NFSv4 Open Source Project Update

Andy Adamson
CITI, University of Michigan
Ann Arbor

NFSv4 Open Source Project

- Sponsored by Sun Microsystems
- Support from Network Appliance
- IETF reference implementation
- Linux 2.6, FreeBSD 5.2, and OpenDarwin(Mac OS 10.3.2)

Outline

- Linux 2.6: we have arrived!
- Linux utilities
- RPCSEC_GSS
- Fedora Core 2
- FreeBSD/OpenDarwin

Linux 2.6 NFSv4 Client

- Share state, Byte-range locking
- ACLs (POSIX), reboot recovery
- Atomic open (VFS change)
- Rewrite of attribute cache
- RPCSEC_GSS framework
  - krb5, krb5i mechanisms
- Idmap upcall

Linux 2.6 NFSv4 Server

- Share state, Byte-range locking
- ACLs (POSIX)
- Reboot recovery
  - OPEN and LOCK reclaim
  - upcall to non-volitile storage designed
- RPCSEC_GSS framework
  - krb5, krb5i mechanisms
- Idmap upcall

AUTH_SYS

- Hardened code
  - very few bugs found at Connectathon 2004
- Connectathon, FSX, Iozone testing
- Performance very close to NFSv3
  - v3/v4 shared read write code path
- Reboot recovery and pynfs testing

Connecthon04 AUTH_SYS

- See billboard for details
- Passed basic, general, special and locking with
  - Solaris, IBM, and Hummingbird client and server
  - Network Appliance server
- EMC and HP server
  - various tests pass, see billboard

ACL testing

- Interoperabiltiy issues with
  - ACL and mode bits
  - POSIX and NFSv4 mapping
- Tested with
  - Network Appliance
  - Hummingbird
  - IBM

Reboot Recovery Testing (AUTH_SYS)

- Testing in progress, some bugs found
- Server Reboot: Clientid, Open reclaim
  - Linux client, Solaris and Linux server
- Client Reboot: Clientid, Open reclaim
  - Solaris and Linux client, Solaris and Linux server

Linux Utilities

- nfs-utils and utils-linux patches
- Idmapd for client and server
  - uses nsswitch
- Patch for mount – nfs4 fs type
- Patch to mountd for v4 export upcall
- GSSD for client and server
- Server init scripts

RPCSEC_GSS

- Bug fixes into kernel since Austin Bakeathon
- GSSD
  - mechanism switching statically bound
  - mechanism libraries dynamically bound
  - /etc/gssapi.conf: mechanims library location: MIT Krb5, SPKM3
- Kernel and user mechanism switching framework tested with Krb5 and SPKM3

SPKM3

- Austin Bakeathon
  - Sun, IBM, Network Appliance, CITI, etc: implementation discussion
- Implemeted Austin changes
  - SPKM3 mutual authentication with DHKeyAgreement is functional in Linux 2.6 (patch, not submitted)
  - GSSD SPKM3 library built against OpenSSL
- This is a work in progress!

Connecthon04 AUTH_GSS/KRB5

- See the billboard for details: still testing
- Passed most tests with krb5, krb5i mechanisms
  - Linux client: Solaris, IBM, Hummingbird,Network Appliance server
  - Linux server: Solaris, IBM, Hummingbird client

Fedora Core 2 and NFSv4

- We are working with Steve Dickson to provide the first Linux NFSv4 distribution
- Based on Linux 2.6 kernel
- NFSv4 AUTH_SYS and AUTH_GSS/krb5,krb5i client and server

Fedora Core 2 and NFSv4

- Will contain features described in this presentation
- Fedora version of nfs-utils and utils-linux, with up-stream integration to follow
- See ***fedora.redhat.com*** for details
- Other Linux distributions to follow

FreeBSD, OpenDarwin Client

- Completed port from OpenBSD to FreeBSD
  - share state complete
  - RPC seperated from NFS Client code
  - porting OpenBSD server to FreeBSD
- FreeBSD client in FreeBSD 5.2
- Ported FreeBSD client to OpenDarwin
- Currently a shared code base

Connectathon 2004 Tests

- FreeBSD Client passes basic, general, and special tests against
  - Linux, Solaris, IBM, Network Appliance
  - see billboard for other server results
- The *brand new* OpenDarwin (Mac OS 10.3.2) client
  - intermittently passes basic tests against Linux and IBM servers

Citi and Open Source

- A measureable value of sponsoring CITI is the acceptance of CITI code in the open source community
- Linux 2.6 kernel
  - Trond Myklebust working at CITI (thanks Beepy, Network Appliance!)
  - Neil Brown accepting/reviewing patches
    - » submitted 34 patches to the Linux 2.6 kernel yesterday!

Citi and Open Source

- FreeBSD 5.2
  - NFSv4 client in FreeBSD 5.2
  - server to follow
  - direct write access to FreeBSD 5.2
- OpenDarwin
  - NFSv4 client in OpenDarwin
  - server to follow
  - commit access to OpenDarwin

Citi and Open Source

- MIT Krb5
  - working with MIT developers to create a new interface for gss_context export needed to import context into the kernel
  - CITI rpc onc rpcsec_gss patch slated for inclusion in MIT krb5 1.4

What's Next

- Finish NFSv4.0 features
  - Delegation, FS_LOCATIONS, Named Attributes
- Begin NFSv4.1 feature development
  - NFS over RDMA work has started
- Explore related technologies
  - global namespace
  - migration/replication
  - NFSv4 extensions for cluster computing

Any Questions?

http://www.citi.umich.edu/projects/nfsv4