NFS INDUSTRY CONFERENCE

AFS* to NFS Migration
Global Data Sharing at Intel

Travis V Broughton

Senior Systems Programmer

Intel Corporation

tvb@intel.com

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Intel's Engineering Environment

- ~80% NFS, ~15% AFS*, ~5% CIFS
  - ~30 AFS cells managed by ~30 loosely-connected IT organizations
  - AFS used for CAD and /usr/local applications, global data sharing for projects, secure access to data
  - NFS used for everything else, gives higher performance in most cases
  - Wide range of client platforms, OSs, etc

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Things AFS* Does Well

- Security
  - Uses Kerberos, doesn't have to trust client
  - Uses ACLs, better granularity
- Performance for frequently-used files
  - e.g. /usr/local/bin/perl
- High availability for RO data
- Storage virtualization
- Global, delegated namespace

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Things AFS* Doesn't Do Well

- Performance on seldom-used files
- High availability for RW data
- Scalability with SMP systems
- Integration with OS
- File/volume size limitations

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

AFS* Usage at Intel: Global Data Sharing

- Optimal use of compute resources
  - Batch jobs launched from site x may land at site y, depending on demand
- Optimal use of headcount resources
  - A project based at site x may "borrow" idle headcount from site y without relocation
- Optimal license sharing
  - A project based at site x may borrow idle software licenses (assuming contract allows "WAN" licensing
- Efficient IP reuse
  - A project based at site x may require access to the most recent version of another project being developed at site y

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

AFS* Usage at Intel: Other Applications

- x-site tool consistency
  - Before rsync was widely deployed and SSH-tunneled, used AFS namespace to keep tools in sync
- @sys simplifies multiplatform support
  - Environment variables, automounter macros are reasonable workarounds
- /usr/local, CAD tool storage
  - Cache manager outperforms NFS
  - Replication provides many levels of fault-tolerance

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Global Filesystem Requirements

- Platform Independence
  - Client interoperability is a must
  - Choice of server platforms also desirable
- Global namespace
- Strong authentication/authorization
- WAN-friendly

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Global Filesystem "nice-to-have"s

- Suitable for local data
  - HA capabilities
  - High performance for "hot" files
  - Replication with consistency assurance
- Data location independence
  - Data migration for performance, space, maintenance, etc.

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Global Filesystem "nice-to-have"s

- Flexible quotas
- Filesystem encryption
  - Transport layer?
  - On-disk?

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

AFS* is going away, now what?

- OpenAFS*?
  - Same issues as AFS, mostly
- NFSv3?
  - Not secure enough for global data sharing
  - NFSv3+Kerb not supported on all platforms
- Caching NFS proxy solutions?
  - Somewhat cumbersome to use globally
  - Security?
  - Possibly a point solution
- NFSv4?
  - Looks promising, concerned about roadmaps

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

NFSv4 Opportunities

- Single filesystem for all data
  - No longer have to ask “AFS* or NFS”
- WAN performance improvements
  - Probably better than AFS, even w/o cachefs
- HA via clustering should just work
- Filesystem tested and released by OS vendors
- NAS cost-effectiveness

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

NFSv4 Challenges

- Integration of lots of enabling technologies
  - Automounter
    - autofs or am-utils for clients?
    - LDAP, NIS, etc for map storage?
  - Security
    - Kerberos KDC (MIT? Heimdal? Active Directory*?)
    - Different ACL flavors on different servers?
    - Group management? Global groups?
- Support for legacy client platforms
  - Access for clients without native NFSv4
- Optional features and/or “enhancements”
  - e.g. cachefs

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

NFSv4 Challenges

- Global namespace management
  - Delegate portions of namespace to different IT orgs
- New security model
  - New authentication mechanisms
- User and group consistency
- Timeline for client availability vs. IBM's AFS* EOL schedule

* Other names and brands may be claimed as the property of others

NFS INDUSTRY CONFERENCE

Conclusions

- AFS* EOL will require changes
- NFSv4 looks like a good fit for global data sharing
- AFS to NFS migration non-trivial
  - Enabling technologies and tools needed

* Other names and brands may be claimed as the property of others