Evolving & Improving Red Hat Enterprise Linux NFS (handout)

Steve Dickson
Consulting Software Engineer, Red Hat
Thursday, June 13

Evolution of NFS

- NFSv2 - RFC 1094, March 1989
  - Side car protocols with mounts and file locking.
    - Not firewall friendly
  - 32 bit file handles
    - No files > 2GIG
  - Slow writes
    - (write, wait, write, wait)
  - Each component is an OTW lookup
    - /foo/bar ==> (lookup foo, lookup bar)

Evolution of NFS

- NFSv3 - RFC 1813, June 1995 (6 years later)
  - 64 file handle file > 2GB
  - Async writes
    - write, write, commit, wait
  - READDIRPLUS
    - readdirs and lookups all in the same call
  - Post attribute ops
  - Entire path is looked up
    - /voo/bar ==> (lookup foo/bar)

Evolution of NFS

- NFSv4.0 - RFC 3010, Dec 2000 (5 years later)
  - Client caching (aka deletions (aka Statefull))
  - Strong security
  - Compound Operations
  - TCP only
- NFSv4.1 - RFC 5661 Jan 2010 (10 years later)
  - pNFS
  - _reliable_ only-once semantics
  - callbacks using 2049 port

Evolution of NFS

- NFSv4.2 - IETF-Draft ver 19
  - Server-side Copy offload
  - Sparse file support
  - Space Reservations
  - Label NFS (SELinux support)

Parallel NFS (pNFS)

- Architecture
  - Metadata Server (MDS) – Handles all non-Data Traffic
  - Data Server (DS) – Direct I/O access to clients
  - Shared Storage Between Servers
- Layout Define server Architecture
  - File Layout (NAS Env) - Netapp
  - Block Layout (SAN Env) - EMC
  - Object Layout (High Perf Env) Pananas & Tonian

Parallel NFS (pNFS) - RHEL 6.4

- First to market with Client support (file layout)
  - Thank you very much Upstream and Partners!!!
- Enabling pNFS:
  - **mount -o v4.1 server:/export /mnt/export**
- RHEL-Next
  - Block and Object layout support

RHEL7 NFS Server Updates

- Red Hat Enterprise Linux 7.0 completes the server side support for NFS 4.1
  - Support for only-once semantics
  - Callbacks use port 2049
- No server side support for parallel NFS ... yet!

Parallel NFS Updates

- Parallel NFS has three layout types
  - Block layouts allow direct client access to SAN data
  - Object layouts for direct access to the object backend
  - File layout
- RHEL7.0 will add support for block and object layout types
  - Will provide support for all enterprise pNFS servers!

Support for SELinux over NFS

- Labeled NFS enable fine grained SELinux contexts
  - Part of the NFS4.2 specification
- Use cases include
  - Secure virtual machines stored on NFS server
  - Restricted home directory access

NFS V4 Referrals - HOWTO

- On redhat-1 Server:
  - Export file system with: **refer=/export@redhat-3**
  - Bind mount file system: **mount –bind /export /export**
  - Start nfs server: **service nfs start**
- On the Client:
  - Mount file system: **mount server:/export /mnt/export**
  - Create the referral: **cd /mnt/export**

FedFS - HOWTO

- Documentation of of this HOWTO is at
  - http://wiki.linux-nfs.org/wiki/index.php/FedFsUtilsDocs
- Install fedfs-utils-0.9 on Fedora 19/RHEL 7
  - On the Client
    - yum install fedfs-utils-client
  - On the Server
    - yum install fedfs-utils-lib fedfs-utils-nsdbparams
    - yum install fedfs-utils-server

FedFS – Setting up client

- yum install autofs
- mkdir /nfs4
- /nfs4 /usr/sbin/fedfs-map-nfs4 ==> /etc/auto.master
- systemctl [re]start autofs

FedFS – Setting up Domain Root Server

- Create Domain Root exports
  - mkdir -p /.domainroot/redhat.com
  - export /.domainroot *(ro,sec=sys,insecure)
- Configure DNS so root server can be found
  - _nfs-domainroot._tcp IN SRV 0 0 2049 server.redhat.com
- Create Junction
  - nfsref add home home.redhat.com /home

Secure NFS – IPA Server

- Install package
  - yum install ipa-server
- Install server bits
  - Ipa-server-install
- All configuration is done with 'admin' ticket
  - kinit admin
- Create 'nfs' keytab principles
  - ipa-getkeytab -k /etc/krb5.keytab -s <ipserver> -p nfs/`hostname`@REDHAT.COM

Secure NFS – IPA Server

- Start Secure NFS server
  - SECURE_NFS=”yes” in /etc/sysconfig/nfs
  - systemctl start nfs-server.service
  - systemctl start nfs-secure-server.service
- Add users
  - ipa user-add ksteved
  - ipa passwd ksteved

Secure NFS – IPA client

- Install package
  - yum install ipa-client
- Install client bits
  - Ipa-client-install
- Start Secure NFS GSS deamon
  - SECURE_NFS=”yes” in /etc/sysconfig/nfs
  - service rpcgssd start
- Secure mount are now possible!!!

Secure NFS – IPA client

- Enable sssd pam & Create Home dirs
  - authconfig --update --enablesssd --enablesssdauth –enablemkhomedir
- service sshd restart
- ssh ksteved@ipa-client.com

Secure NFS – IPA client

- Have SSSD renew the Tickets (see man sssd-krb5)
  - auth_provider = krb5
  - krb5_server = 192.168.1.1
  - krb5_realm = EXAMPLE.COM
  - krb5_renewable_lifetime = 50d
  - krb5_renew_interval = 3600
- service sssd restart

GSS-proxy - HOWTO

- Install IPA
  - yum install ipa-client
  - ipa-client-install
- Install GSS-proxy
  - Yum install gssproxy
- Enable Secure NFS
  - GSS_USE_PROXY="yes" ===> /etc/sysconfig/nfs
  - systemctl restart nfs-secure.service

GSS-proxy - HOWTO

- Creating user keytabs
  - ipa-getkeytab \
    -s ipa-server.redhat.com \
    -p steved@redhat.com \
    -k /var/lib/gssproxy/clients/steved.keytab