

NFS Version 4 Update

Brent Callaghan
Sun Microsystems, Inc
brent@eng.sun.com

NFS v4 @ IETF Status

- **37th IETF BOF @ San Jose December 96**
 - **Good attendance (~70)**
- **Mailing list & Archive**
 - **`nfs4-wg@sunroof.eng.sun.com`**
 - **`http://playground.sun.com/pub/nfsv4/nfsv4-wg-archive`**
- **Change Control Agreement**
 - **IETF must have change control**
 - **Agreement based on RFC 1790 (ONCRPC)**
 - **“Legal Time” is inverse of “Internet Time”**
 - **Will be published as an Informational RFC**
- **42nd IETF Chicago 23-28 August 1998**

Draft WG Charter

“The objective of this working group is to advance the state of NFS technology by producing a specification for NFS version 4 which will also be submitted as an Internet Standard. NFS version 4 will emphasize the following features:”

- **Improved access and good performance on the Internet**
- **Strong security with negotiation built into the protocol**
- **Better cross-platform interoperability**
- **Designed for protocol extensions**
- **Others that may be proposed**

nfs4-wg Mailing List

1 akk@eng	wg@transarc.com	68 tytso@MIT.EDU
2 brent@eng	35 beepy@netapp.com	69 vijay@tagore.wipinfo.soft.net
3 brian.oki@eng	36 boris@nextpoint.com	70 werme@zk3.dec.com
4 bwelch@eng	37 btg@cray.com	71 ed_zayas@novell.com
5 clv@eng	38 Cyril.Plisko@Orbotech.Co.IL	72 jcg161@psu.edu
6 corbin@eng	39 dima@bog.msu.su	73 A.McDermott@ftel.co.uk
7 frankh@eng	40 frank@wins.uva.nl	74 pokallus@ccr-p.ida.org
8 geoff.arnold@East	41 fred@bws.bws.com	75 sulee@ingr.com
9 gww@ebay	42 glowell@engr.sgi.com	76 adam@math.tau.ac.il
10 hacker@eng	43 guy@netapp.com	77 partain@dcs.gla.ac.uk
11 smajee@eng	44 gww@traakan.com	78 Michael.Kuschke@Materna.DE
12 shepler@eng	45 hare@netmanage.com	79 sgard@earthlink.net
13 hai-yi.cheng@eng	46 ira@ftp.com	80 kraemerf@VNET.IBM.COM
14 hisgen@eng	47 jdreyer@pobox.com	81 menze@retrieveinc.com
15 kupfer@eng	48 jes@machine.engr.sgi.com	82 vahalia@emc.com
16 labiaga@eng	49 Jim.Zelenka@cs.cmu.edu	83 nconlwf@nabaus.com.au
17 lling@eng	50 jms@cray.com	84 Damon.Atkins@nabaus.com.au
18 cfs@eng	51 kane@cisco.com	85 Kurt.McDougall@BarwonWa-
19 mahalingam.mani@eng	52 kessler@bigfun.engr.sgi.com	ter.Vic.Gov.AU
20 michael.carney@East	53 kevinw@ancor.com	86 skottie@fa.disney.com
21 mre@eng	54 lance@eco.twg.com	87 conrad@apple.com
22 pen@eng	55 mckusick@flamingo.McKusick.COM	88 george@thursby.com
23 Bill.Pittore@East	56 mohan@netapp.com	89 davidsi@sco.com
24 jim.roberts@East	57 moore+nfsv4@cs.utk.edu	90 phils@sco.com
25 robinson@eng	58 muralir@auspex.com	91 luan@almaden.ibm.com
26 sankar.ramamoorthi@eng	59 nfsv4-wg@austin.ibm.com	92 jefferson@pa.dec.com
27 sohrab.modi@eng	60 nfsv4@hcl.com	93 knvijay@novell.com
28 staubach@eng	61 okir@monad.swb.de	
29 thurlow@eng	62 rajen@distinct.com	
30 traj@eng	63 rick@snowwhite.cis.uoguelph.ca	
31 alex@veritas.com	64 ruthtang@vnet.ibm.com	
32 ampass@ftp.com	65 steved@zk3.dec.com	
33 bae@netapp.com	66 sv@mail.crystaliz.com	
34 bb+transarc.external.nfs4-	67 tg@bosun.bm.lu.se	

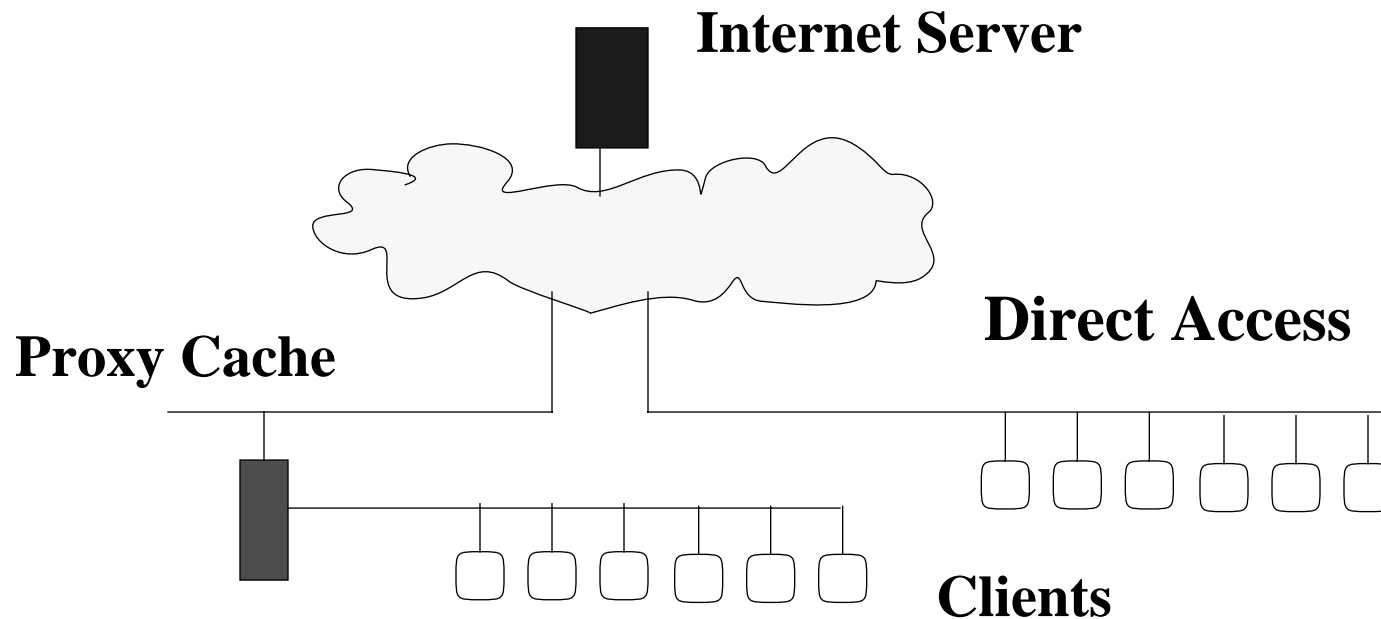
Protocol vs API

- **Does protocol support API or API support protocol ?**
- **Filesystem APIs assume local disk:**
 - **Infinite bandwidth**
 - **Zero latency**
 - **Single cache**
 - **No failures**
- **NFS v4**
 - **Best solution for POSIX ?**
 - **Best serves needs of POSIX & Win32 API's ?**
 - **Best file access for Internet ?**

V4 Discussions

- **File Naming**
 - **Unicode ? UTF-8 ?**
 - **Case sensitivity**
 - **Case preserving**
 - **Case sensitive**
- **Facilitate Operation on Internet**
 - **Single connection, single protocol**
 - **No “sideband” protocols**
 - **Incl MOUNT, NLM, STATMON, RQUOTA**
 - **Compound requests c.f. SMB Chaining**
 - **Reduce round trips**
 - **Simplify protocol**

- **Proxying**
 - **Necessary for Internet scaling**
 - **Faster data access for clients**



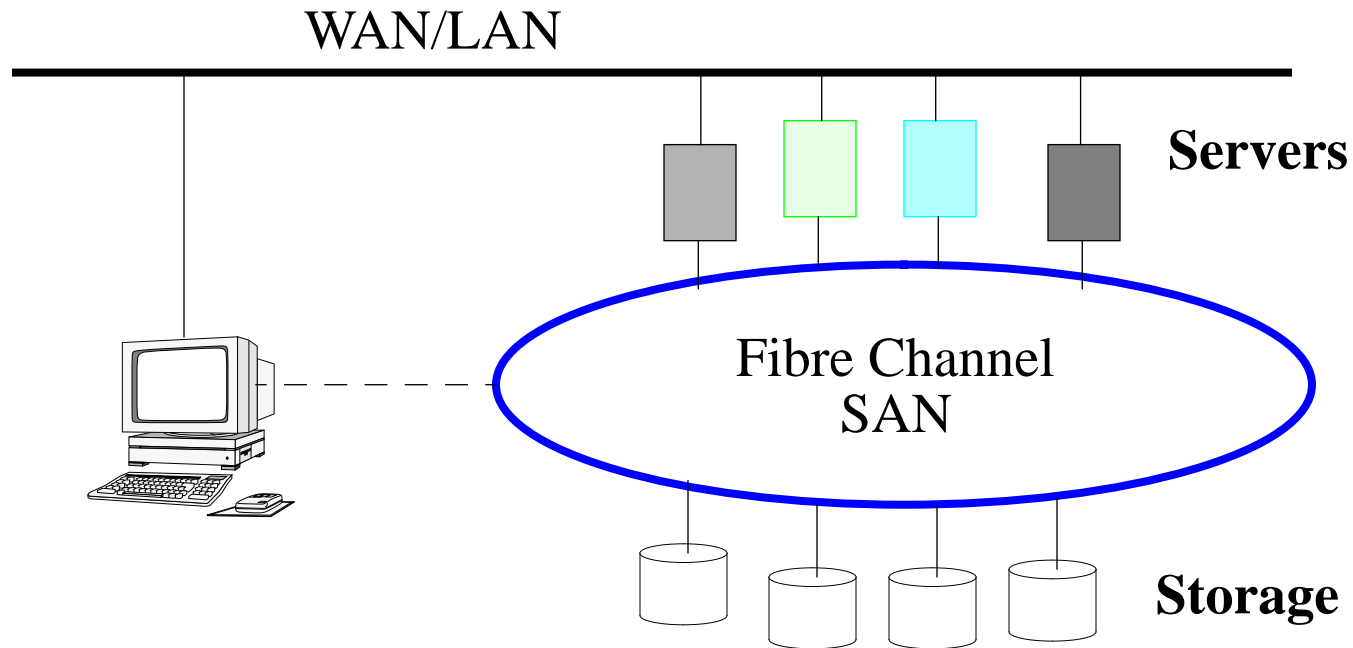
- **Need protocol support for indirection**
- **Access control for cached data**

- **File/Filesystem Attributes**
 - **“Core” vs “extended”**
 - **Extensible attributes**
 - **Access Control Lists (ACLs)**
 - **POSIX, DCE/DFS, WinNT ?**
 - **Attribute naming**
 - **String: “size” “mtime” “type”**
 - **Number: 9 18 32**
 - **Bitmap: 0x50ff**
- **MOUNT protocol equivalents**
 - **Lists of exports and properties**
 - **Lists of clients**
 - **mount/unmount notification**

- **File Locking**
 - **What kind of locking - POSIX or PC share ?**
 - **Used for cache consistency, e.g. oplocks ?**
 - **Callbacks vs leases**
 - **Scalability**
 - **Locking through proxy**
- **Filesystem Model**
 - **Multiple, named data streams (NTFS, UDF)**
- **Volatile Filehandles**
- **Identifying users & groups**
 - **UID & GID: integers or strings or ... ?**

- **Security**
 - **Is RPCSEC_GSS the answer ?**
 - **Negotiation: SPNEGO or within protocol ?**
- **Minor Versioning**
 - **Add new procedures, attributes, filetypes**
 - **Avoid protocol “hacks”**
 - **Standard process**

- **Network Attached Storage**
 - **Fibre Channel SANs**
 - **Server can be bottleneck**
 - **File access vs block access**



Questions ?
Comments ?