

# AFS at USGS

European AFS and Kerberos Conference

March 26, 2014

Ken Dreyer <[kdreyer@usgs.gov](mailto:kdreyer@usgs.gov)>

# What is the USGS organization like?

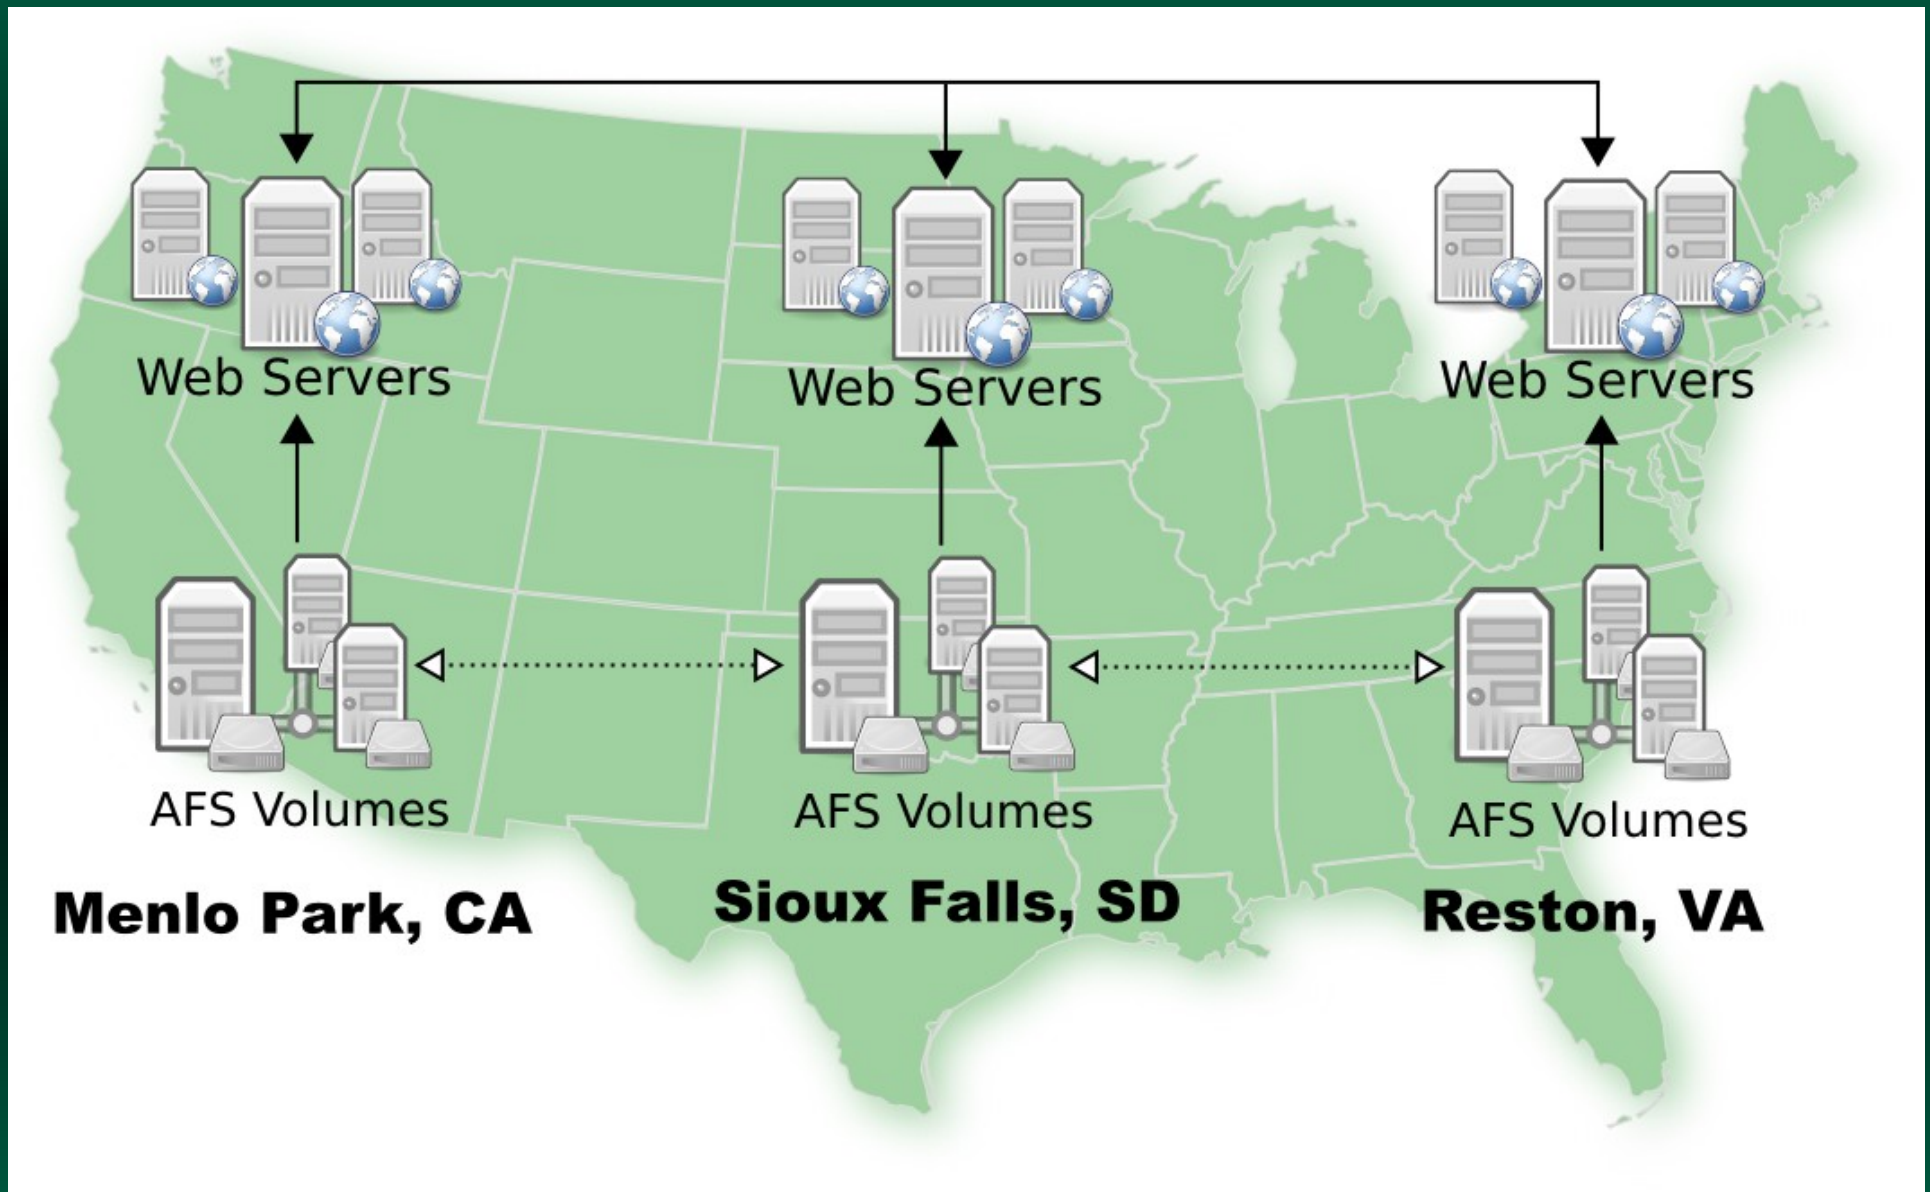
- US Geological Survey: Federal Agency of the US Government
- Research based, fact finding mission, non-regulatory
- Centered around earth sciences

# How does USGS use AFS?

- Website hosting
  - Homebrew automated replication
  - Regionally distributed Web servers: Linux+Apache, and Windows+IIS
- Linux and Solaris packages
- Science research project data storage
- Homedirs for a slim number of users

# What is the size and architecture of the usgs.gov cell?

- Cell is thirteen years old
- Three geographic regions
- Multi-realm authentication (AD and Heimdal)
- ~600 users



# Recent Improvements

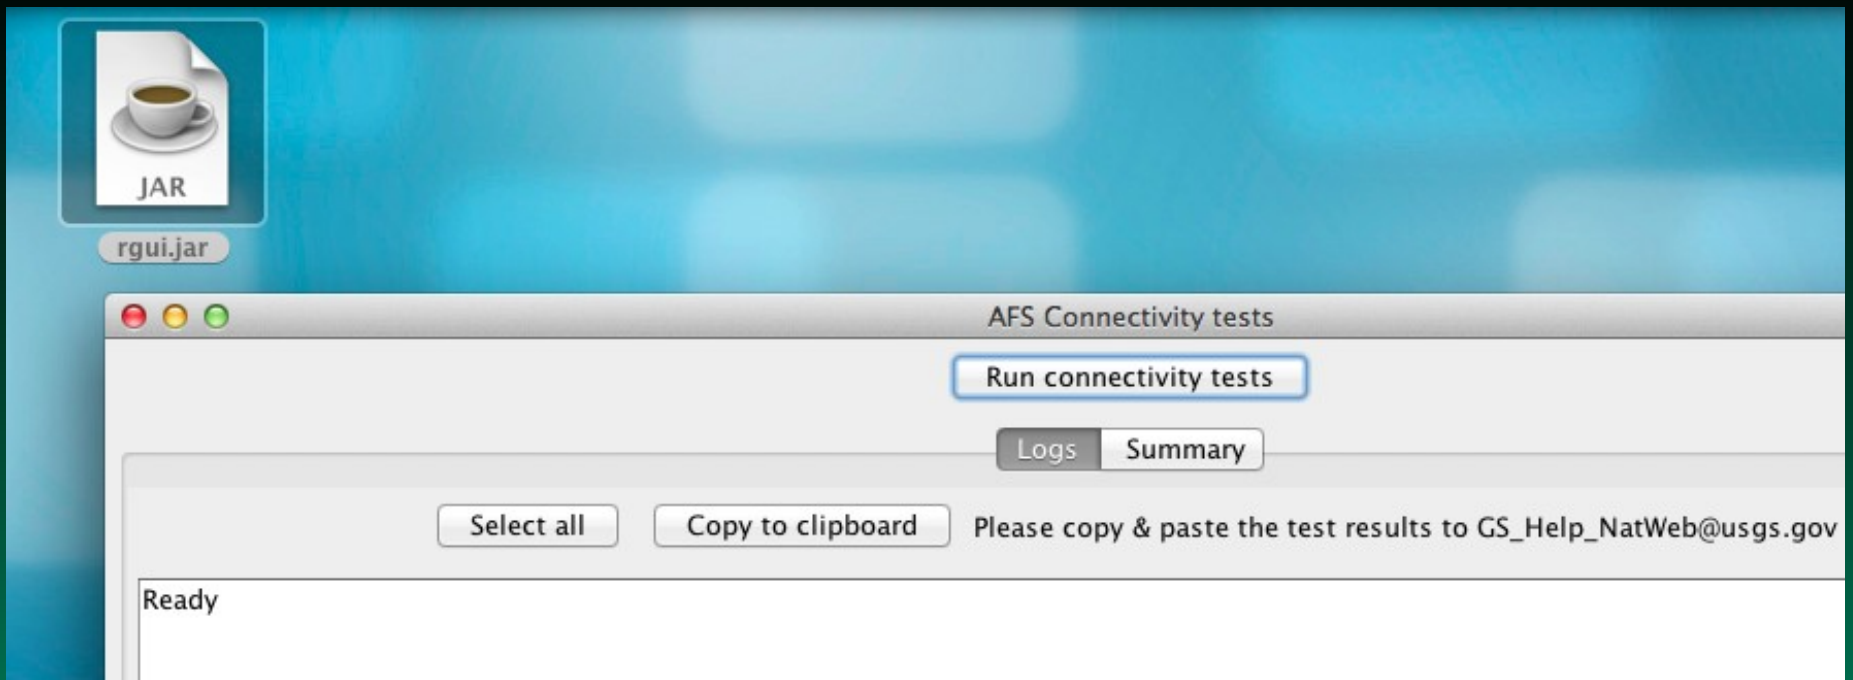
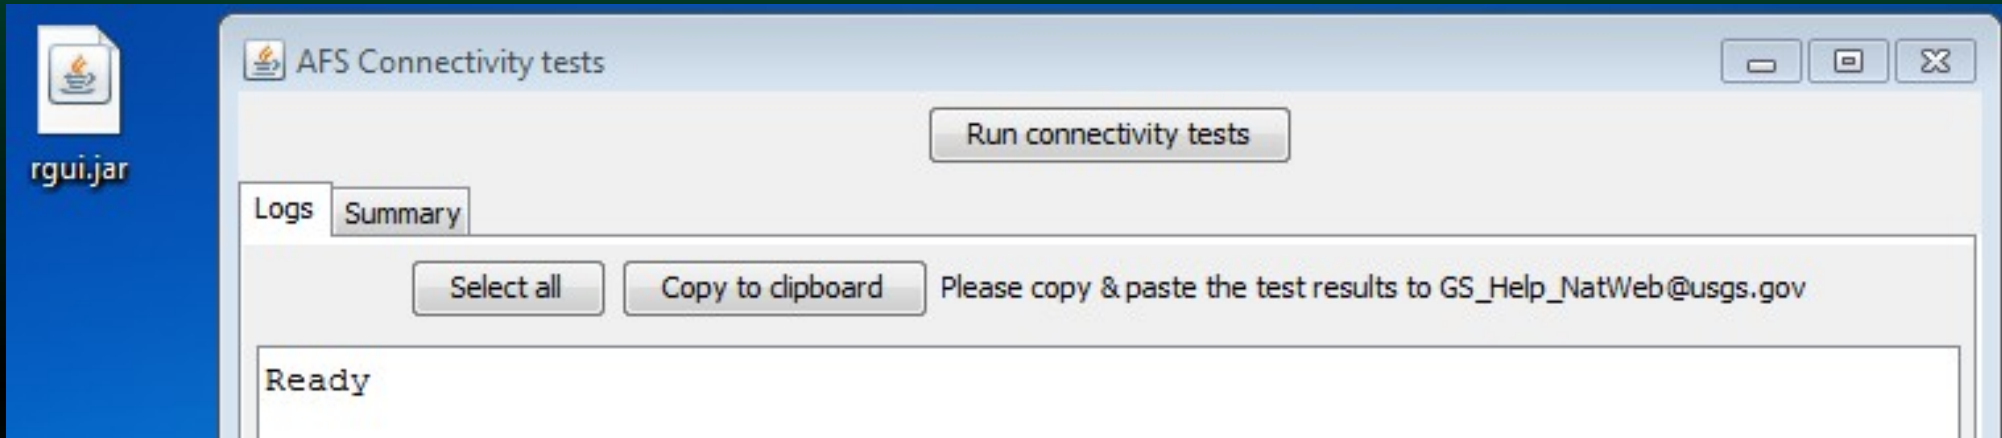
## 2012:

- Moved our web servers (main AFS clients) from Solaris to Linux VMs.
- Introduce Puppet for web servers. "AFS client" in Puppet.
- Ship custom Windows MSIs for OpenAFS to our users - drastically reduced the support load on our team for our Windows users.

## 2013:

- Transitioned AFS file servers from Solaris+Fibre Channel to Linux VMs+iSCSI.
- AFS file servers in Puppet.
- Start shipping Windows OpenAFS to users via Microsoft SCCM.

# GUI for rxdebug



# Recent Improvements (continued)

2014:

- Delegated volume releases via remctl
- Moving our Heimdal KDCs from Solaris zones to Linux VMs.
- Heimdal in Puppet.



# Why does USGS use AFS in 2014

- Mature multi-platform support (Linux, Solaris, Mac, Windows).
- Transparent volume moves / unified namespace.
- Low barrier to entry for web developers.
- Good uptime for read-only content.
- Transition costs.

# Things that are difficult about AFS

- ~~KDCs require single DES encryption (without rxkad-k5)~~
- RW volumes are single points of failure
- Lack of IPv6 support
- AFS errors are hard for users to troubleshoot

# (could be better... continued)

- UDP is not always treated well on our WAN - data moves very slowly between data centers
- Interrupted releases require full release to fix
- Limited threads for vos release on servers
- VLDB contention
- First two releases being performed serially is not optimal for us, and `vos release -stayonline` had problems
- No “progress meter” for `vos release` (would be especially helpful for user delegated vos releases)

# Things that we would love to do in the future

- Better staging environments
- Make more operations self-service (like account creation, or volume management)
- AFS with AWS (Amazon Web Services)