

CloudStor

Getting out of the way of researchers

David Jericho, Solutions Architect, AARNet




Who is AARNet?

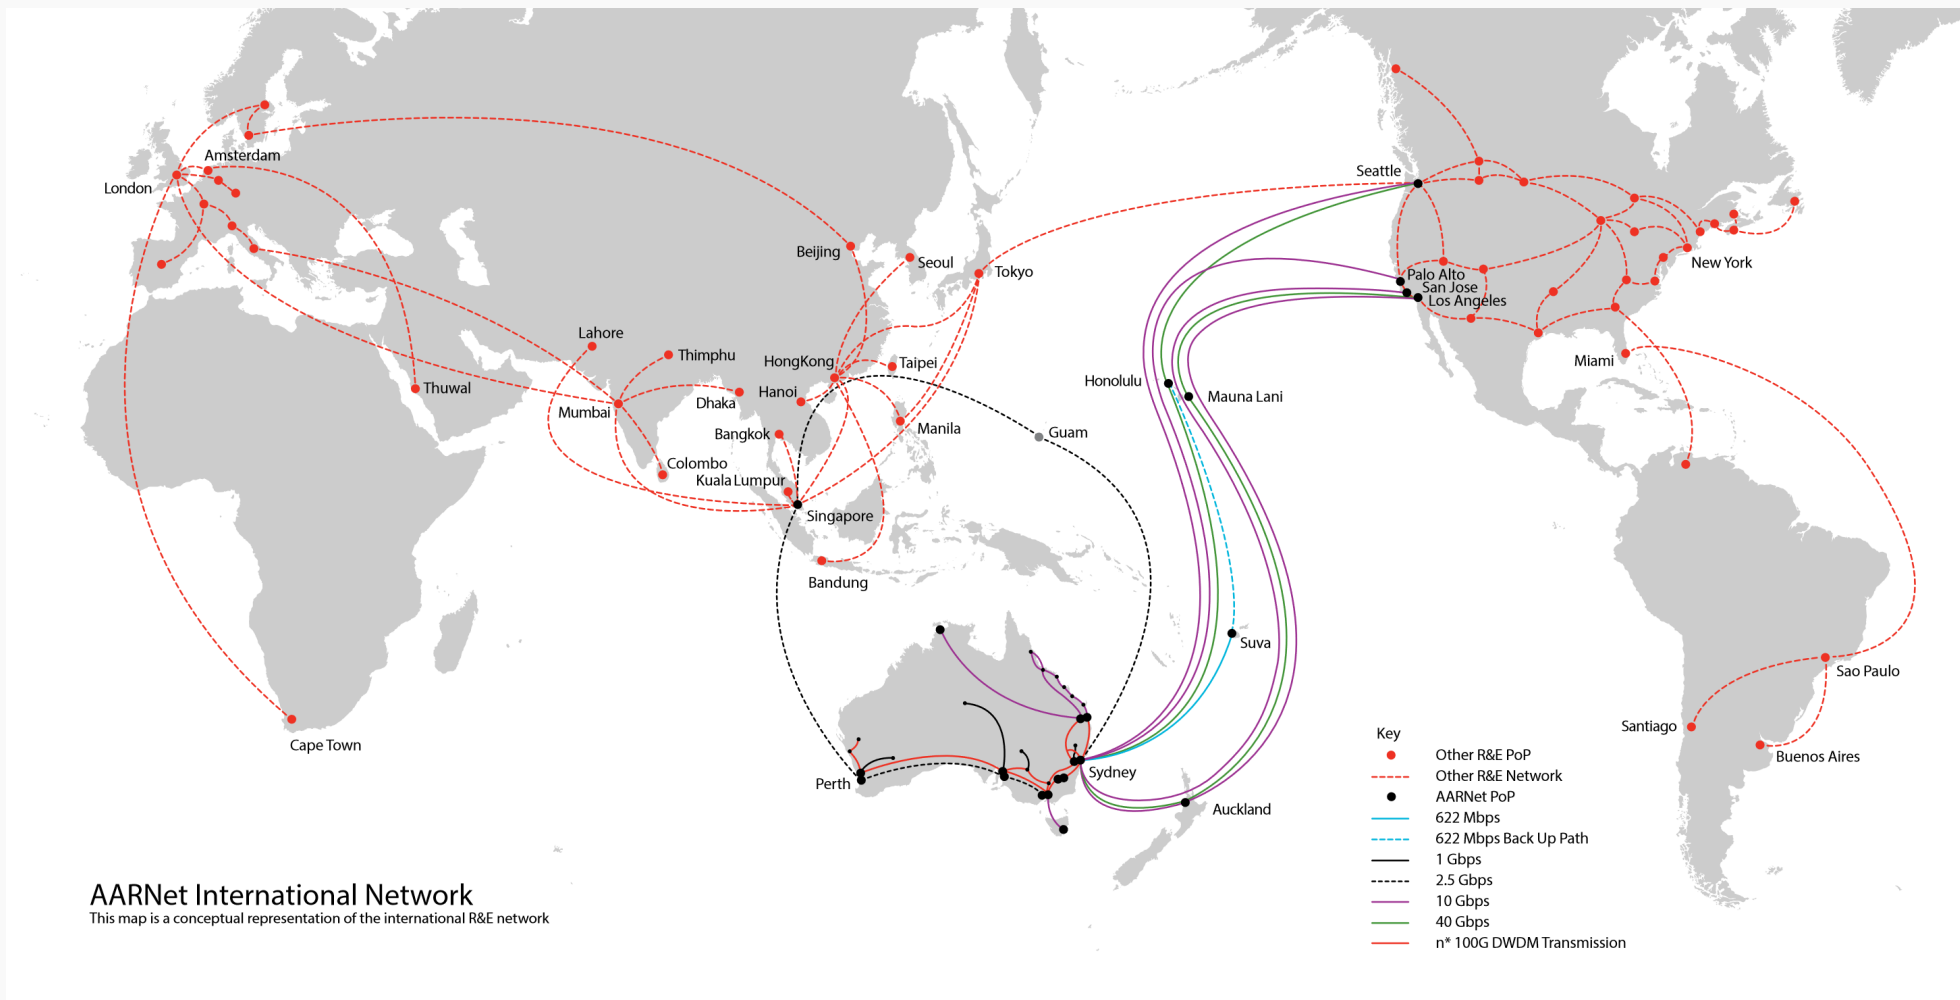
- Australian Academic and Research Network
- 38 universities and CSIRO
- Hundreds of
 - schools
 - galleries
 - libraries
 - museums
 - research centres
 - vocational training sites
- Own and run our own national fibre infrastructure



Australian Population Density



AARNet Global Reach




Australian Population Density Redux

- Geographic size of Brazil, with 10% of the population
- Can fit all of continental Europe inside Australia
- 90ms between extremes of the national network
- Australian researchers travel extensively
- New Zealand use our services too

AARNet's Sync, Share and Transfer



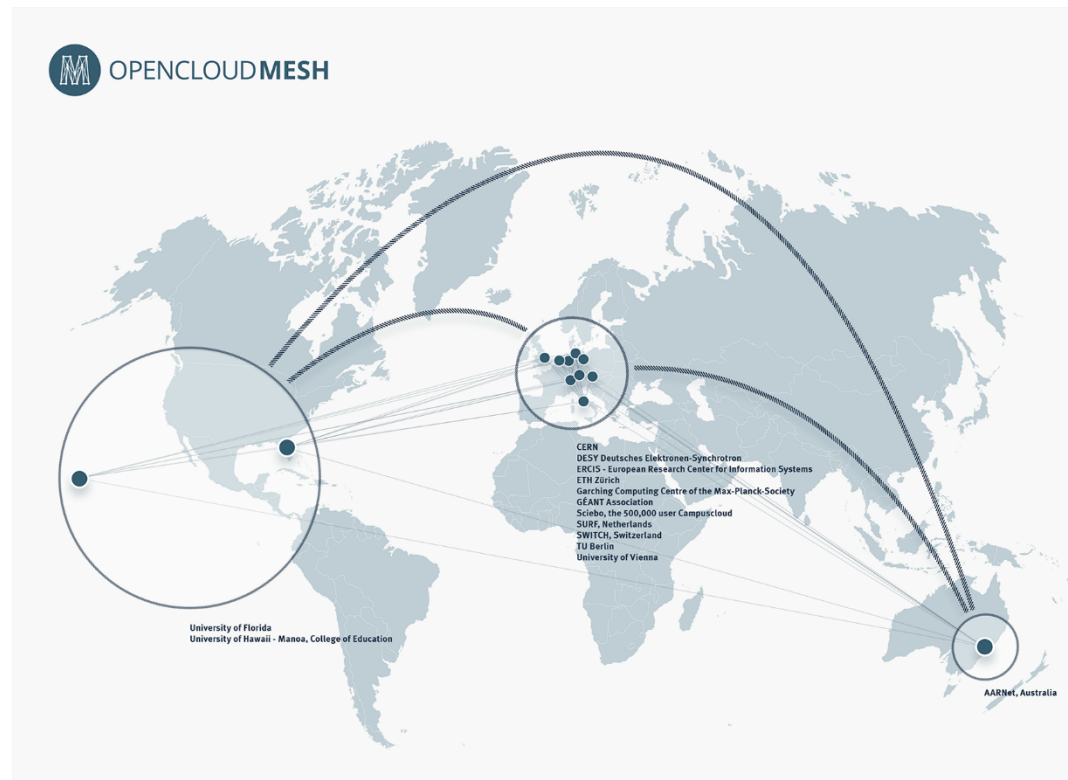
- Named CloudStor
- ownCloud 8.1 based
 - Also providing FileSender
- 17,086 researchers with 100GB quota as on 8th Feb 2016
 - 50% of Australian researchers via word of mouth
- 190 institutional identities present
- It's a concept, not an execution



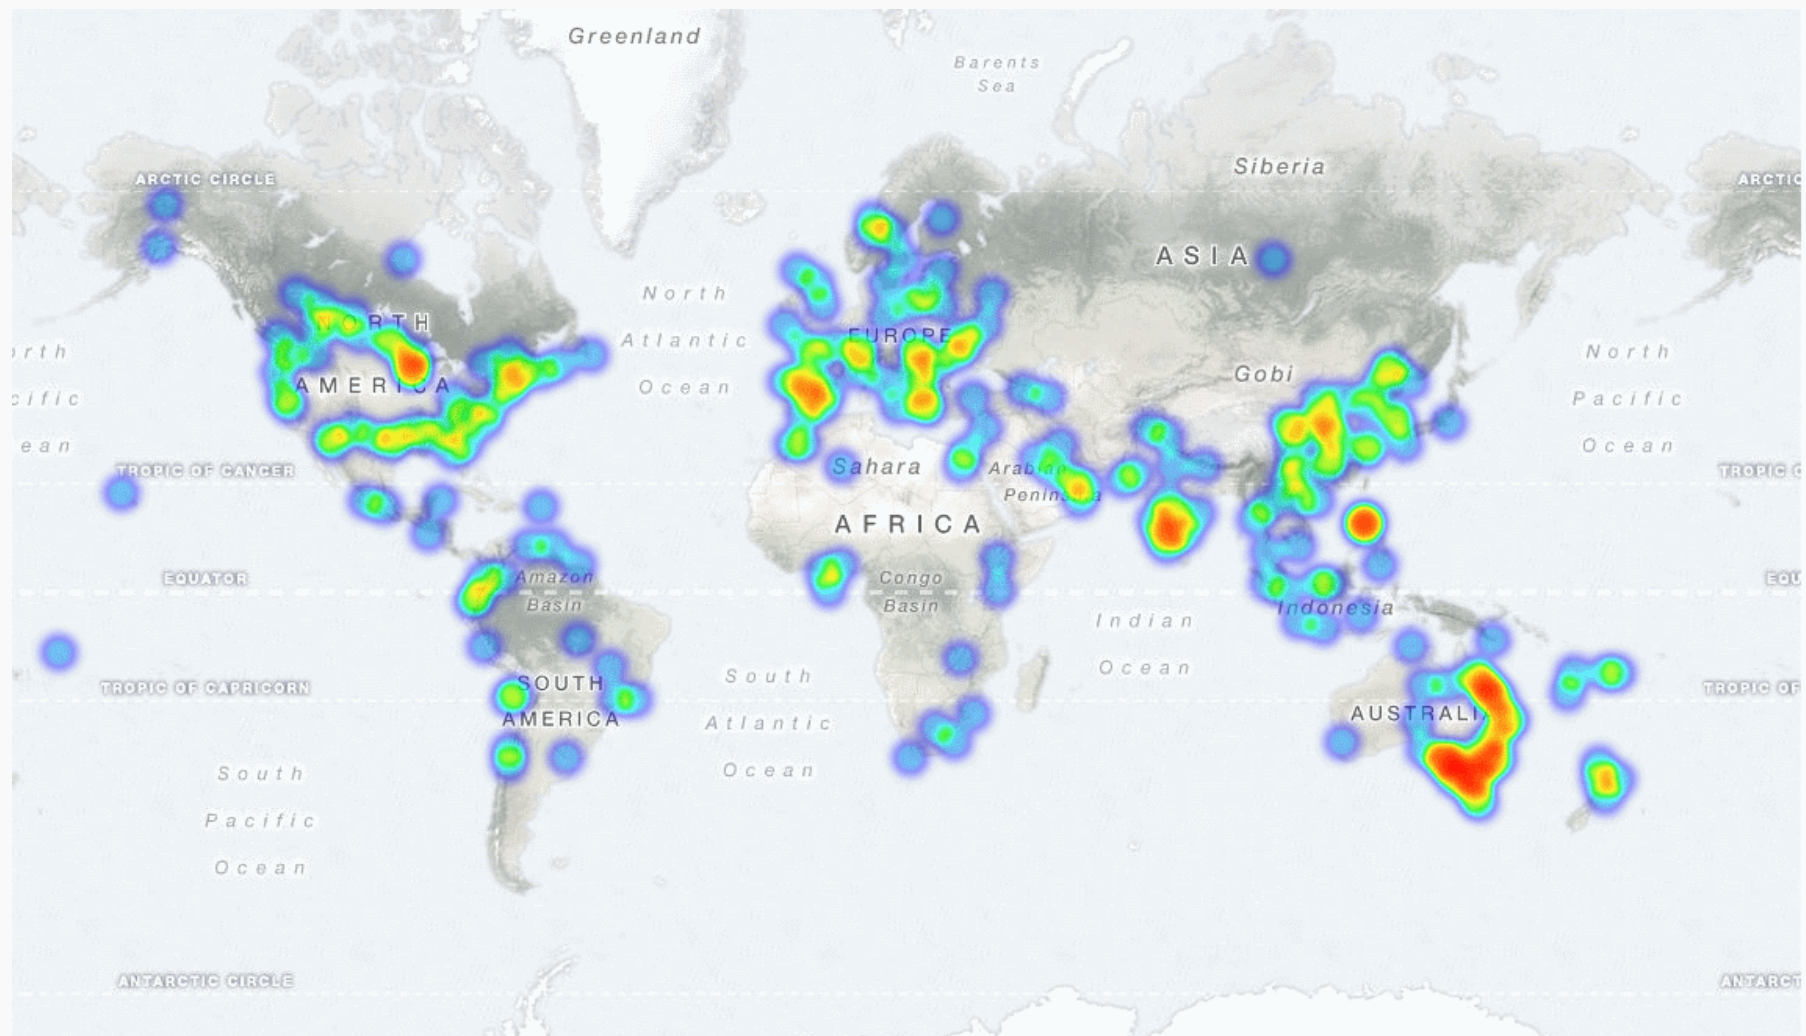
OpenCloudMesh



- Researchers share in the least expected ways
- Millions of R&E and students reachable
- Get on board, we're about including your system
- We do the heavy lifting because users can't



24 hours of CloudStor



That really is a bison genetics researcher in Mongolia

More on Why

- The DropBox problem
 - No control of the data
 - Sovereignty problems
- The CDN problem
 - Providing a way for people to distribute content
- The large data problem
 - Not everybody has big files
 - Most have a lot of small files that they care about
 - We do the heavy lifting



CloudStor Really Is...

- We position ourselves as:
 - The data pump
 - The ad-hoc data storage site
 - The middleware for researcher workflow

- We love
 - We actually celebrate moving data of all sizes
 - ... we won't block accounts because someone used it
 - Finding out what weird and wonderful research we're helping
 - Tooling things up to meet the needs of our users

Because We Like To Help

- Researchers shouldn't have to be IT experts
 - And they shouldn't have to be
- Do what you're good at
 - Let us do what we're good at

The Design Problem

- Highly IT focused people looking in
 - We don't know the disciplines
- We need to be told
 - Which brings us back to the things we love

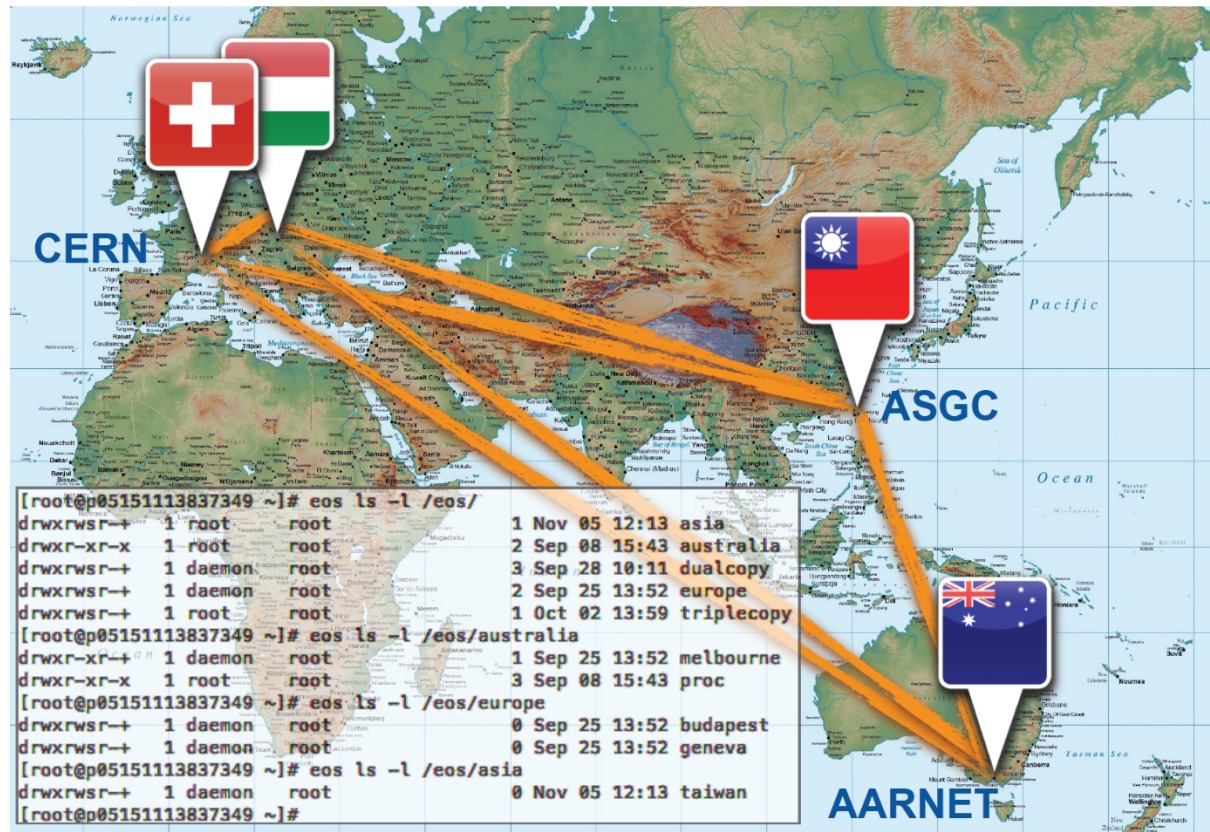
CERN's EOS in Australia For CloudStor



- True high performance single namespace at 65ms
- Data in and out per client at tens of Gbps
- Most scalable solution found after failing with many
 - We keep talking about a paper on failing storage on WANs
- ...and it offers a filesystem for old tools to interact with
- ...and it scales to hundreds of petabytes readily

Ongoing Experiments with CERN

R&D - EOS World-Wide Deployment





Deployments by Quarter 2

- Beta of integration with VicNode's OpenStack volumes
 - Tests have shown 36Gbps ingest at 2,000 km
- AARNET branded client
 - Free on Apple App Store and Google Play
- More storage
 - Ensuring we have both the capacity and IO to cope
 - Tripled user count in less than a year
 - And the requests keep coming

The 2016 Roadmap

- CloudStor under the desk
 - Bringing data extremely close to the user
- CloudStor, er, in the Cloud
 - Containerised user space on OpenStack VM
 - Choose your own file transfer protocol
- CloudStor replicas in other countries
 - True replication, by policy to remote sites
- Institutional administration and control

The 2016 Roadmap 2

- Institutional publishing and archive integrations
 - Integrating ad-hoc use with archives and publishing
- Enabling additional protocols
 - Because this is about getting out of the way of the workflow
- Adding application integrations
 - Enabling users to flow their data between their instruments, the cloud, their supercomputers, and their laptop
- OpenCloudMesh
 - Heterogeneous cloud global sync, sharing and transfer

The 2016 Roadmap Beyond Thunderdome



- Service level agreements
 - For Site Reliability Engineering metrics
 - For customer comfort and assurances
- Compliance testing and certification
 - Government data
 - Medical data

CloudStor: Getting Out of the Way



- Easier to connect
- Choose your own protocol
- Institutional control
- Sharing globally without tools
- On-demand global replication
- Making institutions happy