

# NFSv4 Open Source Project Update

---

Andy Adamson  
CITI, University of Michigan  
Ann Arbor

# A little bit of history

---

- ◆ Sponsored by Sun Microsystems
- ◆ IETF reference implementation
- ◆ Linux and OpenBSD
- ◆ "Full featured" Linux 2.4.18 NFSv4 only client and server - not integrated with NFS v2/v3

# NFSv4 in Linux 2.5

---

- ◆ Stateless client and server
  - Fully integrated with v2 and v3
  - Good performance
  - Uses special stateids
- ◆ RPCSEC GSS framework
  - Client up-call architecture
- ◆ Kerberos v5 GSS mechanism
  - Mutual Authentication

# Linux 2.4 Back Port

---

- ◆ Linux 2.4.18-18 kernel with RedHat 7.3 environment
- ◆ Matches Linux 2.5 NFSv4 functionality
- ◆ RPM install for kernel and daemons
- ◆ First release: Client only
- ◆ Effort led by Chuck Lever (NetApp)

# Release Managemet

---

- ◆ New functionality released in Linux 2.5 kernel
  - Submit patches to Trond Myklebust and Neil Brown
  - They submit patches to Linus
- ◆ New functionality then added to back port
- ◆ Linux 2.4 back port: RPM release
- ◆ User daemons and configuration: RPM release

# In the Pipeline

---

## ◆ Server

- Up-call mechanism
- Name to Id and Security context (GSS) cache
- Clientid state

## ◆ Client

- Name to id up-call and cache
- Share state

# In the Pipeline

---

- ◆ FS\_Locations (client)
- ◆ RPCSEC\_GSS Data Integrity
  - Kerberos v5
- ◆ Server share state
- ◆ Client and Server ACLs
- ◆ Kernel SPKM3

# Daemons

---

- ◆ Idmapd: name to id mapper
  - Uses getpwxxx and /etc/nsswitch
  - Secure AFS PTS nsswitch option
- ◆ GSSD: GSS context initialization
  - Performs RPCSEC GSS Null proc to server
  - Able to service multiple GSS mechanisms



# Related Work

---

- ◆ MIT Kerberos RPCSEC GSS code
  - Based on existing MIT K5 rpc code
  - Kadmin and Kadmin2 interoperates with SEAM
  - submitted to MIT code base
- ◆ User level SPKM3
  - Diffie-Hellman key exchange with target X.509 keys
  - Data integrity
  - Uses OpenSSL libraries

# OpenBSD

---

- ◆ Stateless operations complete
- ◆ Share state complete
- ◆ Byte-range locking being tested
- ◆ RPC Framework in use
  - Name to ID mapping
  - RPCSEC\_GSS framework being coded
- ◆ Port to FreeBSD

# Any questions?

---

<http://www.citi.umich.edu/>

