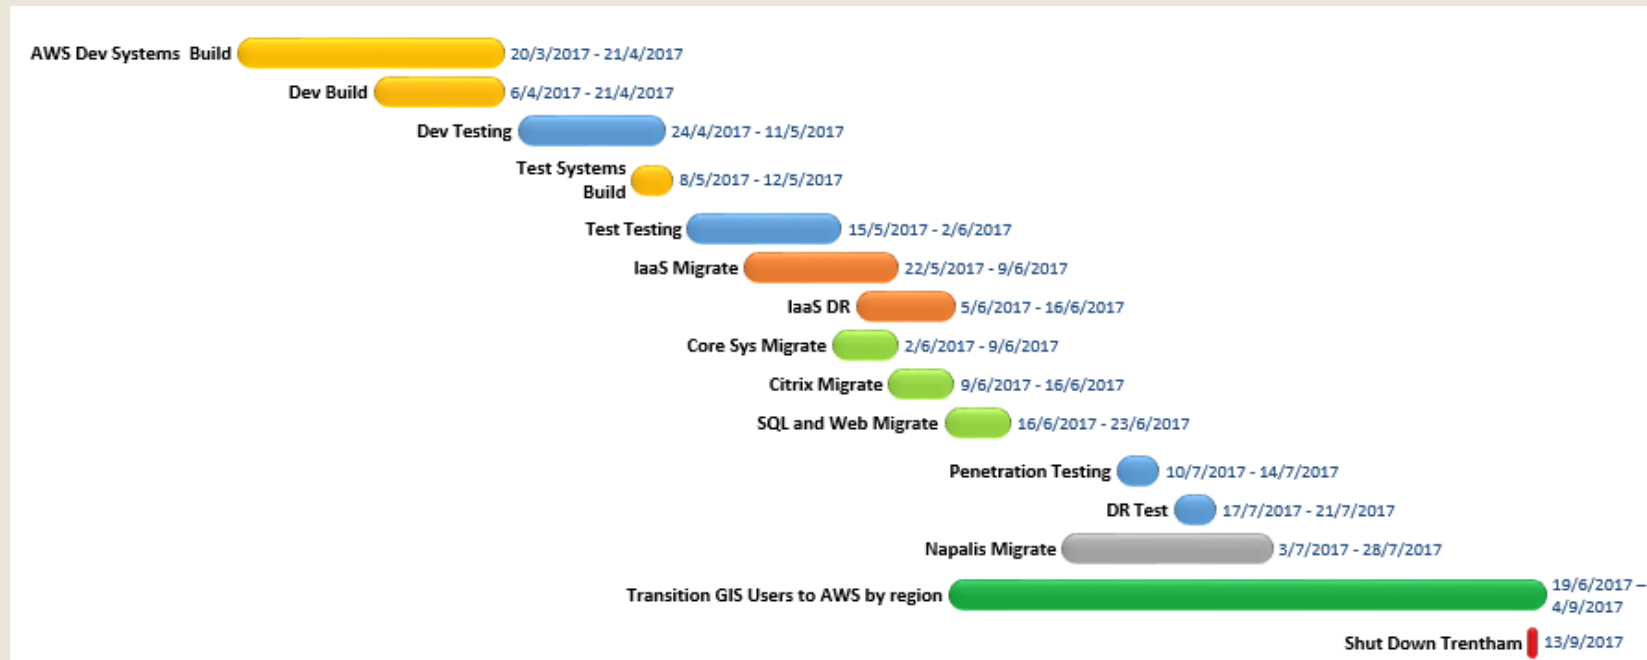




Moving GIS to the cloud: *The process, challenges and lessons learnt*

The journey begins: April 2016

- Migrate DOC's computing storage and infrastructure to the cloud
- Would involve 155 terabytes of data including:
 - 22 terabytes of GIS data (6,417,078 files in 305,700 folders)
 - 69 GIS drives
- GIS migration the final step.



Reason for change?

- Data storage services old and at risk of failing

Why the cloud?

- Zero infrastructure costs
- Flexible billing of resources
- High-availability and disaster-recovery
- Ensure security of data and information
- Scalability
- No single point of failure

Why Amazon Web Services?

Global Footprint

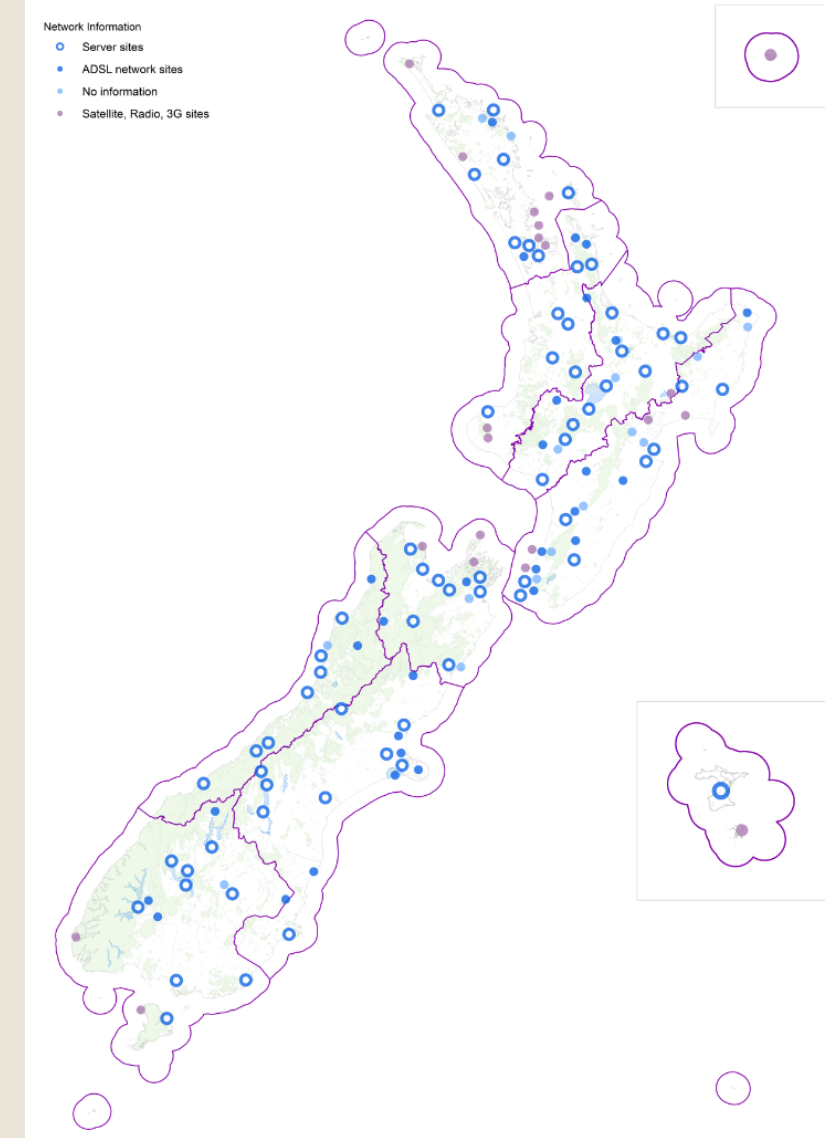
Everyday, AWS adds enough new server capacity to support Amazon.com when it was a \$7 billion global enterprise.

- Over 1 million active customers across 190 countries
- 900+ government agencies
- 3,400+ educational institutions
- 11 Regions
- 28 Availability Zones
- 53 Edge Locations



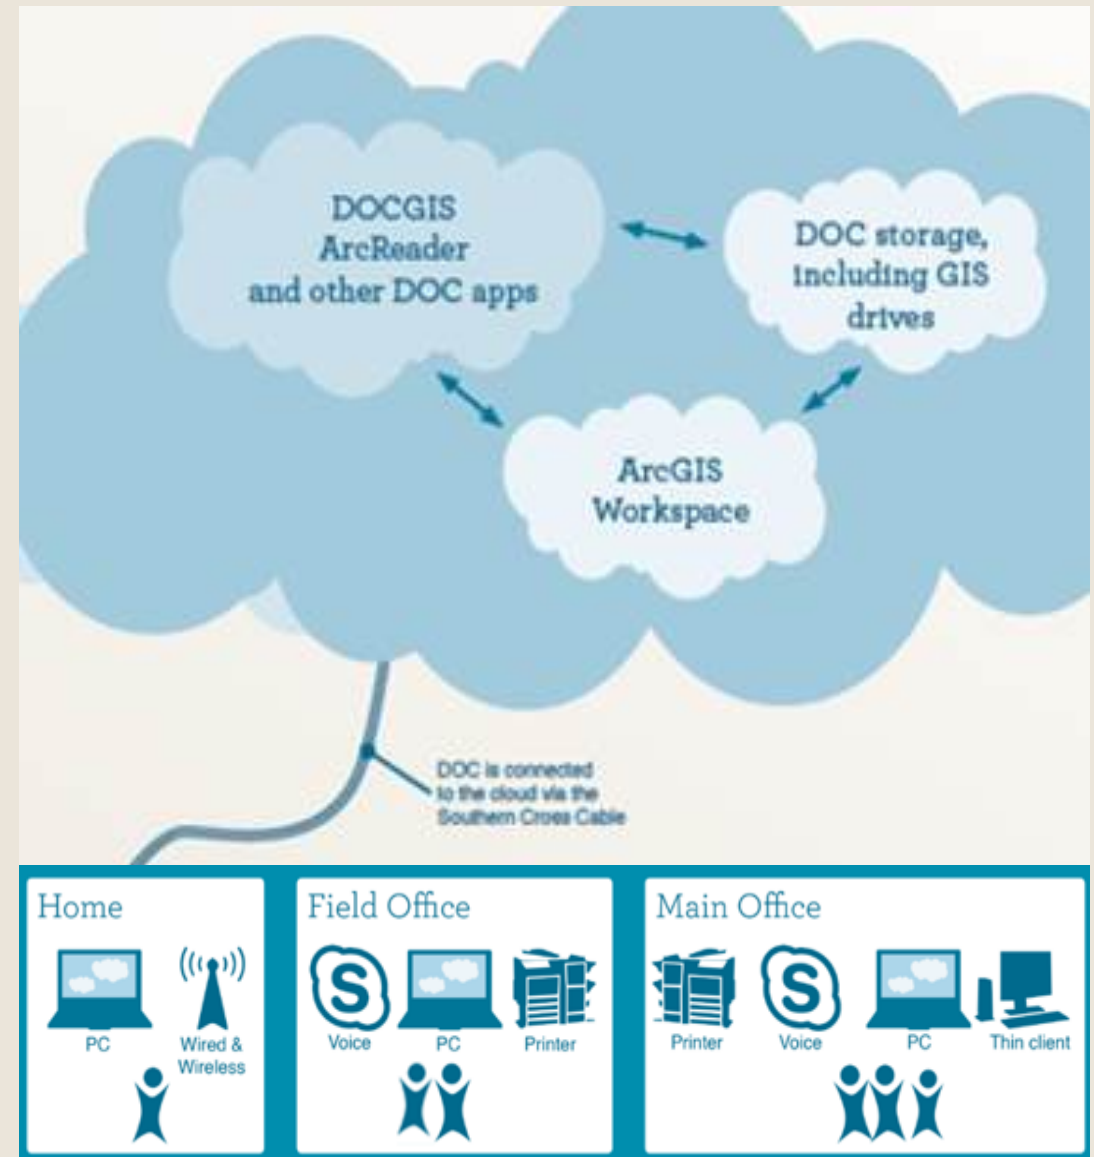
Existing GIS structure

- De-centralised GIS infrastructure
 - 31 staff in 11 offices
 - 69 servers
 - 500 ArcGIS installations
 - 300 active GIS users
 - 3 offices not moving -
 - Aniwa Niwa
 - The Chathams
 - Pureora



Vision for GIS

- Virtual workspaces
- Centralised storage
- Single source of data
- Shared to multiple users
- Structured directory





Preparing for the move



Proof of concept: Apr – June 2016

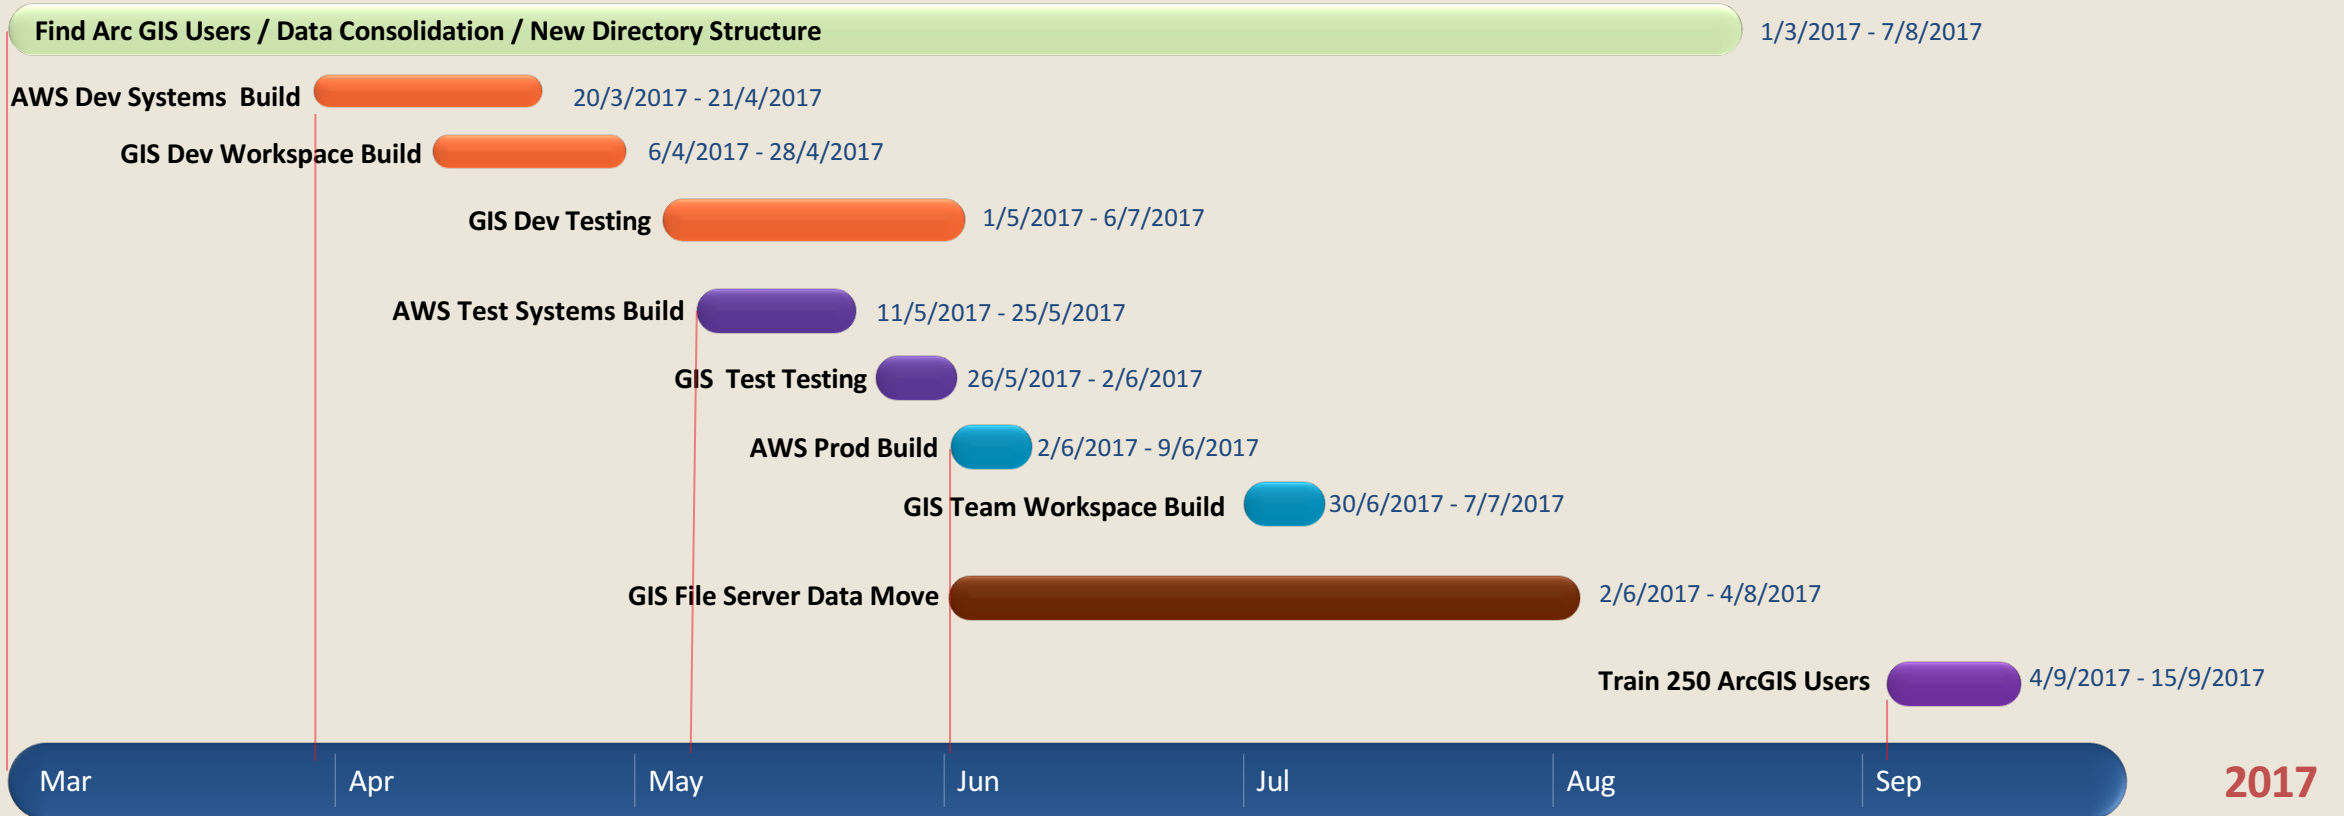
Categories	# of Test Cases	Pass ✓	Fail x
Cartography	5	5 ✓	0 x
Offline	8	7 ✓	1 x
3D	2	2 ✓	0 x
Data Transfer	5	4 ✓	1 x
Data Capture / Editing / Loading	1	1 ✓	0 x
Database Management	5	5 ✓	0 x
Data Analysis	3	3 ✓	0 x
App dev	2	2 ✓	0 x
DOC Apps/Tools	1	1 ✓	0 x



Staff feedback

Location	Comments
Invercargill	<ul style="list-style-type: none">• Very good, fastest time I've ever seen ArcMap open
Christchurch	<ul style="list-style-type: none">• It's noticeable at least 5 times faster if not more when it comes to loading data and programs.
Whangarei	<ul style="list-style-type: none">• All the tools I used worked perfectly and quickly.
Motueka	<ul style="list-style-type: none">• Normally takes 30 seconds to a minute to load a blank MXD in <u>Motueka</u>. It took one second to load in AWS.• Took 1/10 of time to connect to NATIS .• These speeds will have huge positive impact on analysis work.
Whangarei	<ul style="list-style-type: none">• AWS is substantially faster.• <u>ArcReader</u> project opens and runs well.

Gradual process



Start

4/9/2017
AWS Go Live
4/9/2017
▶ G Drive, S/GIS Drive Read Only

Consolidate data

Situation:

- 22 terabytes of data on G:\drive and 4 terabytes on the S:\GIS drive
- Duplicates of images and national data sets, contour lines.

Action required:

- Establish data review group
- Evaluated commercial geospatial search tool
- Python script to find duplicates
- Delete duplicates

New Directory Structure

The central GIS Drive



Q

National datasets



Data

(All data for GIS and Science)



GIS_Analysts



GIS_Development



GIS_Science



GIS_Users



Imagery



Logos_Styles_Symbology_Templates



Old



Secure



Shared

Science users folder



General users folder



All local data and projects


Former G:\ and S:\GIS content



For confidential work. Only seen by those with permission to folders



For confidential work. Only seen by those with permission to folders



Projects and working data

Projects to be used across the department e.g. ArcReader



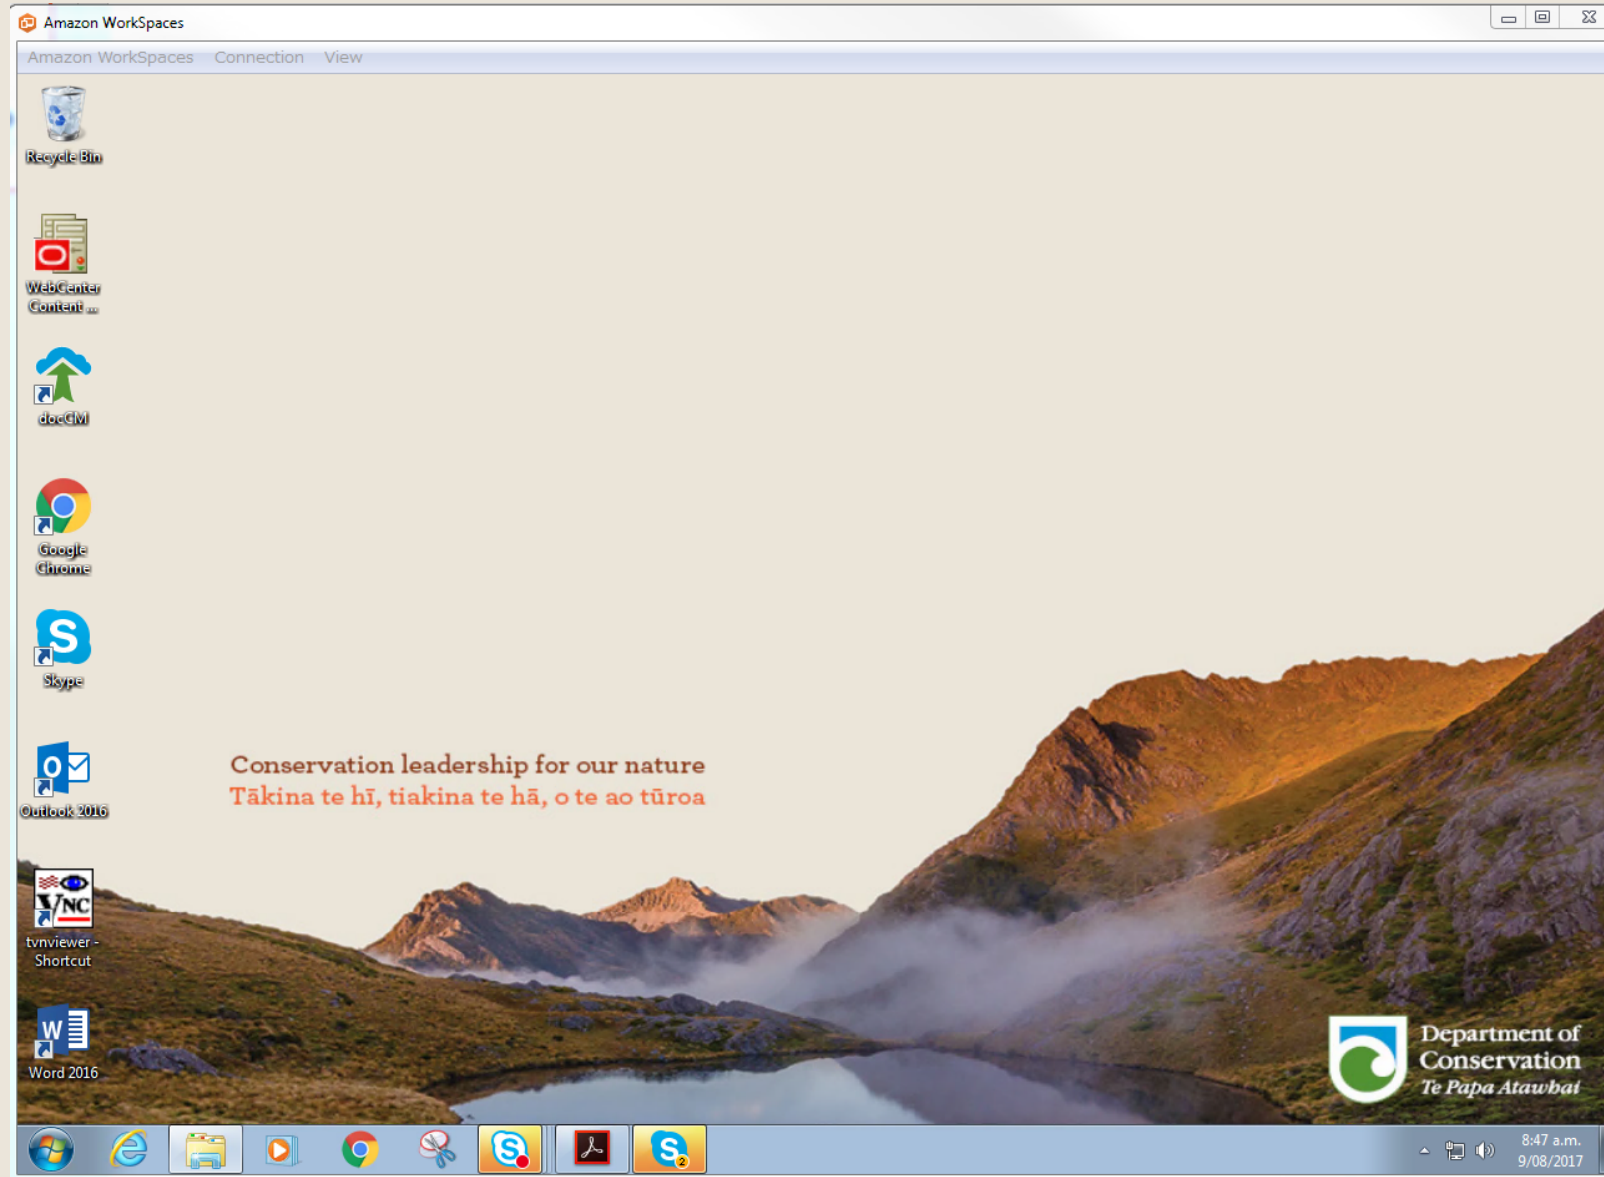
Department of Conservation
Te Papa Atawhai

AWS Workspace Types











	Value	Standard	Performance	Power	Graphics
vCPUs	1	2	2	4	8
Memory GiB	2	4	7.5	16	15
vGPUs					1
Video Memory GiB					4
SSD Root Volume GB	80	80	80	175	100
SSD User Storage GB	10	50	100	100	100
Software	Utilities software bundle	Utilities software bundle	Utilities software bundle	Utilities software bundle	Utilities software bundle



AWS Workspace



Sample Test Case

Tasks	Completed on AWS	Performance
Open Manawatu.A0.mxd on AWS		1.55 seconds, 2.01 seconds, 2.03 seconds, Average = 1.86 seconds
Open Manawatu.A0.mxd on local PC		1.43, 1.32, 1.39, Average = 1.38 Seconds
Add in boundaries from NATIS 1		Slow to open initially (30seconds) but after initial opening then instantaneous
Add in Topo250 from NATIS3		Slow to open initially (30seconds) but after initial opening then instantaneous
Add in shape files from local AWS project folder(P:\...)		DOC symbols are available in AWS
Add legend		
Add DOC logo		No DOC logo available , not copied over from G drive
Clip to boundary function		
Export function to PDF at 300dpi		Took 16.29 minutes. Topo250 didn't export properly, lost the cache images. Need to create cached images in AC when image data copied onto AWS. Couldn't view map in Adobe as not installed
Sent PDF to A0 plotter in Wellington		Had to install printer drivers on workstation and print from there

User experience

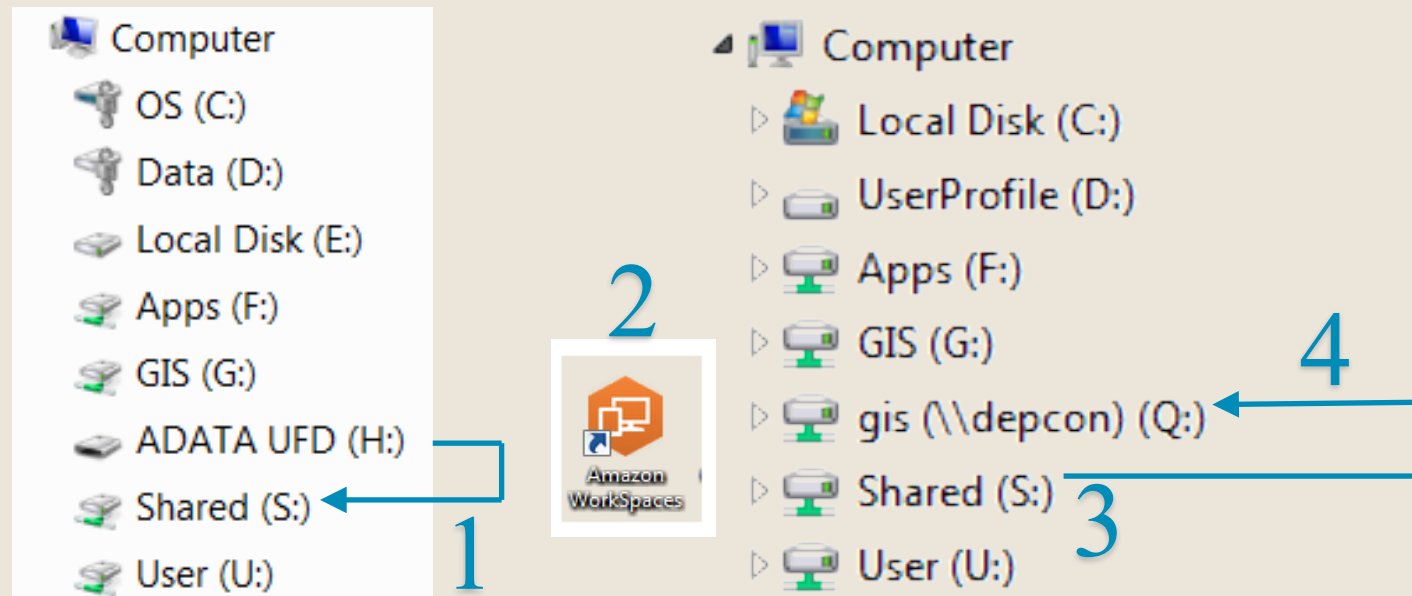
Evidence	User Narrative	Next Steps
Feedback provided from Analysts/scientists	Straight forward to use. Response speed overall was good. Only issues, raster caches not coming through and printer not available	<ul style="list-style-type: none">• Make sure to either copy over raster caches or rebuild in AWS.• Need to install (AO) printer driver to get access to all print features.
Progress against planned tasks	Worked as expected. Not much difference in response, slightly slower to open MXDs in AWS. Manawatu_AO(1) and Wairarapa_AO were copied on both local and AWS machines and all data sources fixed. Wairarapa tool 1.49 minutes to open in AWS and 1.20 on the local machine.	Trevor Hart from Eagle to investigate
Skill gap	Had all skills needed.	
Ease of deployment	Good	
Capabilities required from development partners/vendors	Needed ISS to install AO printer drive and Adobe Acrobat Reader on AWS	<ul style="list-style-type: none">• Install Adobe Acrobat Reader on all AWS workstations and desktops.

ArcGIS Testing Outcomes

- 88 test cases
- 18 testers(GIS analysts/developers)
- 10 categories
- 154 page document
- USB not supported
- ArcGIS Pro failed on Standard, Performance and Power Workspaces

Categories	# of Test Cases	Pass ✓	Fail x
Cartography	23	21 ✓	2 x
Offline / GPS	12	11 ✓	1 x
3D	5	5 ✓	0
Data Transfer	7	7 ✓	0
Data Capture / Editing / Loading	3	3 ✓	0
Database Management	5	5 ✓	0
Data Analysis	6	5 ✓	0
App dev	9	9 ✓	0
DOC Apps/Tools	11	11 ✓	0
Access	7	7 ✓	0
	88	85	3

Workarounds – Moving Data from USB to AWS



Consequences of move

- Broken MXD links
- Republishing ArcReader projects
- Access ArcReader projects through Citrix on AWS

Training

- Over 250 staff in 9 regions
- Train the trainer sessions
- Running 23 classroom sessions
- 13 trainers
- 2 week time frame
- ½ day courses



Training Materials

Quick User Guide: ArcGIS in AWS

Logging on

1. Go to your start menu
2. Click on Amazon WorkSpaces
The first time only, a registration screen will appear.
3. Enter DOC registration wssyd+FFF8F7
4. Click "Register"
5. Standard log in window appears
6. Enter your DOC username and windows password.
7. Click 'Yes' to remember me question. You only need to register once, thereafter the standard log in screen appears.

Data migration best practice

- Work in Q drive from now on
- Copy data from Q:\OLD to new Q drive structure
- Name files appropriately
- Rename based on data management best practice

How to fix a data source in an MXD

1. Click ! against the broken link to open "Set Data Source"
2. Navigate to the new location

Bulk update of data sources in MXD

1. Open ArcCatalog
2. Right click on the MXD
3. Click 'Set Data Source'
4. Click 'Replace' or 'Replace all'
5. Specify the connection path you want to replace
6. Enter the new location

Fix MXD

1. Copy mxd to the Q:\GIS_Users\Projects folder
2. Open and identify broken data links
3. Copy and rename relevant data
4. Check new data source has metadata attached to it.
5. Reprint mxd to new data location(s)
6. Repeat for each layer.

Data management best practice

- Meaningful file names – location, subject, ID number, species, date, your initials.
 - ✓ Dates: YYYYMMDD
 - ✓ No gaps in names – use underscores
 - ✓ Limit file length to 160 characters
- Include metadata (who, what, where, why, when).
- Where possible combine data
- Avoid duplicating (check if data already copied or under different name)

Move data from USB/local PC to AWS

1. Copy data from USB stick to local S:\Transfer
2. Open Windows Explorer in the AWS workspace
3. Navigate to S:\Transfer
4. Copy the data to Q: Drive

Map Toaster in AWS

1. Plug GPS into your computer
2. Go to S:\GIS_Mapping
3. Double click 'Map Toaster' icon
4. Click 'MyPlaces'
5. Go to GPS Config tab
6. Click 'AutoDetect'
7. To download data click 'From GPS'
8. Symbology for waypoints and tracks can be changed in 'My Places'
9. 'Save'
10. Navigate to S:\Transfer, 'Save'
11. Open Windows Explorer in the AWS workspace
12. Navigate to S:\Transfer
13. Copy the data to Q: Drive



ArcGIS in the cloud (Amazon Web Services): Training for ArcGIS users

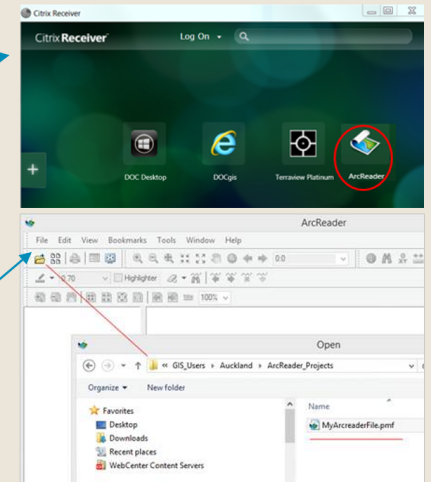
www.doc.govt.nz

Department of Conservation
Te Papa Atawhai

Quick User Guide: ArcReader in Citrix

ArcReader is now only accessible via AWS. If you don't have access to AWS, follow these steps:

1. Launch Citrix
2. Click on the ArcReader icon
3. Click "Open File" to show your drives, then open the project.
4. Navigate to Q:\GIS_Users\{Region}\ArcReader_Projects to show your drives.
5. Then open the .pmf file



- Detailed guide for ArcGIS users
- Facilitation guide to ensure consistent approach

Projects that have been updated will be found in the new Q:
Projects that still need to be updated and migrated will be found in Q:OLD



Challenges: Environment and People



AWS Environment

AWS environment provided many challenges:

- Webcams not supported on AWS
- No support 4K resolution
- Skype for Business limitations
- YouTube resolution and audio issues

Project / People

- Testers were busy with own work
- Data migration delays
- DOC network speed
- PM, BA, GIS analysts & data review team lead on holiday
- GIS Technical lead resigned
- Lack of communication
- Breaking new ground – steep learning curve



Lessons learnt



Lessons learnt

- Pick IT service partner (external) based on cloud experience
- Learn AWS terminology
- Establish core project team early
- Release project team from BAU
- Beware of move consequences
- Dedicated communications resource
- Don't leave too long between pilot project and implementation
- Beware of costings – egress costs
- Always takes longer than you think



What we've achieved



Success?

- Increased performance in DOC in general
- GIS perspective
 - Virtual workspaces
 - Centralised storage
 - Single source of data
 - Shared to multiple users
 - Spring clean of data
 - Performance increase for regions
- Uptake - will know more after training
- Still on the journey – watch this space...



*I wonder which one has
my GIS data in it?*

Questions?

