

# AFS and Linux Containers

*Wednesday, May 16, 2018 11:30 AM (20 minutes)*

One future model of software deployment and configuration is containerization.

AFS has been used for software distribution for many decades. Its global file namespace, the @sys path component substitution macro which permits file paths to be platform-agnostic, and the atomic publication model ("vos release") have proven to be critical components of successful software distribution systems that scale to hundreds of thousands of systems and have survived multiple OS and processor architecture changes.

The AuriStorFS security model consisting of combined-identity authentication, multi-factor authorization, and mandatory security policies permits a global name space to be shared between internal, dmz and cloud; and to store a mix of open and restricted data.

The combination of Linux Containers, the global AFS namespace, and the AuriStorFS security model is powerful permitting the development of container based software deployments that can safely bridge internal, dmz and cloud with reduced risk of data leaks.

This session will discuss the most recent updates to AuriStorFS and the Linux kernel implementation of AF\_RXRPC socket family and (k)AFS filesystem. A demonstration will be included consisting of:

- Containers with binary executable files stored in /afs
- Containers mounting private AFS volume for scratch space
- AuriStorFS and (k)AFS file system implementations running side-by-side
- Linux namespaces for /afs

AuriStorFS milestones since HEPiX Spring 2017 include:

1. Successful migration and replication of volumes exceeding 5.5TB. The largest production volume so far is 50TB with a 250TB volume
2. Deployment of a single AuriStorFS cell spanning an internal data center, AWS and GCP with more than 25,000 nodes for distribution of software and configuration data.
3. Meltdown and Spectre remediation. In response to nearly 30% performance hit from Meltdown and Spectre the AuriStor team optimized the Rx stack, Ubik database and fileserver to reduce the number of syscalls by more than 50%
4. AES-NI, SSSE3, AVX and AVX2 Intel processor optimization of AES256-CTS-HMAC-SHA1-96 cryptographic operations for kernel cache managers reduces computation time by 64%

AF\_RXRPC and kAFS highlights:

- IPv6 support for AuriStorFS
- dynamic root mount -o dyn
- @sys and @cell support
- multipage read and write support
- local hero directory caching
- per file acls (for AuriStorFS)
- server failover and busy volume retries

## Desired length

20 minutes

**Primary authors:** Mr ALTMAN, Jeffrey (AuriStor, Inc); Mr DIONNE, Marc (AuriStor, Inc); Mr HOWELLS, David (Red Hat)

**Presenter:** Mr ALTMAN, Jeffrey (AuriStor, Inc)

**Session Classification:** Storage and file systems

**Track Classification:** Storage & Filesystems