

IPv6 support for Linux NFS

Chuck Lever, Oracle Corporation

Use cases

- Today some applications can run natively on IPv6-only networks; NFS can't
- IPv4 / IPv6 mixed environments
- IPv6-only environments

Project status

- Kernel timeline
- User space
- What's missing

Kernel timeline

- Pre-2.6.23 kernels have:
 - Some RPC server support for IPv6
 - Client-side support for rpcbind versions 3 and 4

Kernel timeline

- 2.6.23 introduces:
 - String-based NFS mount option parsing
 - Subsumes legacy binary `nfs_mount_data` mount option structure
 - NFSv2/v3 mountd client in-kernel
 - Needed for many advanced NFS features including IPv6, NFS/RDMA, cache FS

Kernel timeline

- 2.6.24 adds:
 - IPv6 support in the in-kernel RPC client
- 2.6.25 will have:
 - IPv6 infrastructure in the NFS client (but not in NLM or NFSv4 callbacks)

Kernel timeline

- 2.6.26 may have:
 - IPv6 support in the in-kernel NLM and NSM
 - Remaining patches to support IPv6 in the in-kernel NFS server

User space progress

- Development components
 - libtirpc
 - Provides IPv6-enabled user-space RPC facilities
 - Collides with legacy RPC facilities already in glibc

User space progress

- Daemons
 - rpcbind replaces portmapper
 - rpc.statd
 - rpc.mountd

User space progress

- Client-side command line utilities
 - mount.nfs & friends
 - NFSv4 support is easy: umount is local-only, no need for getport
 - NFSv2/v3 require version and transport discovery for NFS and NLM
 - nfs(5) updates

User space progress

- Server-side command line utilities
 - exportfs must support specifying IPv6 addresses in export rules
- exportfs(8) updates

What's missing?

- RPC pipefs changes
- IPv6 support for NFSv4 callbacks and referrals
- Unknown requirements for advanced security flavors
- Significant test capabilities

Milestone

- Expect basic IPv6 support in all NFS components to become available for distributions to begin integrating during 2H08

Discussion topics

- What NFS on IPv6 use cases are a priority?
- Who can help finish the implementation?
- Who can test?