



PEcAn to the cloud

Moving ecological forecasting from
high performance computer to cloud


Christopher Tate

Principal Software Engineer in Red Hat Research



Software tools


What kind of software tools do forecasting researchers need?




Open source



Accessible



Reusable



Scalable



Meet the Ecological Forecasting Team




Jeff Simeon - BU SAIL



Greg Frasco - BU SAIL



Shashank Karthikeyan - BU SAIL



Dongchen Zhang - BU PHD Student



Michael Dietze - BU Professor



Christopher Tate - Red Hat

Red Hat Collaboratory with Boston University

A partnership advancing research and learning in emerging technologies





2023

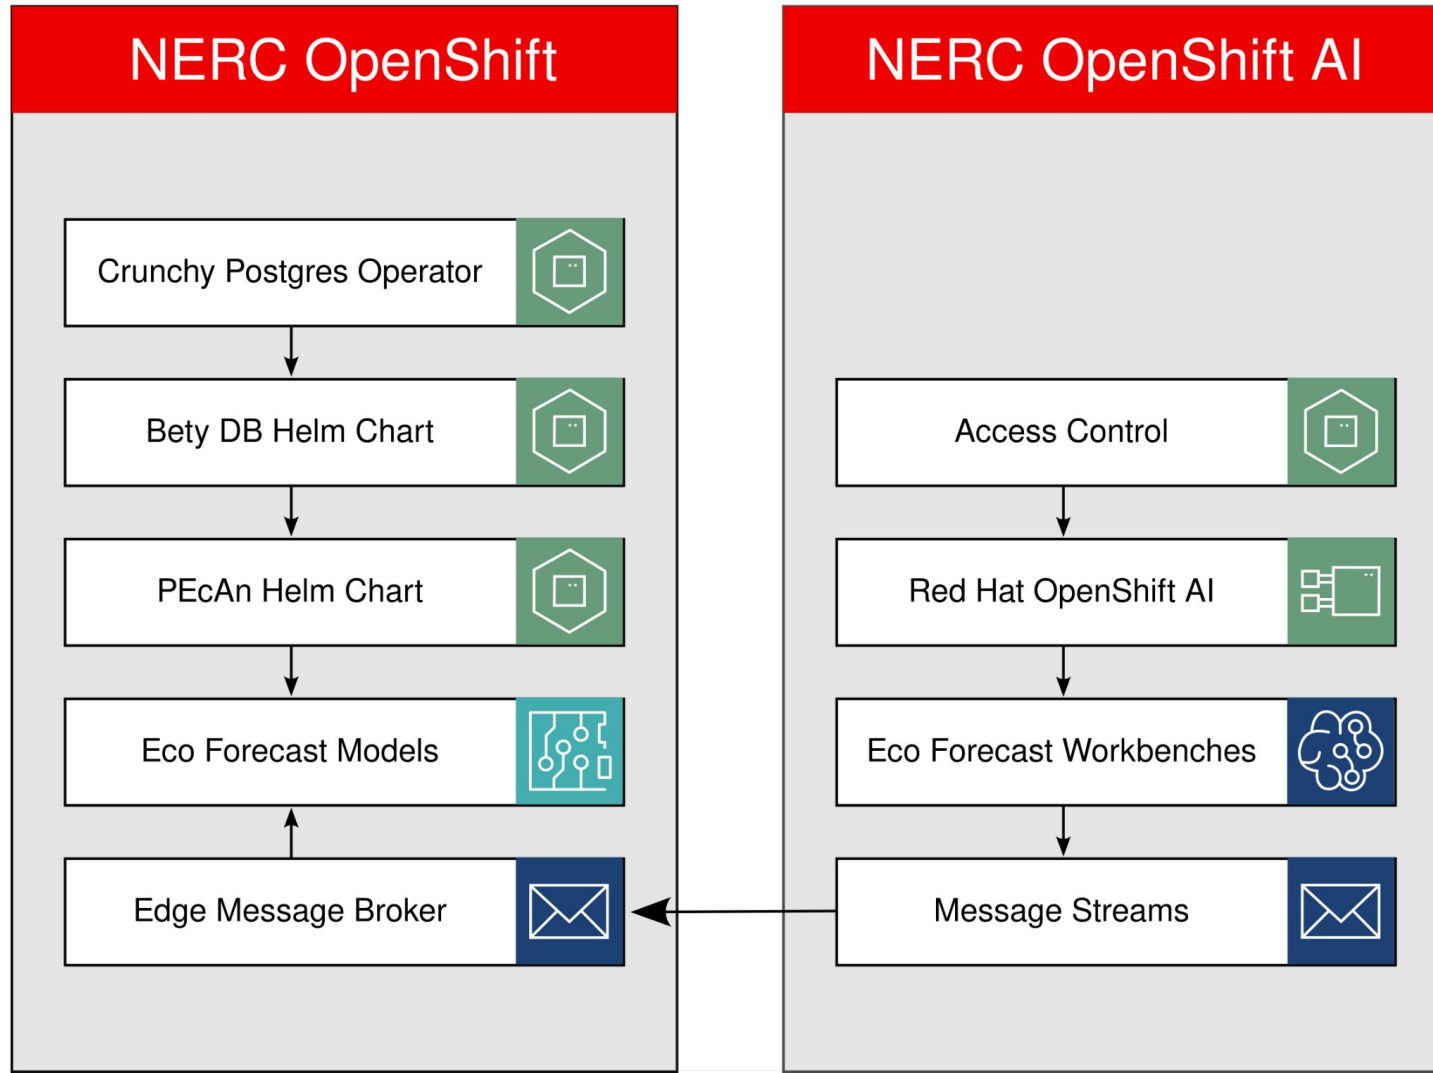
Prototyping a Distributed, Asynchronous Workflow for Iterative Near-Term Ecological Forecasting

Michael Dietze (BU), Christopher Tate (Red Hat), Yannis Paschalidis (BU), Atefeh Hosseini (BU) will prototype an accessible community infrastructure to generate ecological forecasts at scale, focusing on the development of a cloud-native workflow that can handle an asynchronous, event-driven, distributed approach to execution.

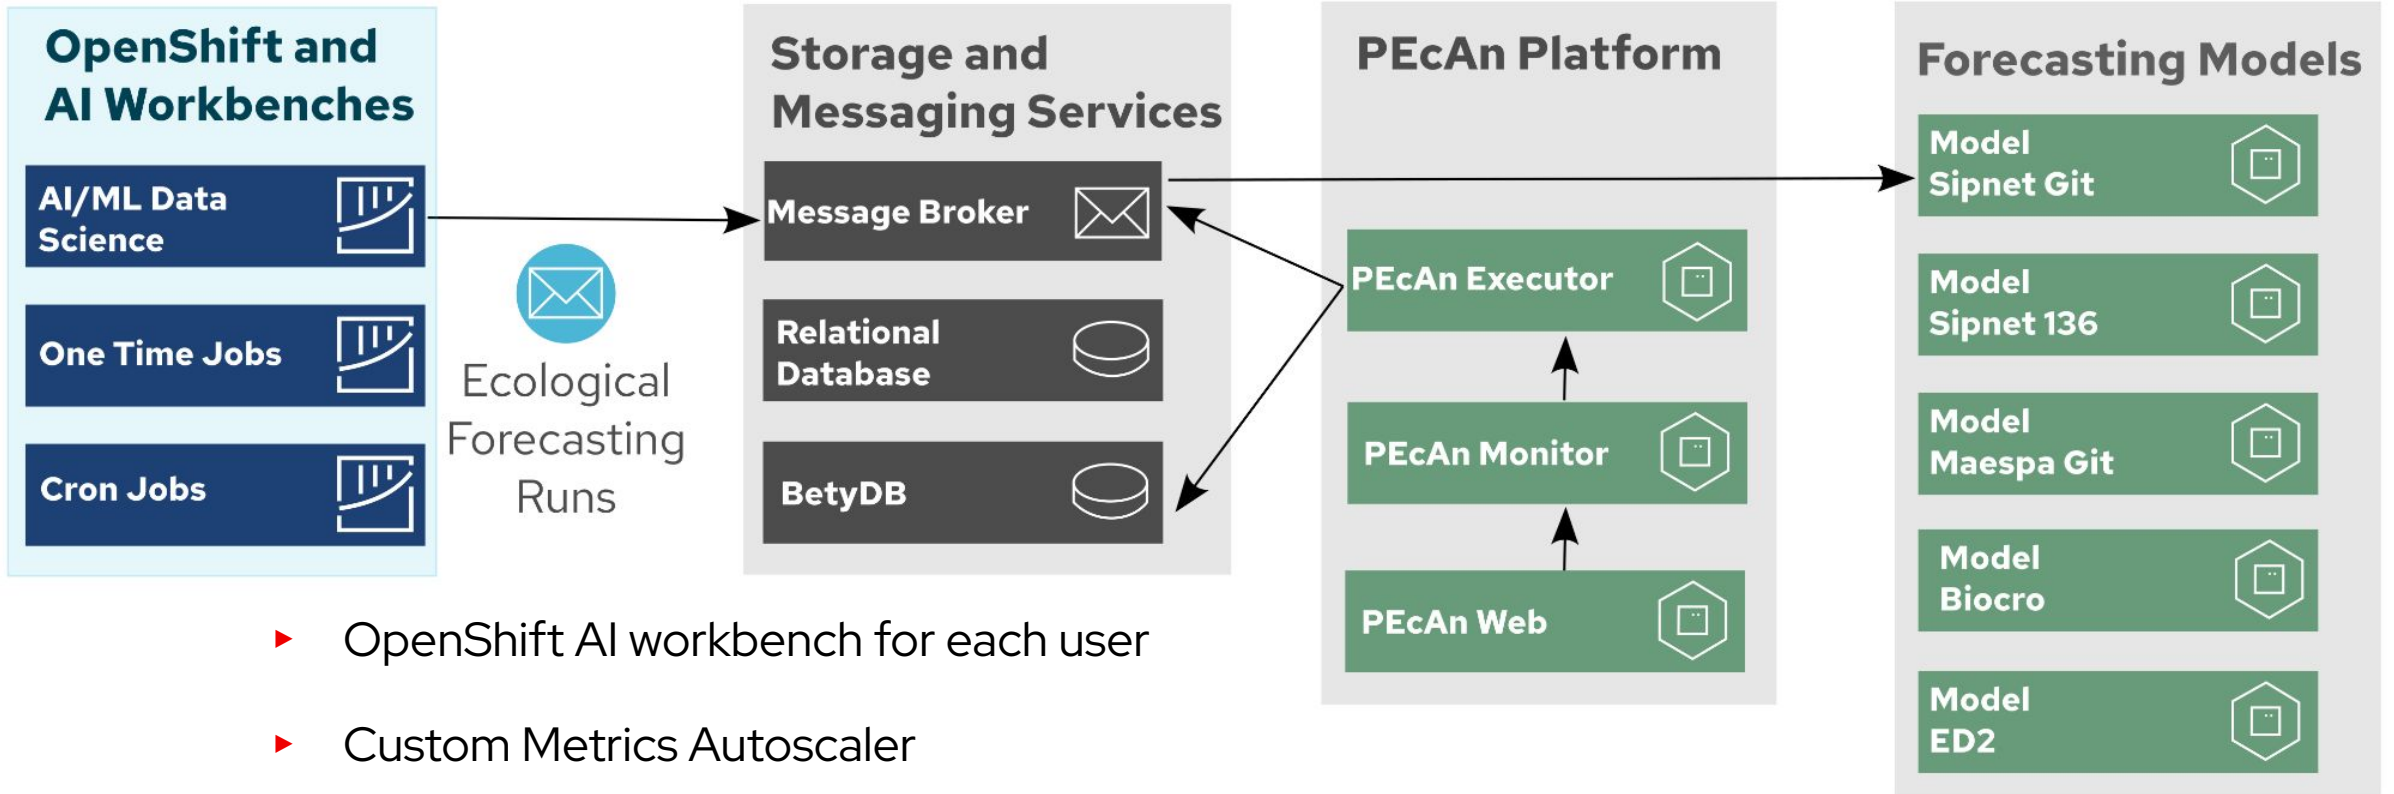
New cloud architecture for PEcAn

PEcAn

PEcAn Unconstrained Forecast
R-Studio Data Science



PEcAn in the New England Research Cloud



- ▶ OpenShift AI workbench for each user
- ▶ Custom Metrics Autoscaler
- ▶ One Time Jobs
- ▶ Cron Jobs
- ▶ Object Storage
- ▶ Microservices in the same project

Open cloud solutions

PEcAn's new event-driven, distributed design

Open source branching

Prototyping new cloud-native software relies on branching

Distributed storage

Consider object storage, ephemeral storage, and persistent storage needs

Pod autoscaling

Scale the number of running model pods based on the number of messages

Rsync strategy

Pull forecast data from one pod and back with ephemeral storage



Why ecological forecasting on Red Hat OpenShift?

Back to “What kind of software tools do forecasting researchers need?”

Open source

The trusted enterprise open source cloud platform

Accessible

- Test your forecasts in the free [OpenShift Developer Sandbox](#)
- Develop open source models and microservices with [OpenShift Local](#)
- OpenShift Local microshift preset runs on laptops with fewer resources

Reusable

- Easily migrate your forecasting to OpenShift clusters On Premise, or on AWS, Azure, or Google
- Also deploys to Red Hat MicroShift at the Edge

Scalable

Scale from your laptop to a datacenter like the New England Research Cloud



OPENSIFT



New England Research Cloud

[ABOUT](#) ▾[SERVICES](#) ▾[DISCOVER](#) ▾[SUPPORT](#) ▾[COSTS & BILLING](#) ▾

NERC Services


NERC offers virtual machines with diverse computing resources (CPU/GPU), a containerized platform for microservices workflows/pipelines, an AI/ML analytics platform, and persistent data storage via Block and Object storage, manageable through S3-compatible APIs.

[Our Services](#)

Future milestones



Asynchronous, event-driven scheduler




Elastically launch data ingest containers



More sites, data constraints, and models




Who makes a successful research collaboration?



Subject matter experts

Support from ecological forecasting community experts




University professors

Proposals, and high level design, support, and organization



Graduate students

Academic knowledge and technical domain expertise working with the data



Software engineers

Industry experience in cloud software development and deployment



New England Research Cloud

Computing resources for innovation for academic researchers and technology professionals