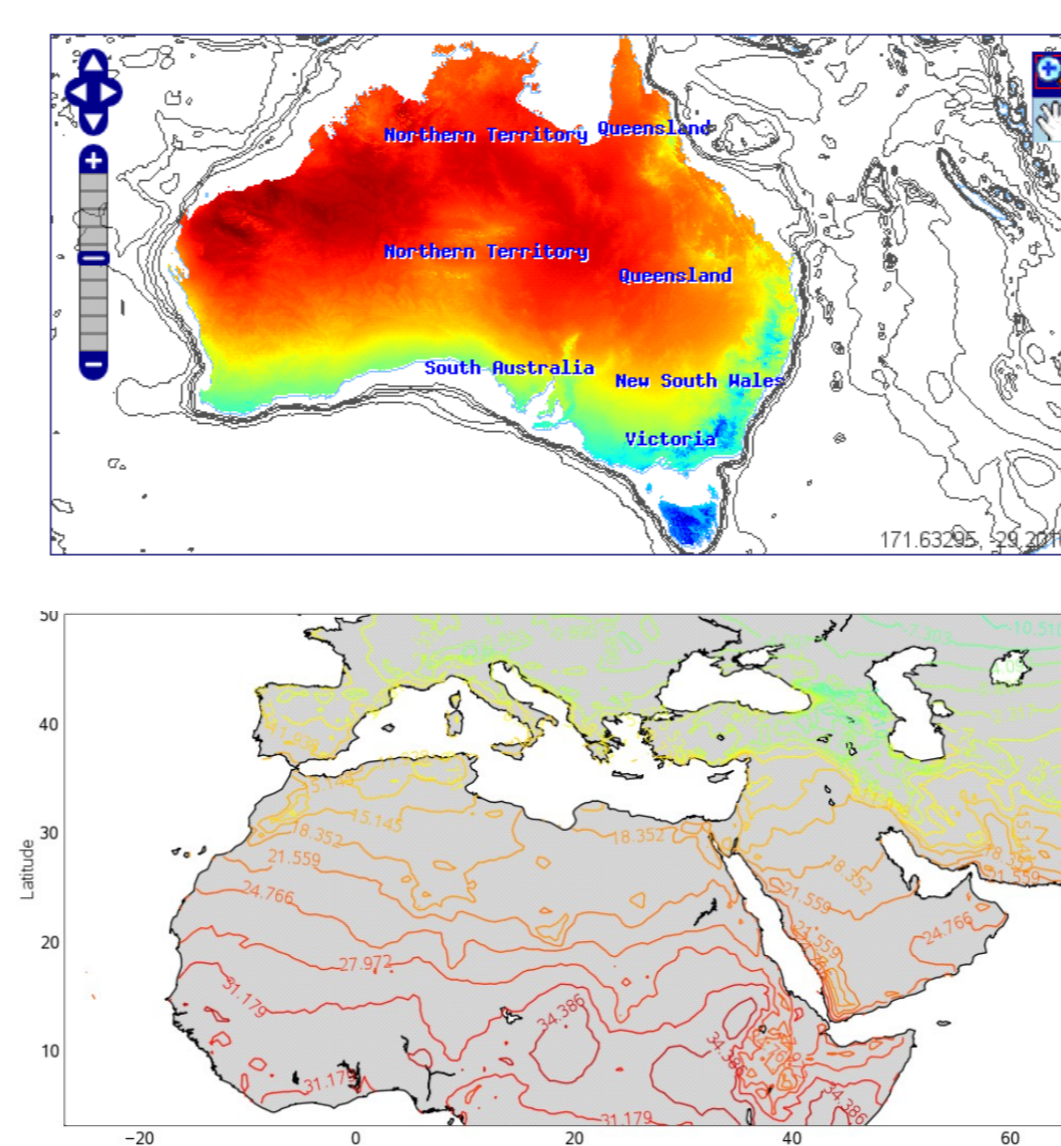


COWS server development proceeds by customising template server implementations, allowing highly customised services to be developed following the OWS architecture. We have used this feature to develop services extending OGC standards that meet the specific requirements of the Atmospheric Science community.

Many of these services leverage the GML application schema Climate Science Modeling Language (CSML) which provides a semantically rich description of underlying CF-NetCDF datasets. COWS provides WMS, WCS and WFS implementations utilising CSML.



QESDI map view showing WMS layer selection and styling options discovered from WMS GetCapabilities extensions.



The contour plotting feature of COWS server facilitates overlay of overlapping layers.

Application: UK Climate Projections '09 User Interface

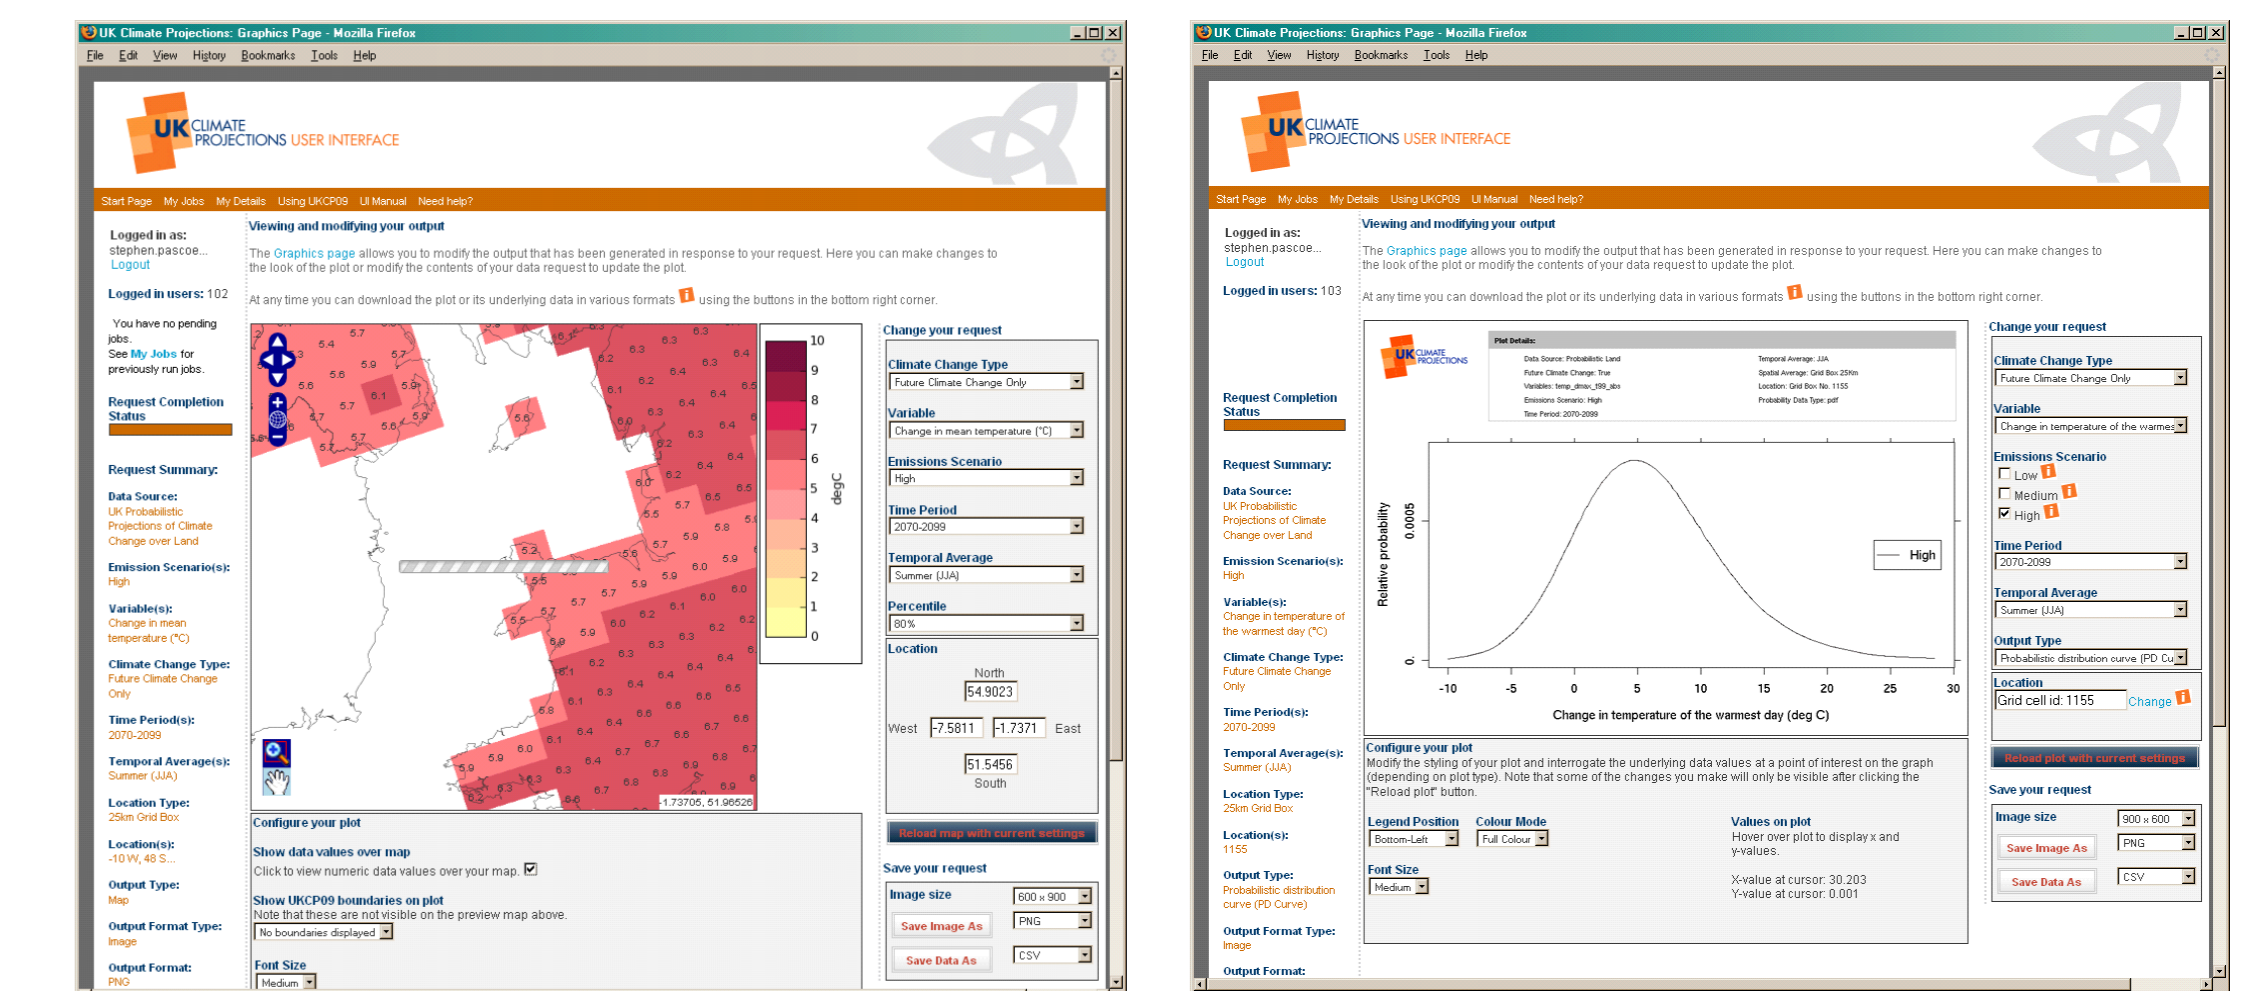
<http://ukclimateprojections-ui.defra.gov.uk>

The UK Climate Projections 2009 are a set of probabilistic data products for the UK climate over the 21st century. The UKCP09 UI provides a rich, interactive user experience including map-based selections and outputs. The implementation integrates and extends Open Source GIS technologies including GeoServer, PostGIS, OpenLayers and Tilecache. Where necessary new components have been developed using the COWS framework.

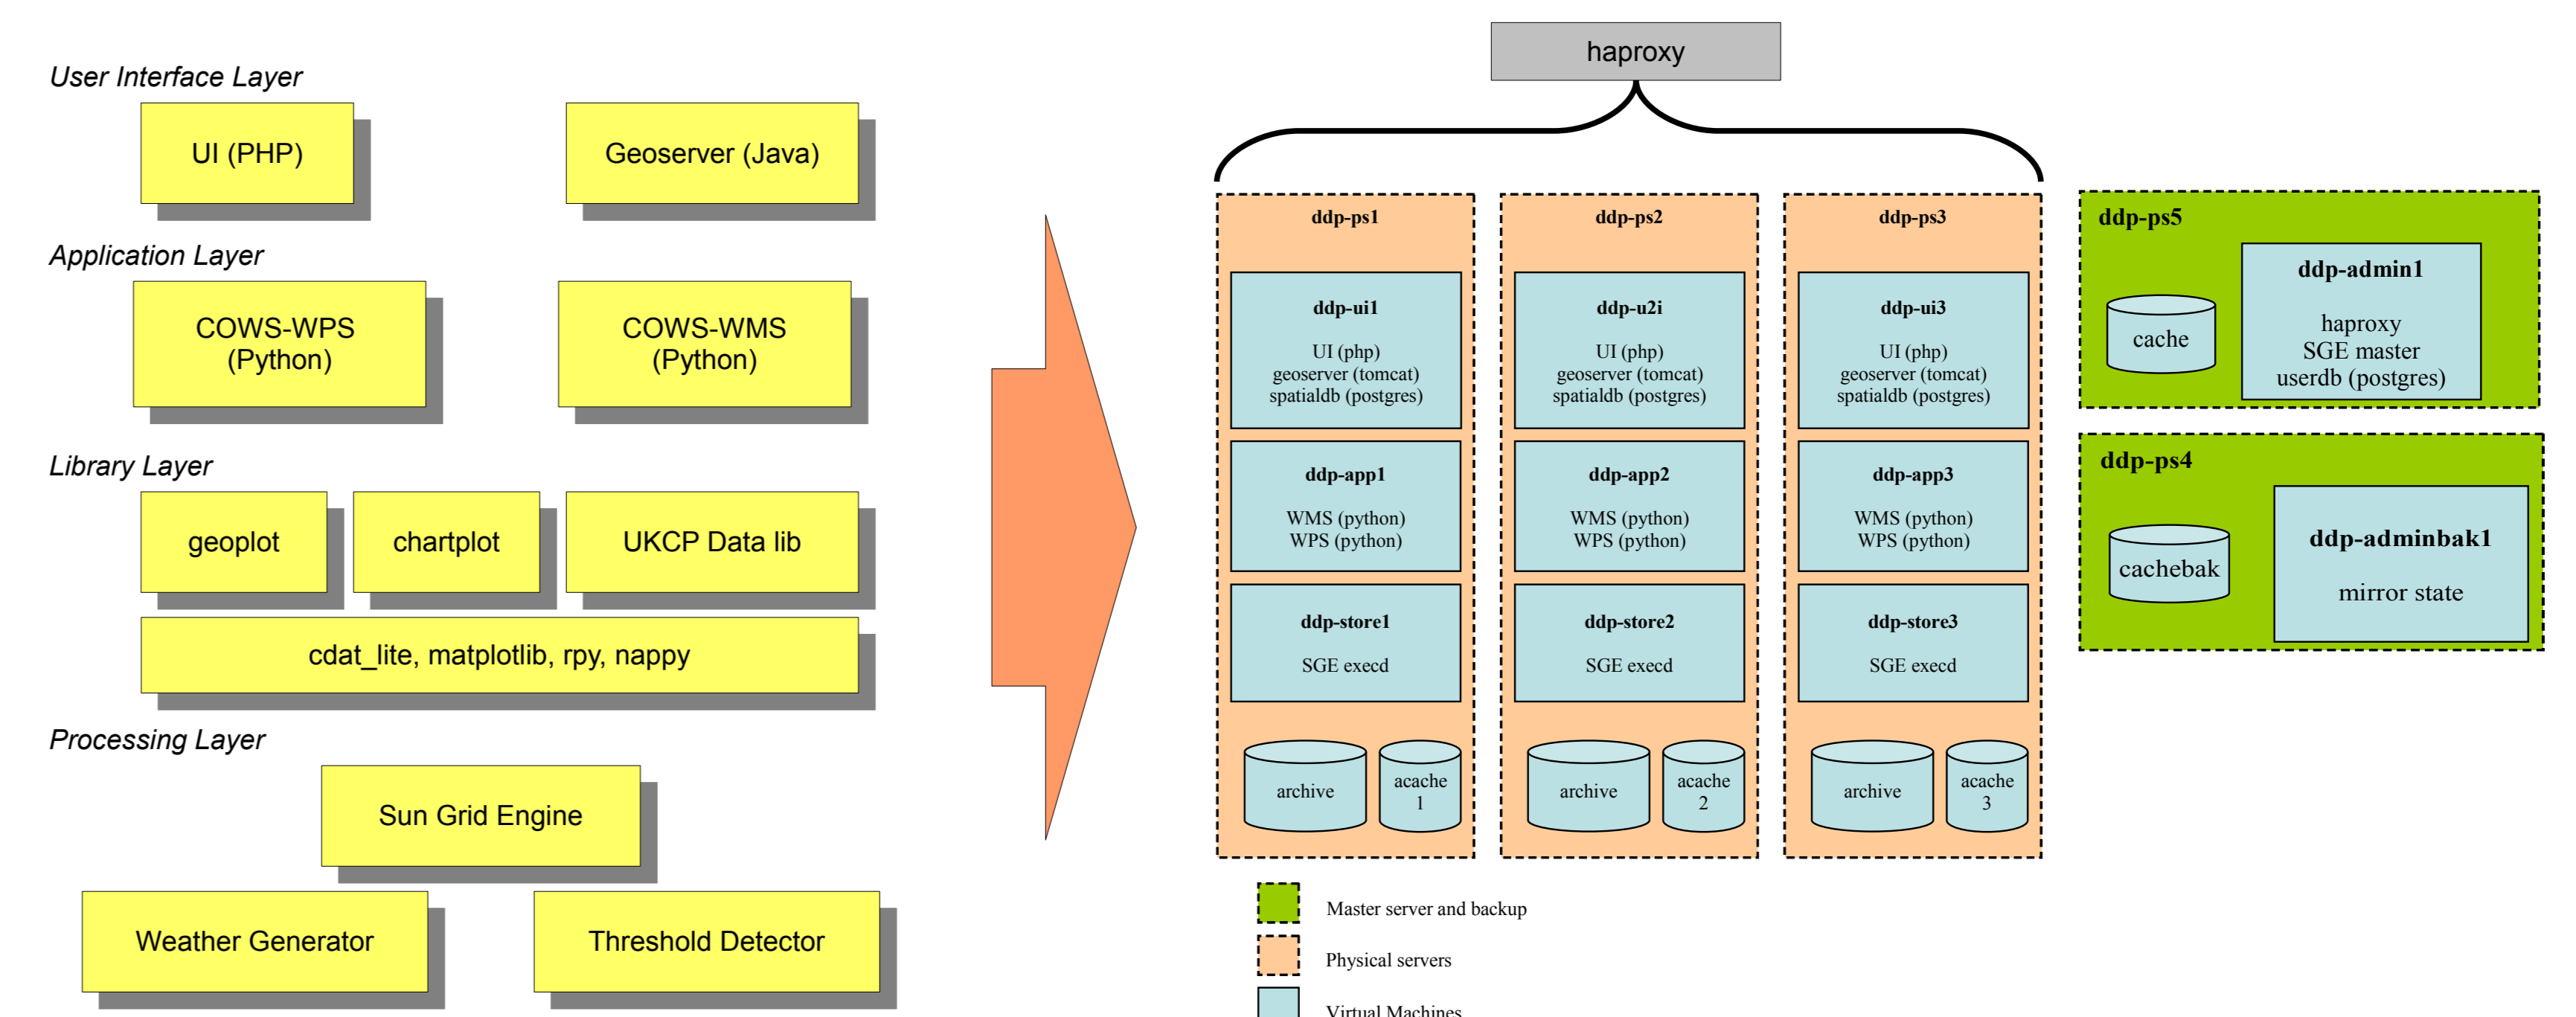
A bespoke WMS dynamically extracts data from the UKCP09 datasets. The UKCP09 UI presents these layers from the COWS-WMS along with coastlines and geometries served by GeoServer.

Back-end services for data extraction, plotting and processing are implemented as a WPS built on COWS. COWS-WPS processes range from short jobs, such as plot generation and small data extractions, to complex Weather Generator simulations that may take hours to complete.

Each WPS Process is implemented as a Python module that executes in its own environment interacting with the WPS server through a well defined API. The COWS WPS server supports a variety of process execution modes within this API. WPS Processes that require fast response times are executed synchronously within the application server or in a sub-process to avoid thread safety issues. For longer running WPS Processes COWS-WPS executes jobs asynchronous by passing them to a back-end scheduler (Sun GridEngine), allowing execution on a cluster of nodes.



The UKCP09 UI graphics page allows geospatial and probabilistic plots to be customised by the user. Requests for publication quality plots or the underlying data can then be sent to the WPS for processing before arriving in the user's download area.



The UKCP09 UI was designed from the outset to be highly scalable and fault tolerant. The system is deployed as a cluster of XEN Virtual Machines, dividing the system into a set of functional VM types. These VMs are spread over multiple physical servers to provide both physical and virtual system failover. Load-balanced VM types are deployed in parallel allowing the system to scale up and down according to load, whereas the VMs managing state employ a separate hot-backup failover strategy. A VM image server manages the backup and restoration of VM images allowing rolling upgrades to the system.

The deployment configuration allowed the UKCP09 UI to quickly increase its capacity in response to predicted demand during launch period from 11 to 23 Virtual Machines (Figure 2). In the 23 VM configuration the system was independently verified to support up to 3000 simultaneous sessions and over 1000 simultaneous active users making requests to the UI, WMS and WPS layers.