

Hands-on with AWS Ronin and Spot Instances, a very user- and budget-friendly gateway to HPC and compute work on the AWS cloud

*Convener Name*¹

Kevin Jorissen¹, Ben Thurgood², Haofei Feng³

¹ Amazon Web Services, Seattle, USA, jorissen@amazon.com

² Amazon Web Services, Sydney, Australia, btgood@amazon.com

³ Amazon Web Services, Sydney, Australia, haofeif@amazon.com

GENERAL INFORMATION

- Is this workshop half-day or full-day? *Half-day*
- Does the workshop include a hands-on component? *Yes*
- Are there any constraints on the number of attendees? *No, we can accommodate a large number of participants if there is a lot of interest. Amazon will provide the necessary cloud computing resources.*
- Are there any special seating or table requirements (e.g. for breakouts, teams)? *No, except that participants will need to work on a laptop.*
- Are there any technical requirements beyond AV and access to wireless network? *No, but power strips would be nice so people can plug in their laptops.*

DESCRIPTION

Participants will learn how to start and manage AWS compute servers and parallel clusters using a very user-friendly platform offered by RONIN in Australia. The Ronin platform is already successfully deployed in production at several Australian research institutions. It minimizes the learning curve and provides safeguards for the majority of researchers, who need access to HPC and compute resources but don't want to have to learn to become Linux sys admins to do so. (Because science is more interesting.) Workshop participants will have an opportunity to try it all out for themselves. The workshop will also be of interest to central HPC/eResearch support folks, who may discover in the Ronin platform a way to provide their user community with access to much-needed HPC resources - with the administrative tools to keep an eye on things.

Along the way, participants will receive a brief overview of Amazon's HPC services, and of using discounted "Spot instance" servers to get more compute cycles out of their research budget. This will help them decide whether the AWS cloud is right for their work, and where to find more information.

1 AWS HPC Introduction: AWS has a complete spectrum of HPC resources for researchers working on large computational problems, from parallel compute clusters and a suitable server type for every problem; to container and serverless solutions; and high-performance storage and network elements. We'll start with a quick overview of the art of the possible and how to choose the right tool for the (HPC) job.

15 minutes

2. **AWS Ronin:** AWS Ronin is an incredibly easy-to-use web application that allows researchers and scientists to launch AWS compute servers and parallel compute clusters for HPC. It's a life saver for domain scientists who need a simple and safe way to access compute power without a steep learning curve. It's also a blessing for HPC/IT/eResearch staff and RSEs who want to tailor and offer these simple-but-powerful HPC tools to the institution's end user community, while using the Ronin admin dashboard to keep an eye on the activity and resource consumption.

30 minutes

3. **AWS EC2 Spot Instances:** With up to 90% discount compared to On-Demand prices, AWS EC2 Spot Instances allow research scientists to take advantage of unused computing capacity in the AWS Cloud, thereby getting more scientific output for your fixed research budget. It is commonly used for genome sequence computation¹, machine learning, containerized workflows, and more.

10 minutes

4. **Scientific tutorials:** We'll log into the Ronin platform in a web browser. We'll work through how to run a workload on AWS EC2 Spot instances, and how to do an MPI (HPC) job on a parallel cluster. Additional tutorials will be available for fast learners. AWS support will be available in the classroom to help participants. You can also bounce your ideas off them for applications in your own lab or classroom.

120 minutes

WHO SHOULD ATTEND

The skills we teach will be useful to a wide range of practitioners: researchers, educators, research/IT support staff, research software engineers, data scientists, and more.

WHAT TO BRING

Attendees need to bring their own laptop of any kind. Workshop activities will take place on cloud computing resources accessed in a web browser. Attendees should be familiar with the (very) basic functionality of the Linux shell. The workshop content will be available in English only.

REFERENCE:

1. Karczewski K, Fernald G, Martin A, Snyder M, Tatonetti N, Dudley J (2014) : STORMSeq: An Open-Source, User-Friendly Pipeline for Processing Personal Genomics Data in the Cloud [<https://doi.org/10.1371/journal.pone.0084860>]