Australian Government
Geoscience Australia

Real-time streaming from the Cloud

Geoscience Australia's data centre migration to Amazon Web Services

Ryan Ruddick and Brandon Owen

Introduction

As part of the Australian Government's National Positioning Infrastructure capability, Geoscience Australia is leading a test project for a Satellite-Based Augmentation System (SBAS) in the Australasian region. Underpinning this project is a network of continuously operating Global Navigation Satellite System (GNSS) ground stations streaming dual frequency multi-constellation observations in real-time to the satellite uplink facility in Uralla, New South Wales. To achieve the high level of availability and continuity required for this project we chose to re-architect our exiting NTRIP broadcaster within Amazon Web Services (AWS).

This poster details how we utilised the AWS tools to produce a cloud based operational system within Geoscience Australia's regional IGS data centre.

Requirements

For Geoscience Australia's ground network to successfully support the SBAS testbed the new system had to meet the following requirements:

- Continuous provision of GPS (L1, L2 and L5) and Galileo (E1 and E5a) observations from 50+ ground stations across the Australasian region.
- Data streams need to be broadcast to the uplink station with minimum latency (< 6 seconds round trip).
- The system needs to be highly available with automatic failover and low time to recovery.

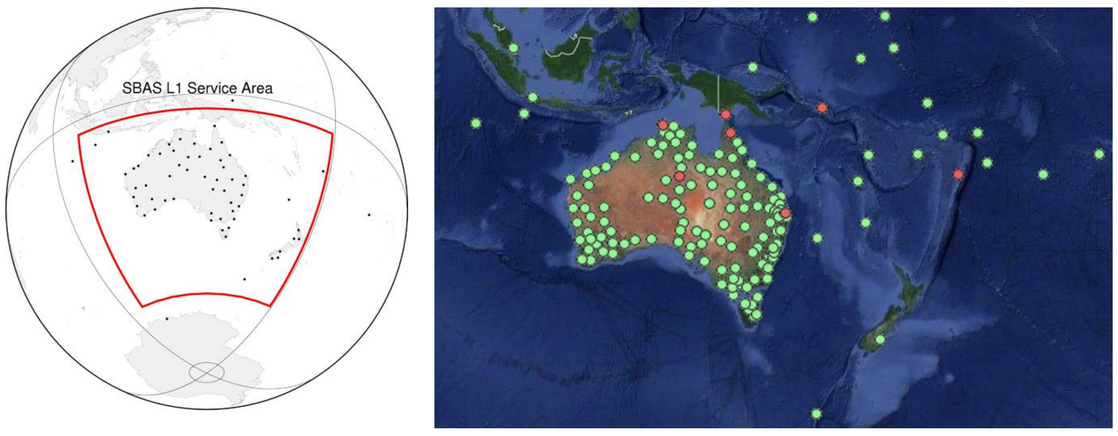

Figure 1: SBAS L1 service area and CORS stations that are being used to generate the corrections (left) and the stations across Australia that are available from the AUCORS NTRIP Broadcaster (right).

Design

To address the requirements we opted to utilise the tools available through Amazon Web Services (AWS). The design is based around the GNSS ground stations streaming RTCM 3.2 messages to two separate NTRIP casters (server stacks) which share a common configuration. Each caster operates independently and both are auto-recovering. Client connections are load balanced across both stacks, and in the event of an outage to one stack the other will take the full load. For the NTRIP caster we are using BKG Professional NTRIP Caster (2.0.26).

Compute

Elastic Compute (EC2) provides the compute capacity in the cloud. We are currently using T2.medium instances. Should our user base grow we can dynamically upgrade to a larger compute capacity.

In the event of an outage Auto Scaling Groups (ASG) provide a mechanism to restore the servers across multiple Availability Zones. In this case the ASG is set to min size 1, max size 1.

Elastic IP provides the server stacks with a static public IP address for use when DNS resolution or firewalls apply.

Elastic Load Balancer orchestrates the ASGs by running health checks against the servers.

Amazon Machine Image (AMI) provides a static image of the working server in a known state for faster recovery times.

Networking

Route 53 load balances traffic between the two server stacks.

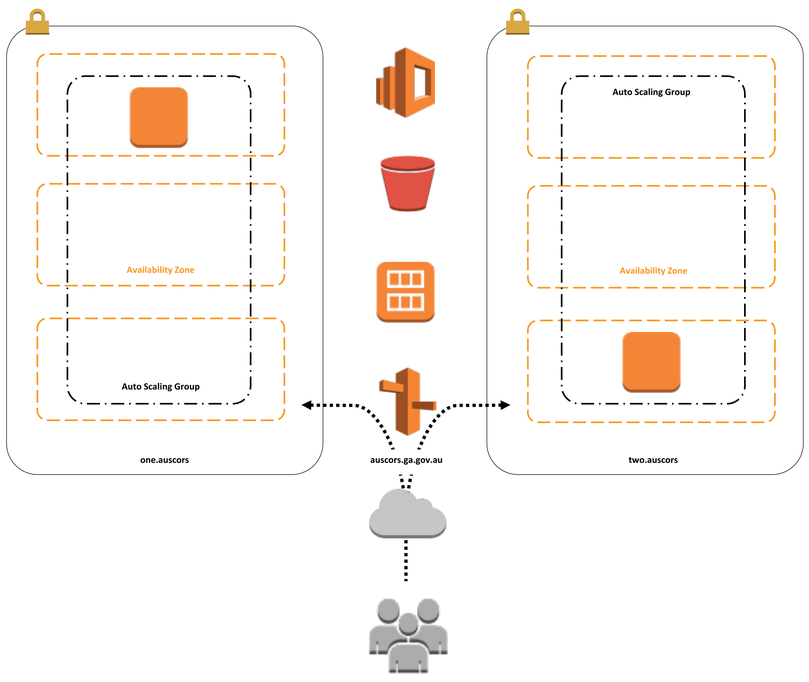

Figure 2: Architecture diagram of AUSCORS NTRIP Broadcaster. This does not include the user management system.

User Management

The AUSCORS user management system implements a serverless architecture for NTRIP user account creation. To do this, a registration webpage is hosted using S3 and Cloudfront. API Gateway and Lambda are used to implement the registration and user acceptance. The registration details are then stored in DynamoDB.

Analytics

Amazon's Elasticsearch Service makes it easy to ingest the NTRIP caster log files and perform analytics in near-real-time. Using this service we are able to monitor user connections, network bandwidth, station outages and data latencies. These metrics can be plotted or used to trigger alerts.

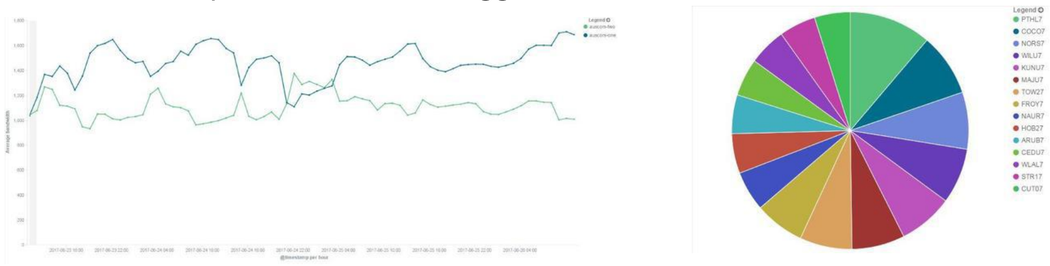

Figure 3: Analytics showing bandwidth per server stack (left) and the top 15 mount-points by bandwidth usage (right).

Limitations

- We have not yet been able to achieve scalability within our service. This is mainly due to limitations with the implementation of the NTRIP protocol and bandwidth limitations within EC2.
- Due to the high cost of data transfer from AWS there is some unpredictability in our monthly billing is spikes in client connections occur.

Summary and Next Steps

The move to AWS has provided our users with increased stability and reliability when connecting to our real-time data streams. In addition we are now able to make use of new analytics to further improve the service.

The new design has also allowed us to become a regional broadcaster in the Australasian region for the IGS and MGEX casters. This will help share the load from the global casters.

284	288	1.7 TB per month	7.5 TB per month
Streams	Subscribers	In	Out

Figure 4: Since deploying the service into AWS we have seen a significant increase in use. The above statistics show our usage for June 2017.

To register for an account or find out more information please visit https://www.auscors.ga.gov.au.

Acknowledgements

James Kingsmill and Tom Butler from Geoscience Australia's Cloud Enablement Team

For Further Information: Ryan Ruddick
Email: geodesy@ga.gov.au
Ph: +61 2 6249 9426 **Web:** www.ga.gov.au