



eReefs Relocatable Coastal Modelling Service

high-resolution operational
GBR modelling at scale, and
on demand

Sharon Tickell | 2023-10-19

Australia's National Science Agency



Great Barrier
Reef Foundation



Australian Government
Bureau of Meteorology



Australian Government



AUSTRALIAN INSTITUTE
OF MARINE SCIENCE



Queensland
Government



eReefs GBR4

- “GBR Regional Model”, ~4km grid
- Extends into coral sea to include SEC, avoid continental shelf features
- 220 x 500 x 44 cells (4.84M)
- 1m vertical resolution at surface
- **Software:** EMS-SHOC
- **Compute:** HPC at ~150:1
 - Hindcast scenarios from 2010
 - NRT operation since 2014
- **Results Storage:** 133TB+
- **Does not resolve reefs, islands**





eReefs GBR1

- “GBR Shelf model”, ~ 1km Grid
 - 510 x 2390 x 48 cells
 - Only 50% of surface and 22% of all cells are wet...
 - => still 12,871,584 cells to calculate!
- Software: EMS-SHOC
- Compute: HPC at ~30:1
 - Too slow for hindcasts
 - NRT operation since 2014
- Results Storage: 166TB+
- Still unable to resolve small reefs, islands, estuaries...





RElocatable COastal Model (RECOM)

Solution: Shard it!



- Only model the interesting bits
- Same EMS-SHOC model code
- GBR1 or GBR4 boundary forcing
- Custom curvilinear grid (down to 100m resolution, any orientation, any cell-size)
- Custom bathymetry, point-load sources
- Web interface to assist with model setup
- Portable model execution environment
 - (reproducibility FTW)

eReefs III: RECOM1 (Prototype), on Bowen Cloud

eReefs IV: RECOM1 Results Visualisation added

eReefs V : RECOM2 + public cloud deployment

Scalable => more users => more impact



eReefs RECOM v2^{ALPHA}

RECOM model configuration: eReefs GBR4 Hydrodynamics v2.0
Configure your run using the options below, then press Submit Run.

Model configuration

- Grid definition **required**
- Grid masking and interpolation **required**
- Show interpolated bathymetry
-

Temporal extent

- Run params
- Label

Map view | Run status | Add map layers

Legend:

- steemap
- topographic
- satellite
- dark canvas
- GBR1km Boundary
- GBR1km Grid
- GBR4km Boundary
- GBR4km Grid
- GBR6km Boundary
- GBR6km Grid

Step 1: Configure & Submit

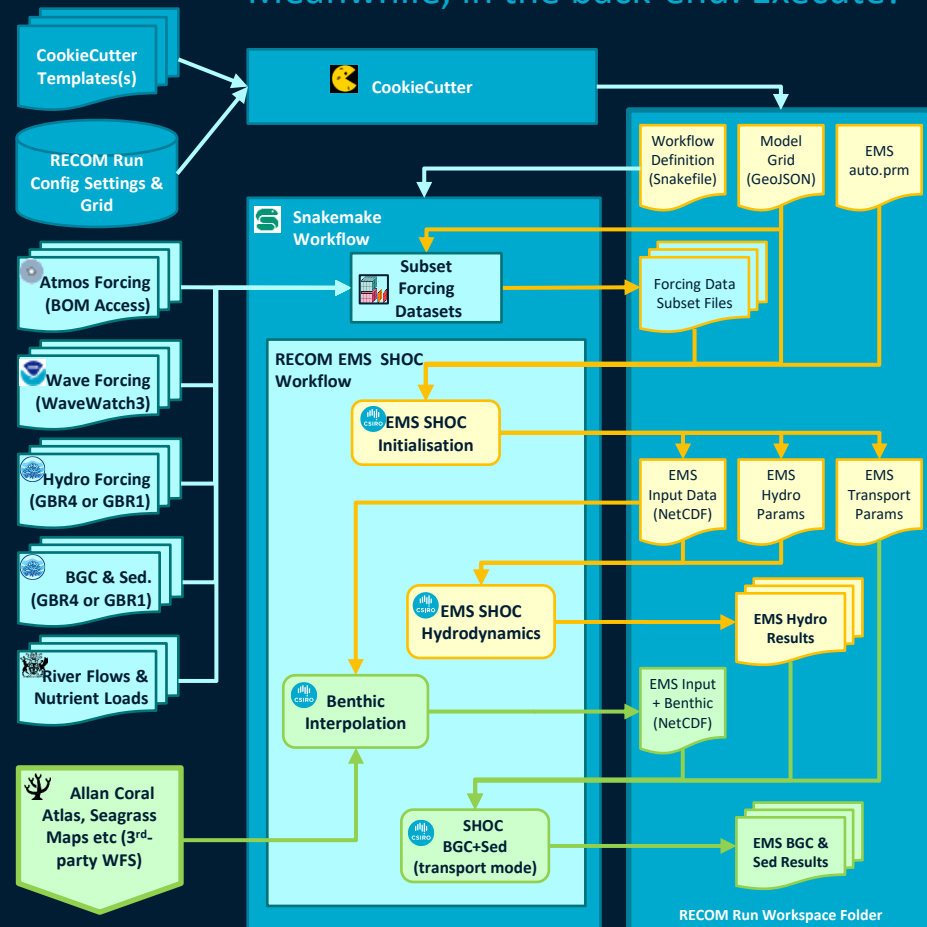
Step 2: Wait & Monitor...

eReefs RECOM v2^{ALPHA}

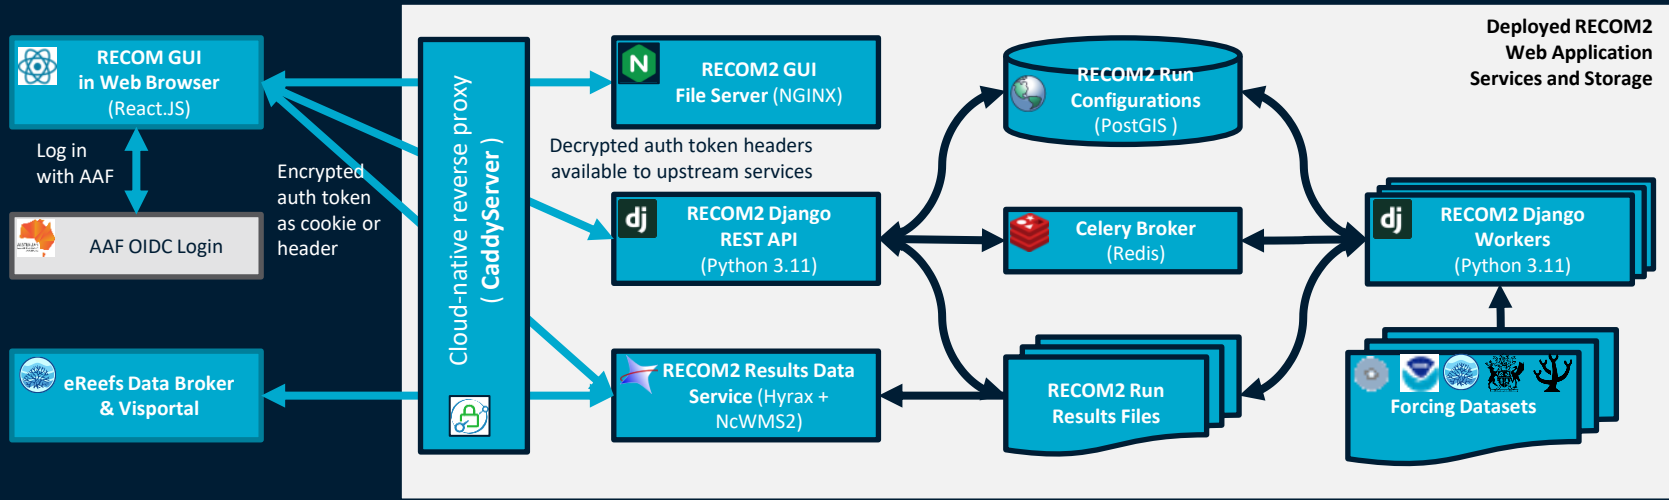
RECOM Runs

Run id	Model config id	Run name	Owner	Model start	Model stop	Status	Published to	Actions
Redacted						Finished	Data Broker	Messages Workflow log Open run Grid JSON Delete
Redacted						Finished		Messages Workflow log Open run Grid JSON Delete
Redacted						Finished		Messages Workflow log Open run Grid JSON Delete
Redacted						Finished	Data Broker	Messages Workflow log Open run Grid JSON Delete
Redacted						Finished	Data Broker	Messages Workflow log Open run Grid JSON Delete
Redacted						Finished	Data Broker	Messages Workflow log Open run Grid JSON Delete

Meanwhile, in the back-end: Execute!



Workflow diagram for a generic RECOM Model Run, showing the input dependencies, processing steps and output files required at each workflow stage. Log files and other diagnostic outputs are omitted for clarity.



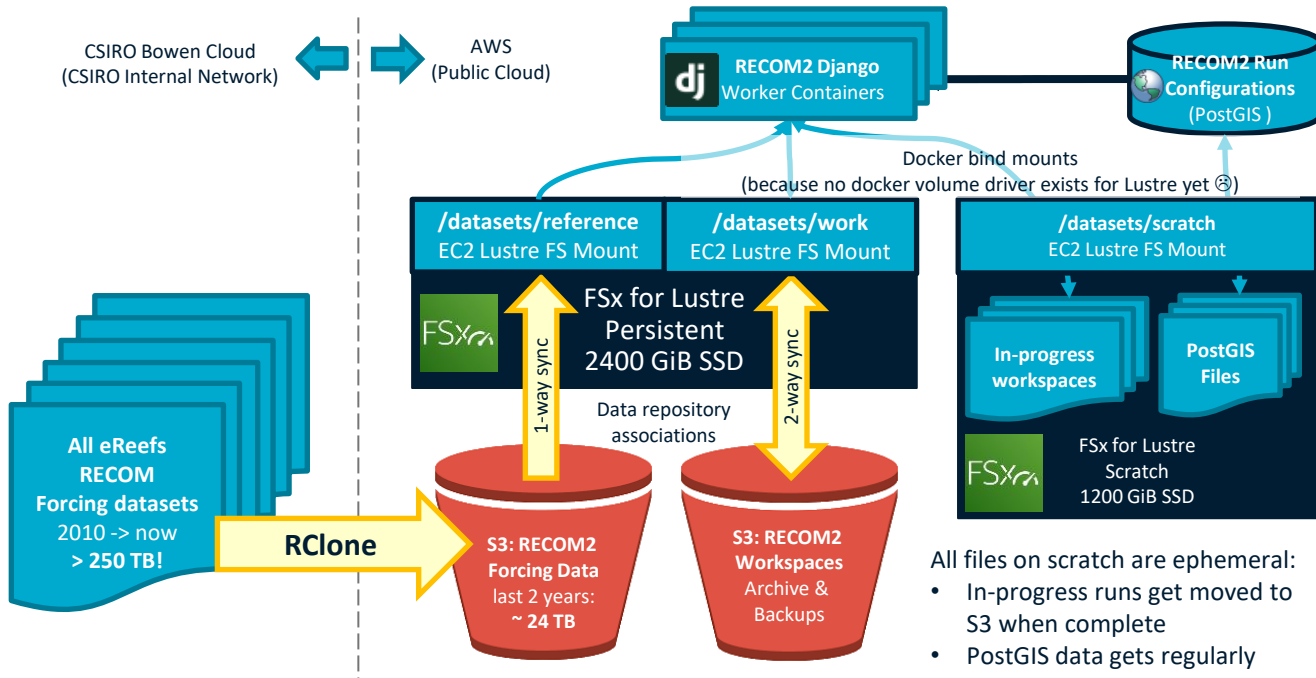
What are the moving parts?

RECOM2 uses a common web application deployment pattern: it is a browser-based user interface in front of a Django REST API that uses a message queue to communicate with back-end Django worker processes. Application data lives in a PostGIS database, while forcing data is file-based.

We use the Australian Access Federation's OIDC service for authentication, handle authorisation at the API layer, and deploy docker swarm services for a flexible, cloud-agnostic runtime environment.



The S3 Bucket Challenge



All files on scratch are ephemeral:

- In-progress runs get moved to S3 when complete
- PostGIS data gets regularly backed up to S3 with pgdump

Cloud Storage Budgets...

Making forcing data accessible to the subset step has been our biggest challenge, because:

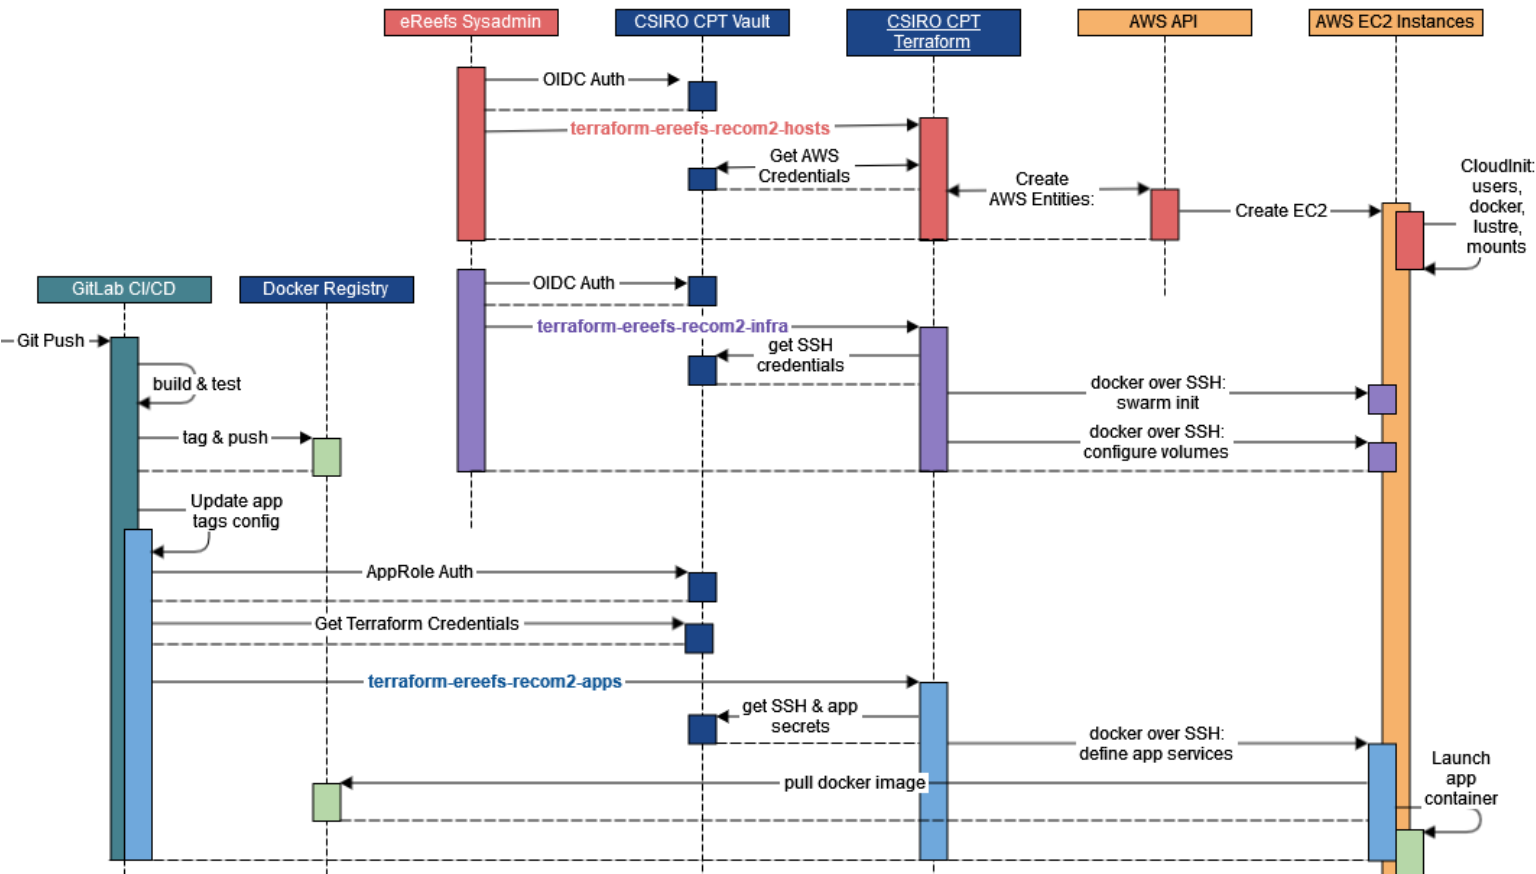
- RECOM requires about 12TB of NetCDF forcing data files for each simulation-year.
- That gets expensive fast! On AWS, S3 storage is the only sensible (and affordable) option.
- BUT EMS software was designed to run on HPC and to access data from a posix filesystem...

Our Solution: use the AWS “FSx for Lustre FS” product to make a just-in-time parallel posix cache.

Figure 3: RECOM2 persistent data is stored in S3 buckets and mounted to docker containers via a FSx for Lustre cache that has data repository associations configured for those buckets. Ephemeral data is stored in a FSx for Lustre scratch filesystem that can be easily shared between multiple containers. Only a 2-year subset of all possible RECOM forcing is available to the Cloud deployment for now – that may change if future budgets permit.



Automate All The Things



Step 0: Create S3 buckets, rsync

Step 1: Create AWS Entities

Step 2: Container Cluster Init

↑ Sysadmin Setup

Automated Deployment ↓

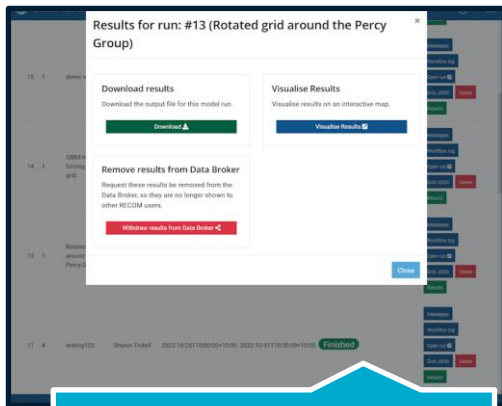
Step 3: Build & register application image

Step 4: Deploy application container



RECOM2 Results: In-situ vs offline analysis

RECOM2 has limited support for in-situ analysis of completed run results, but allows you to leverage the full power of the eReefs visualisation tools if you're willing to share:



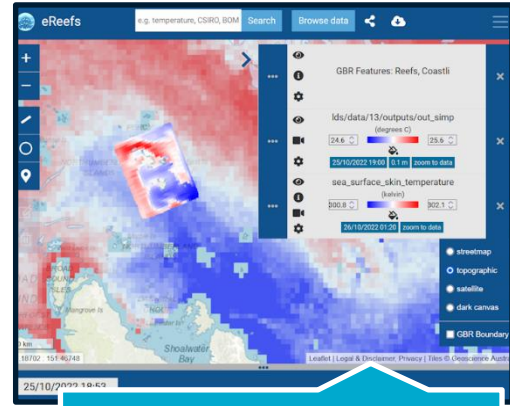
Export and Analyse:

- Download just your results, or your entire run workspace folder
- Analyse your results offline using any method you like



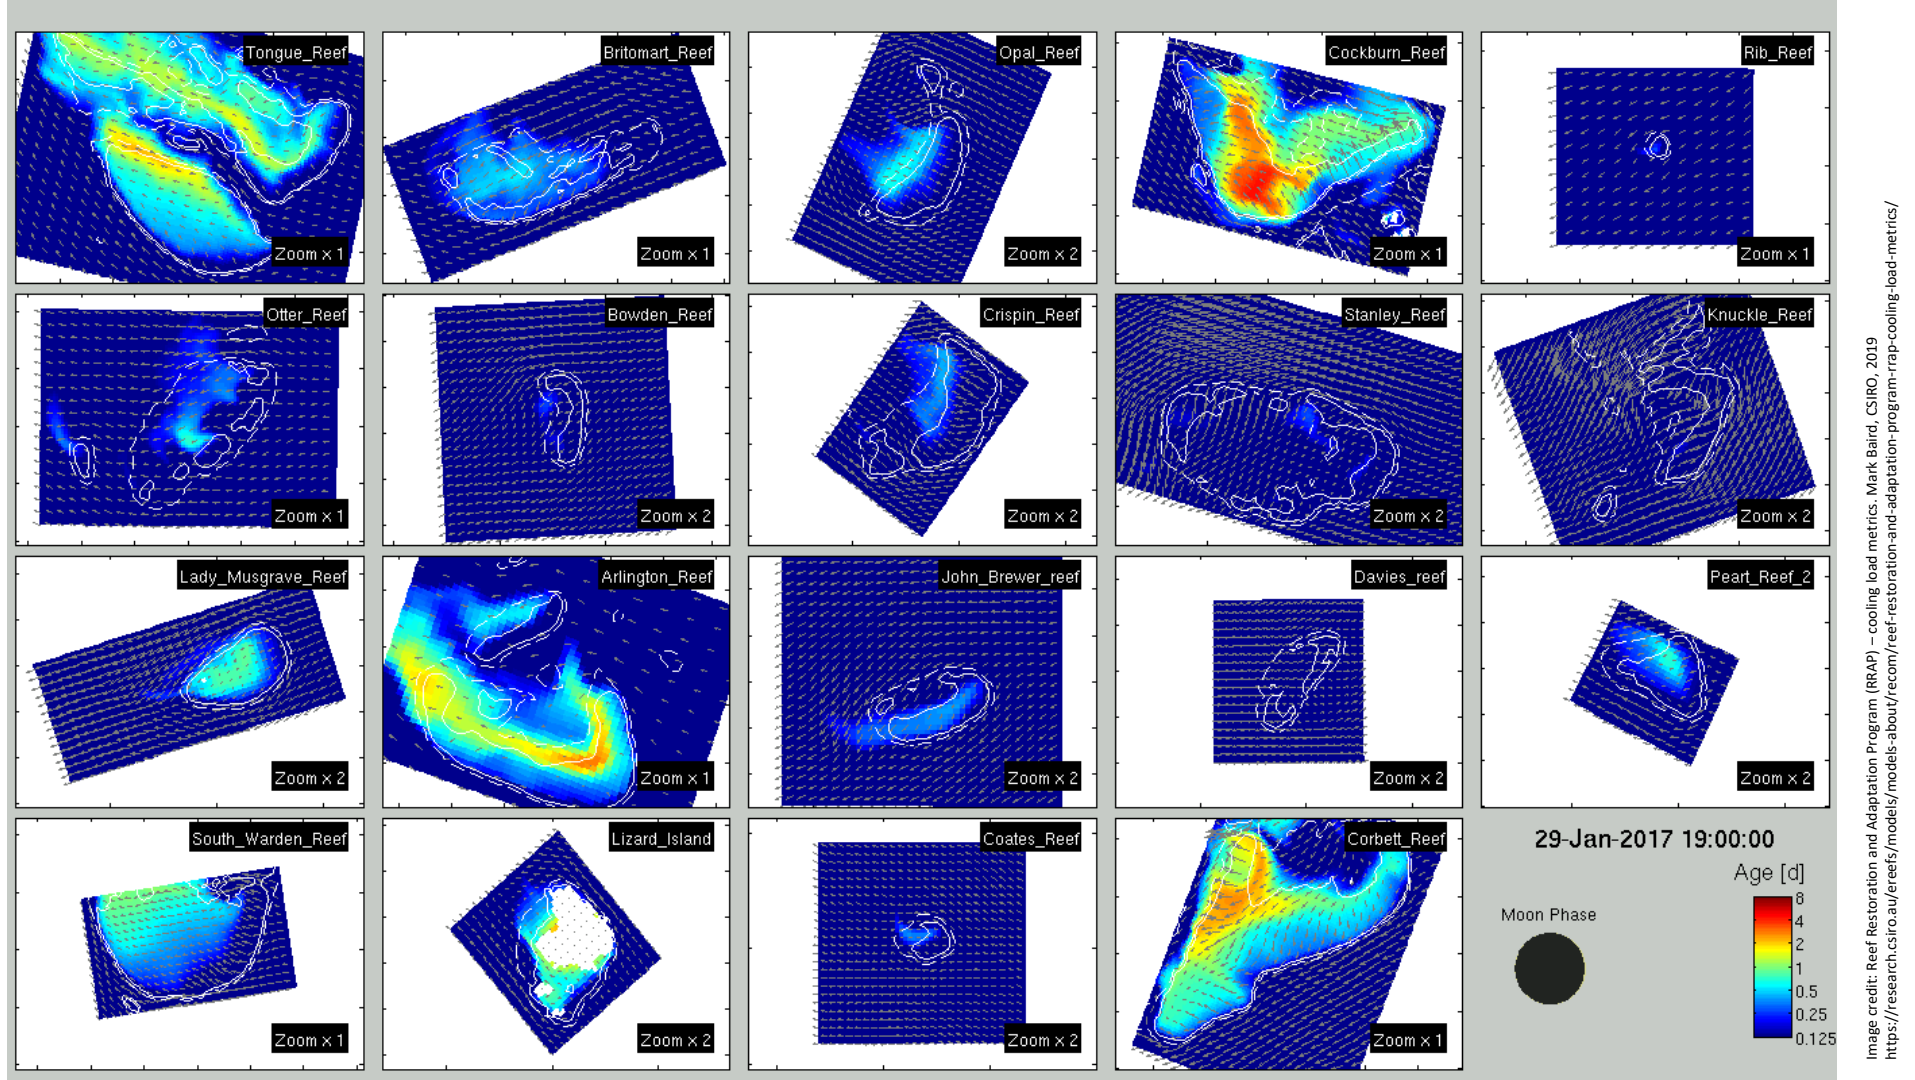
In-Situ Visualisation:

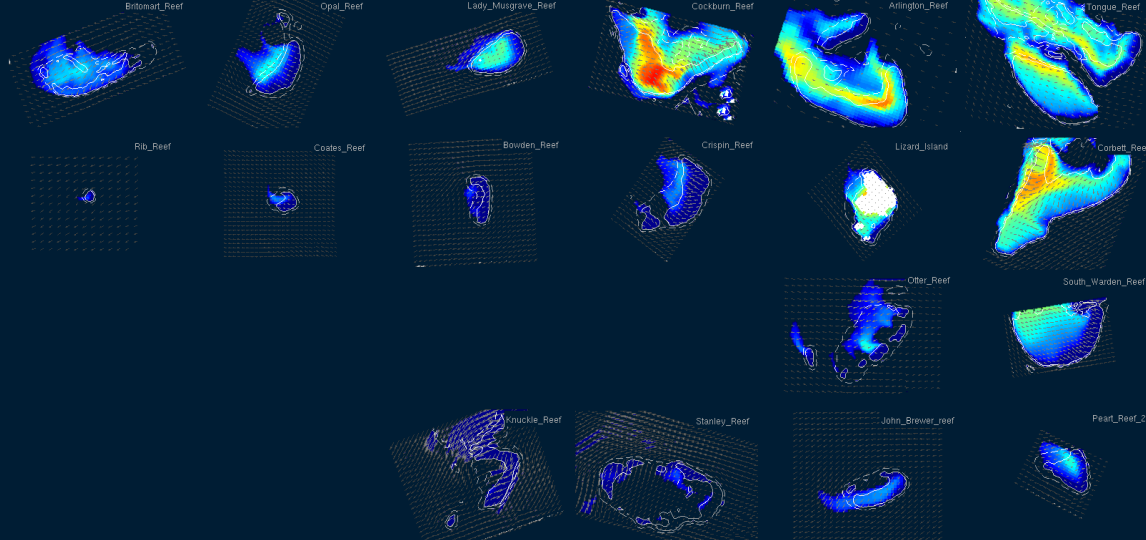
- Using Hyrax Godiva3 Viewer (NcWMS) from RECOM2 GUI
- Sanity check your RECOM results prior to publication



Publish to eReefs (WIP):

- Opt to publish your data as CC-BY
- Your results will be available as layers in the eReefs Visportal
- Combine, compare, analyse, and animate online!





Thankyou

Speaker:

Sharon Tickell


Senior Software Engineer, CSIRO Environment

Sharon.Tickell@csiro.au

RECOM2 Development Team:

Sharon Tickell, Farhan Rizwi, Jack Beardsley,

Tim Heap, Matthew Wildie, Daniel Wild



Further Information:

RECOM2: <https://recom.ereefs.info/>

Science: <https://research.csiro.au/ereefs/>

Support: team@ereefs.info

Community: <https://forum.ereefs.info/>