



Evolution Of NFS

Steve Dickson



Agenda

- History of NFS
- Why NFS v4
- Feature list
- NFS v4 Architecture
- Secure NFS
- RHEL4 Feature list
- Testing Matrix
- Debugging tips



History of NFS

- Mid 1980's Sun Developed NFS
 - ➔ Version 1 was never released.
- Mid 90's NFS version 3 was released
 - ➔ Large file support – 64bit file sizes
 - ➔ Asynchronous I/O - commits
 - ➔ Read Directories with Attributes - READDIRPLUS



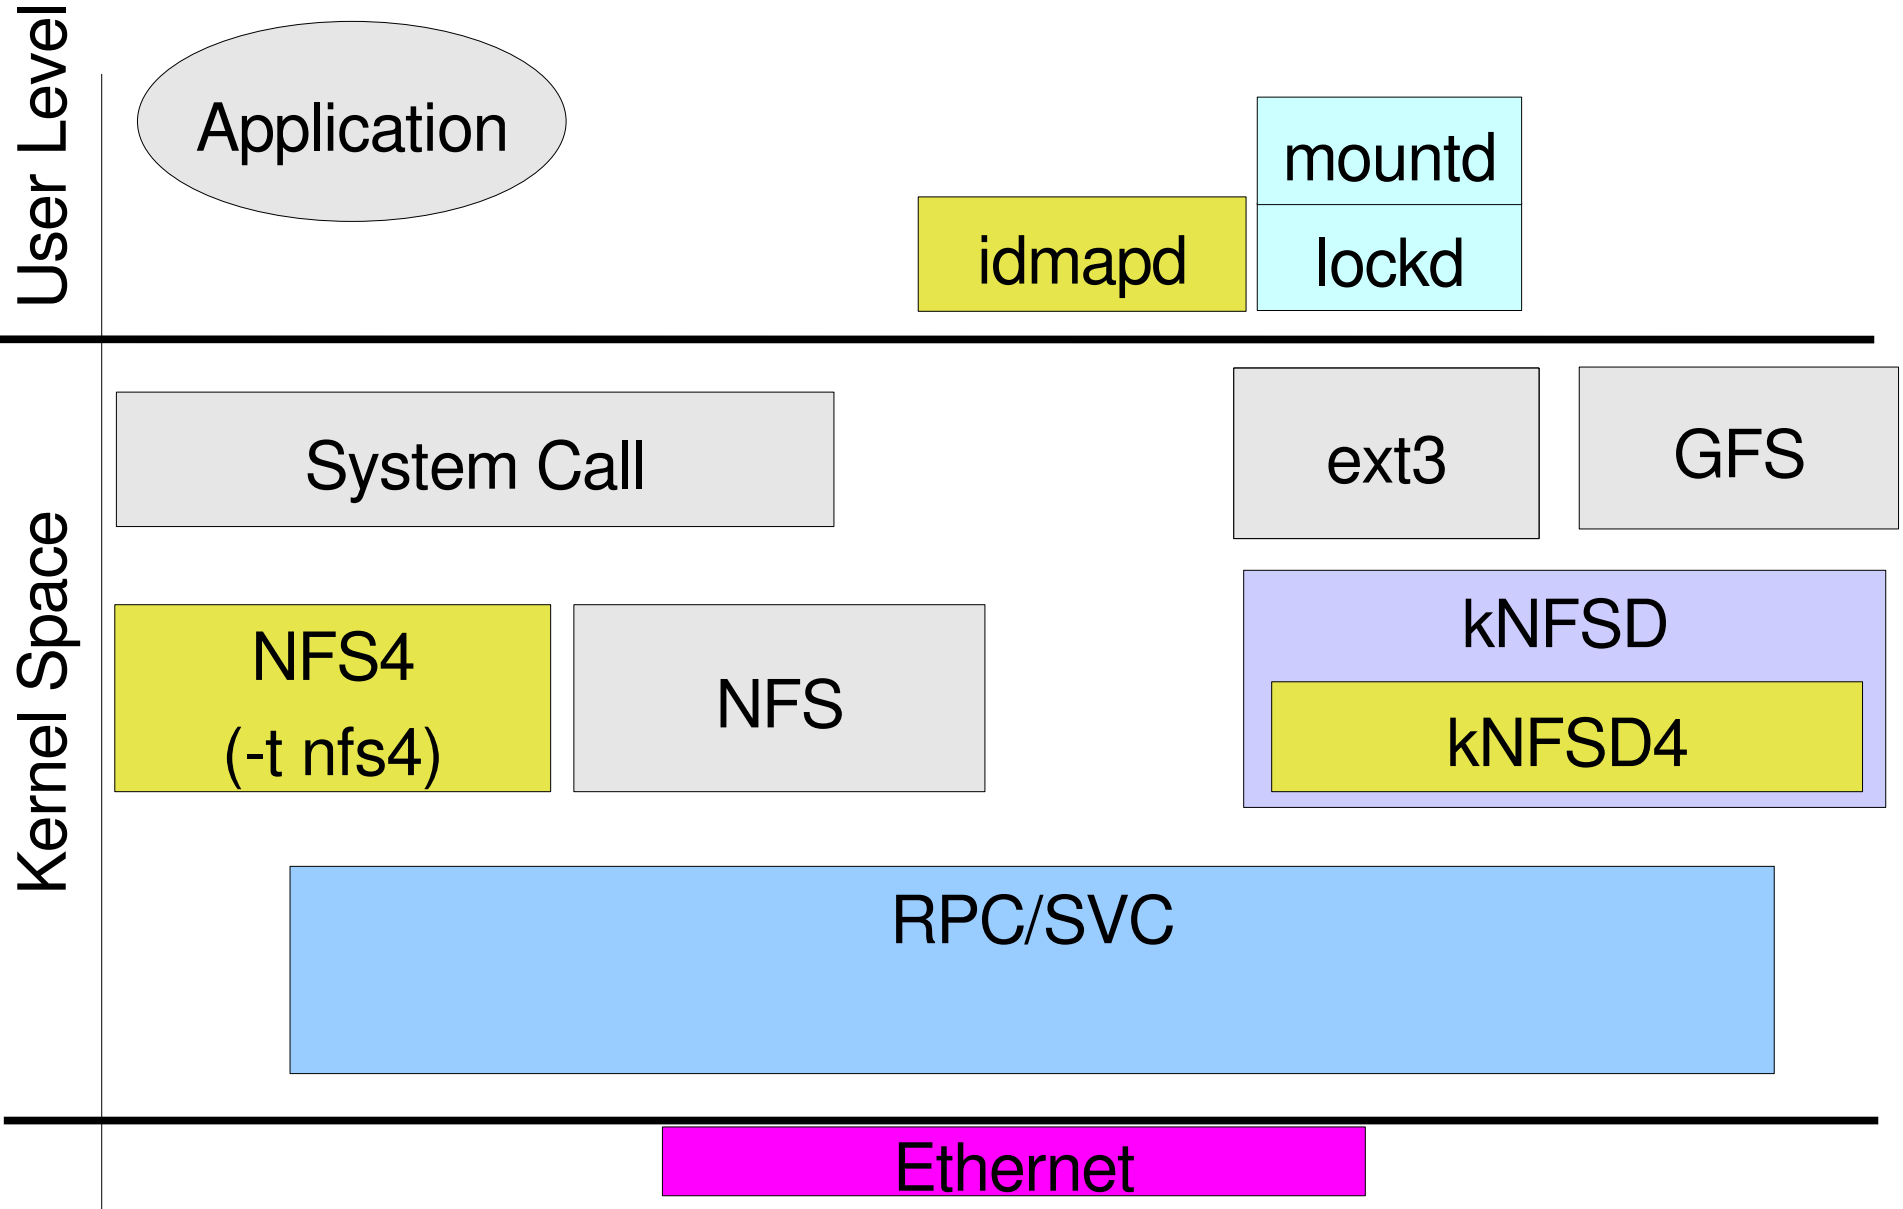
Why NFS Version 4?

- Performance
- Elimination of 'side-car' protocols
- Multi-Component Lookups
- Mandates strong security architecture
- Server maintains client state



NFS Architecture

NFS4 Architecture



NFSv4 Protocol Feature List

- **Compound Procedures**
 - ➔ Multiple operations sent in one Over-The -Write message.
- **Firewall Friendlier**
 - ➔ Mount and locking protocols are integrated into protocol
 - ➔ Only TCP is supported
- **Open and Close Operations**
 - ➔ Atomic creates supported
- **Pseudo File System**
 - ➔ Shared server namespace
- **File Delegation**
 - ➔ Allowing clients to aggressively cache both data and metadata



NFSv4 Feature List (cont'd)

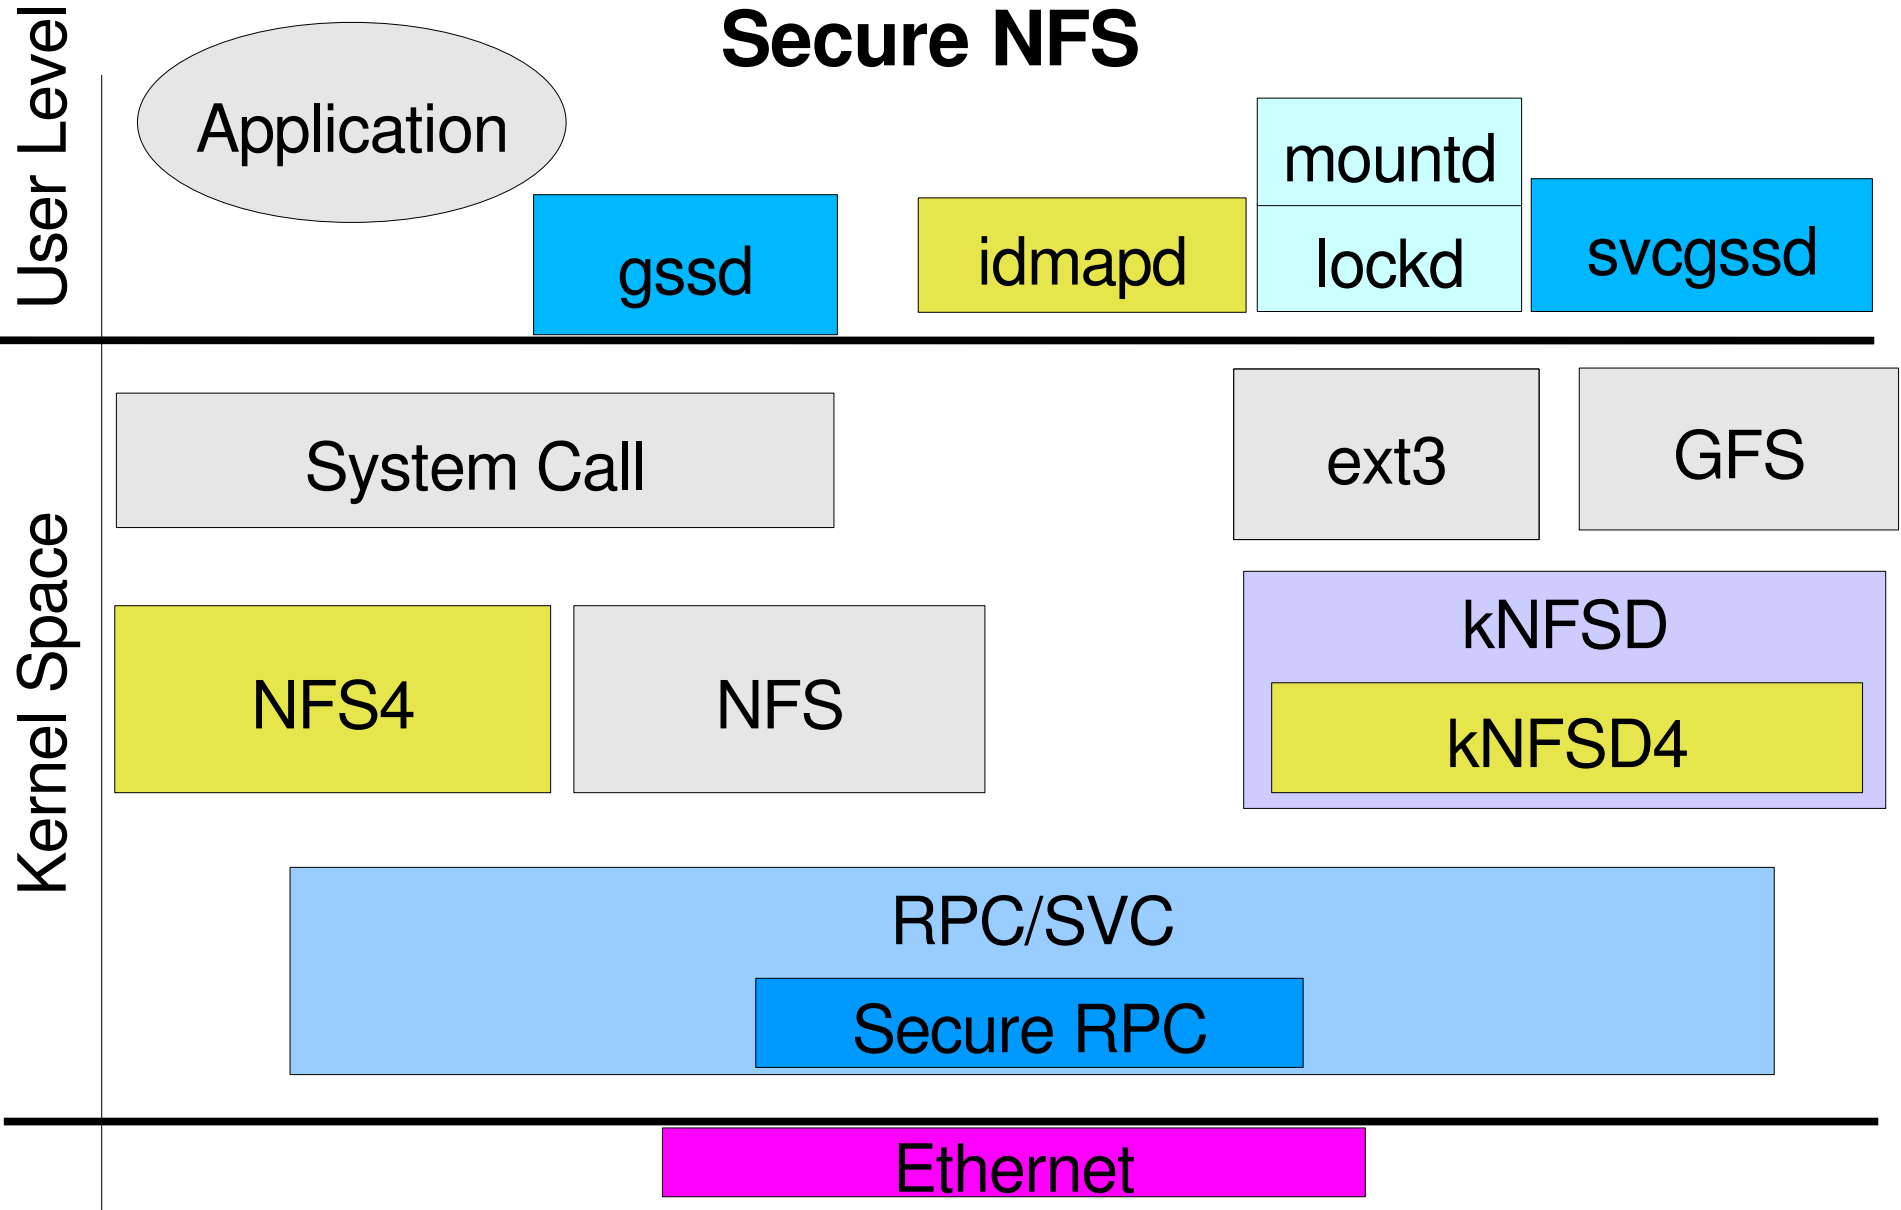
- UTF-8 Strings are used for User/Group ids
 - ➔ Allow for Internationalization support
 - ➔ rpc.idmapd – maps `user@domain` to Linux UIDs on server and client.
- Integrated Access Control List (ACL) support
 - ➔ NT style ACLs
- File Migration and Replication
- Named Attributes
- Designed for future protocol extensions



NFS Architecture

NFS4 Architecture

Secure NFS



Secure NFS

- Used by ALL three NFS versions
 - Use the '-o sec=krb5' mount option
- Uses GSS-API cryptographic method.
- Three Kerberos 5 security levels
 - Authentication (RPC header is signed)
 - Integrity (Header and Body are signed)
 - Privacy (Header signed. Body encrypted)




Secure NFS (cont'd)

- User level Daemons used to Handle complicated context initiation phase
 - ➔ `rpc.gssd` – Client security contexts
 - ➔ `rpc.svcgssd` – Server security contexts
- Set `SECURE_NFS` in `/etc/sysconfig/nfs`
- Both use files in the `rpc_pipefs` filesystem to do “upcalls” to the kernel.



Secure NFS (cont'd)

- Security Context Data flow



RHEL4 NFS4 Feature list

- Compound Procedures
- Firewall Friendlier
- Open and Close Operations
- Pseudo File System (export w/ nohide)
- File Delegation (client only)
- Strong Security



Future RHEL Feature list

- File Delegation (server support)
- Integrated ACL support
- File Migration and Replication
- Named Attributes



NFSv4 Testing

- OSDL Test Matrix
 - A test plan defining over 300 areas to test.
- Connectathon
 - Interoperability with other NFS-es.



Debugging Techniques

- Ethereal is your friend!
 - Binary captures only... Please...
- nfsstat
 - Good to see to look at retransmissions
- ifconfig eth?
 - Look for driver errors



Debugging Techniques (cont'd)

- Hung processes
 - System trace (echo t > /proc/sysrq-trigger)
 - Look for hung nfsd on server



References

- CITI NFSv4 Project – Univ of Michigan
 - ➔ <http://www.citi.umich.edu/projects/nfsv4>
 - ➔ <http://www.citi.umich.edu/projects/nfsv4/gssd/>
- NFSv4 Test Maxtrix - OSDL
 - ➔ <http://developer.osdl.org/dev/nfsv4>
 - ➔ <http://developer.osdl.org/dev/nfsv4/testmatrix/>
- The NFS version 4 Protocol
 - ➔ Presented at SANE 200. Written by Pawlowski, Shepler, Beame, Callaghan, Eisler, Noveck, Robinson and Thurlow.
 - ➔ <http://www.nluug.nl/events/sane2000/papers/pawlowski.pdf>



References (con't)

- Linux NFS Version 4: Implementation and Administration
 - ➔ Presented at OLS 2001.
 - ➔ Written by William A Adamson (CITI) and Kendrick M. Smith
 - ➔ http://lwn.net/2001/features/OLS/pdf/pdf/nfsv4_ols.pdf



Questions

