Checksums for NFSv4?

Alok Aggarwal
Sun Microsystems
alok.aggarwal@sun.com

Agenda

- Data Integrity Domains in a System
- NFS Data Integrity
- Checksums for NFSv4
- Requirements
- Q & A

Data Integrity Domains

System-wide

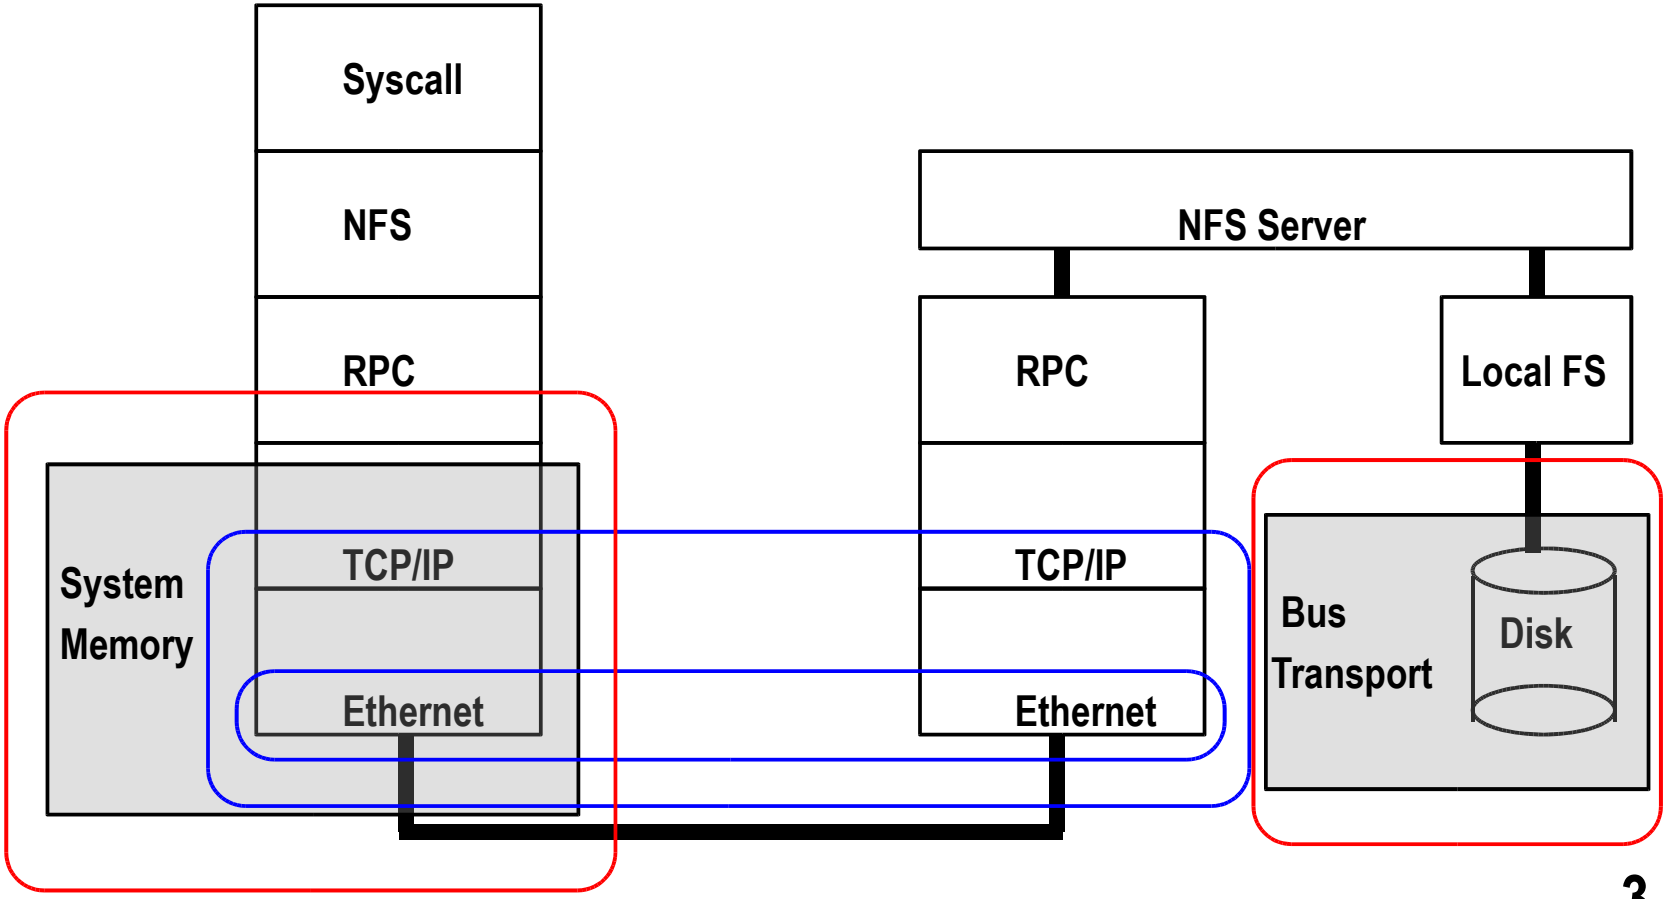

Data Integrity Domains

System-wide

- Parity protection for System Memory
  - Weak, double bit errors might cancel out
- Parity protection for Bus Transport
  - Strong, some double bit errors might still cancel out
- Parity for Memory/Bus – No Integration

NFS Data Integrity

Without RPCSEC_GSS

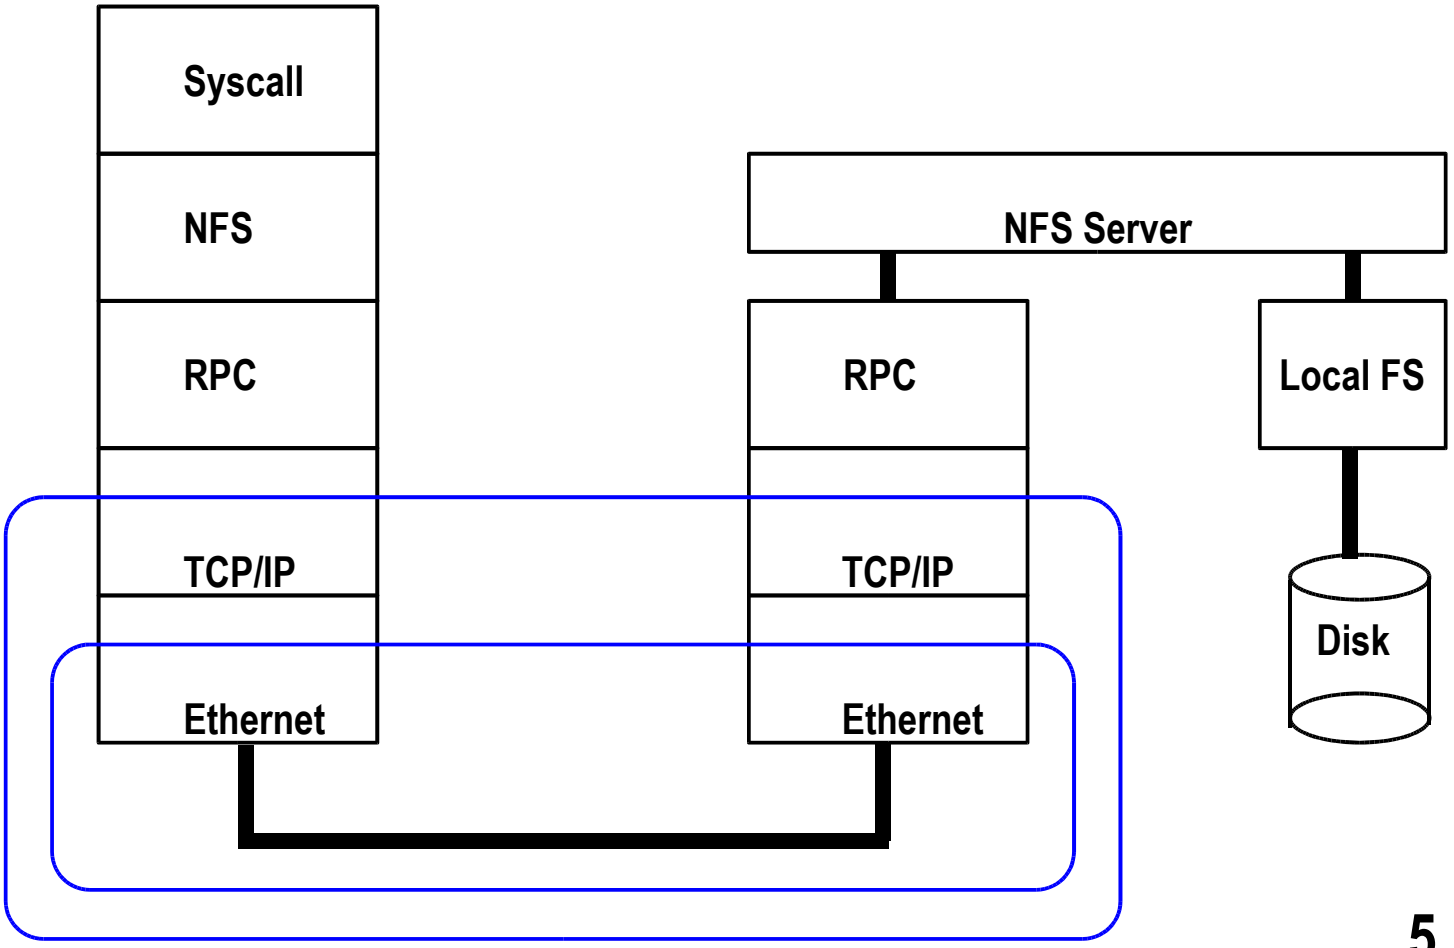

NFS Data Integrity
Without RPCSEC_GSS

- Reliance on TCP checksums
  - Standard checksum is one's complement of 16bit integers
    - Designed for speed, byte order independence
    - Double bit errors in a datagram might cancel out
    - Inability to detect transposition of octets/words in a datagram
    - That is, one's complement of the following datagrams is the same (each box is 16bits)

  - Alternate checksum algorithm optional (RFC 1146)
    - Google doesn't find any implementations that support it

NFS Data Integrity
Without RPCSEC_GSS

- Reliance on Ethernet checksums
  - CRC32 used as the standard checksum
    - Efficient in error detection and error correction
    - Downside is it's not end-to-end and gets re-computed at every hop

NFS Data Integrity

What's left?

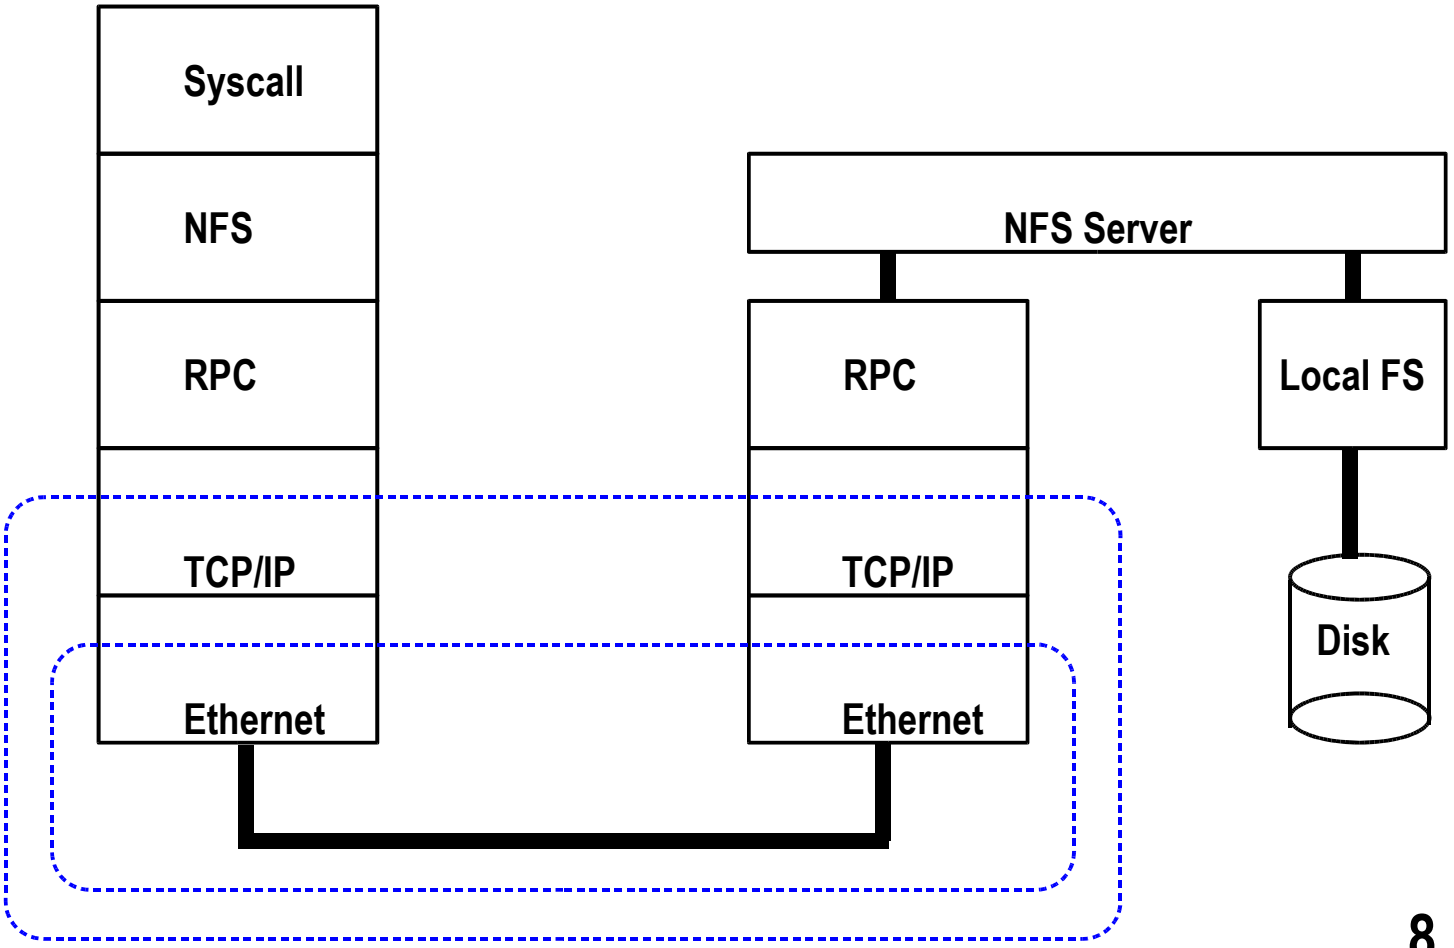

NFS Data Integrity
With RPCSEC_GSS

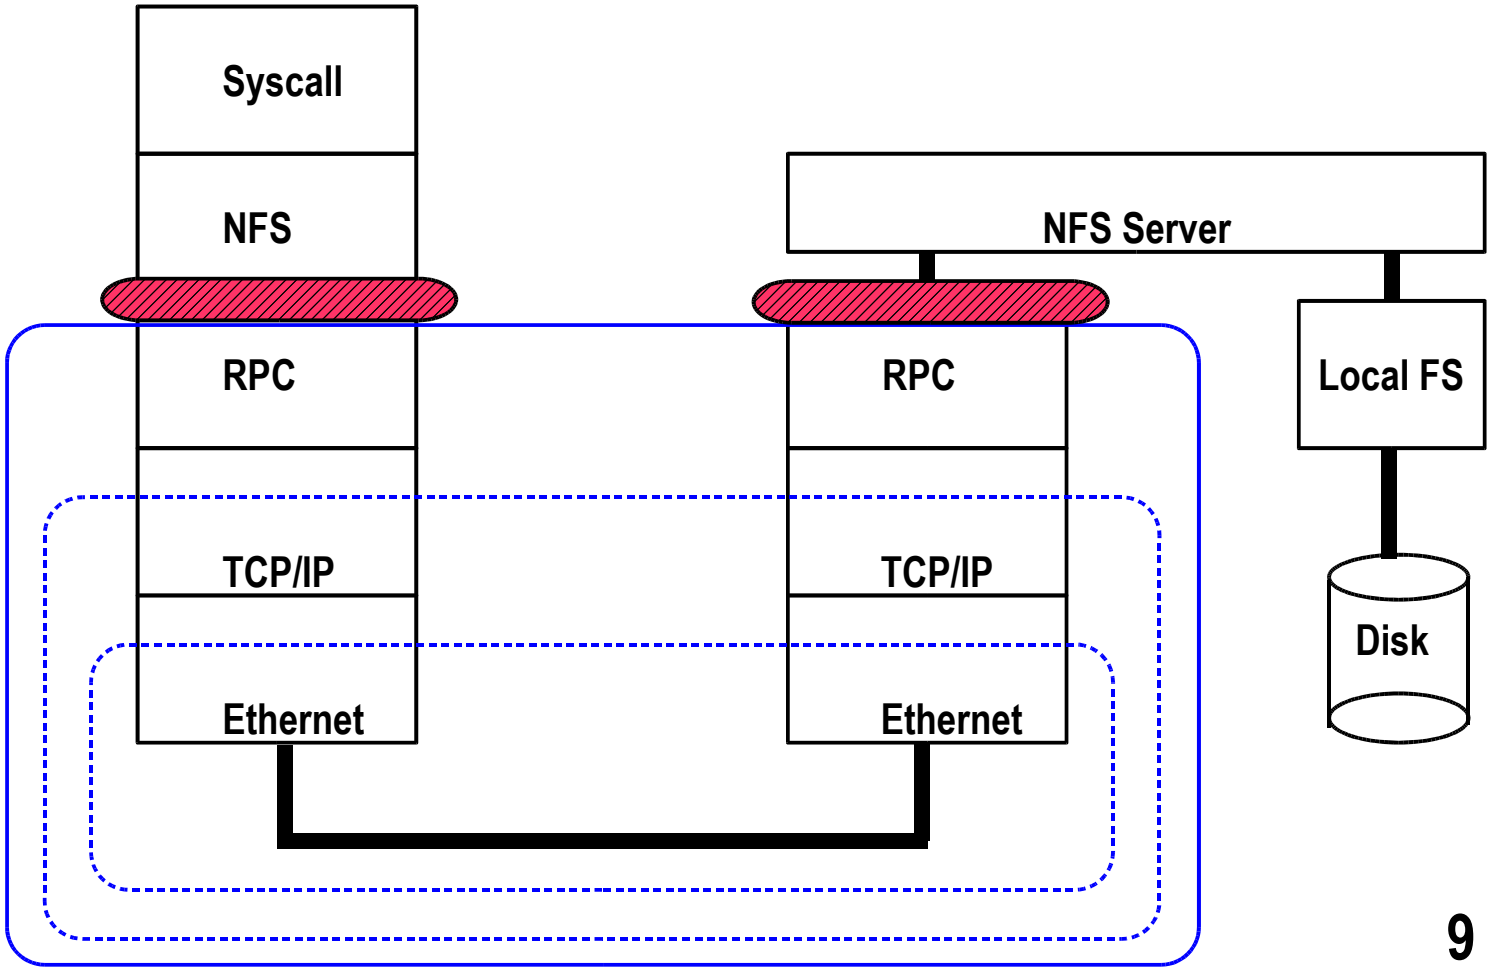

NFS Data Integrity
With RPCSEC_GSS

- Entire RPC payload protected
  - Strong protection
- Data is vulnerable after its decoded
  - No Integration between the NFS and the RPC layer
- Security Mechanism (Kerberos) needs to be configured!

Checksums for NFSv4

End-to-End Integrity

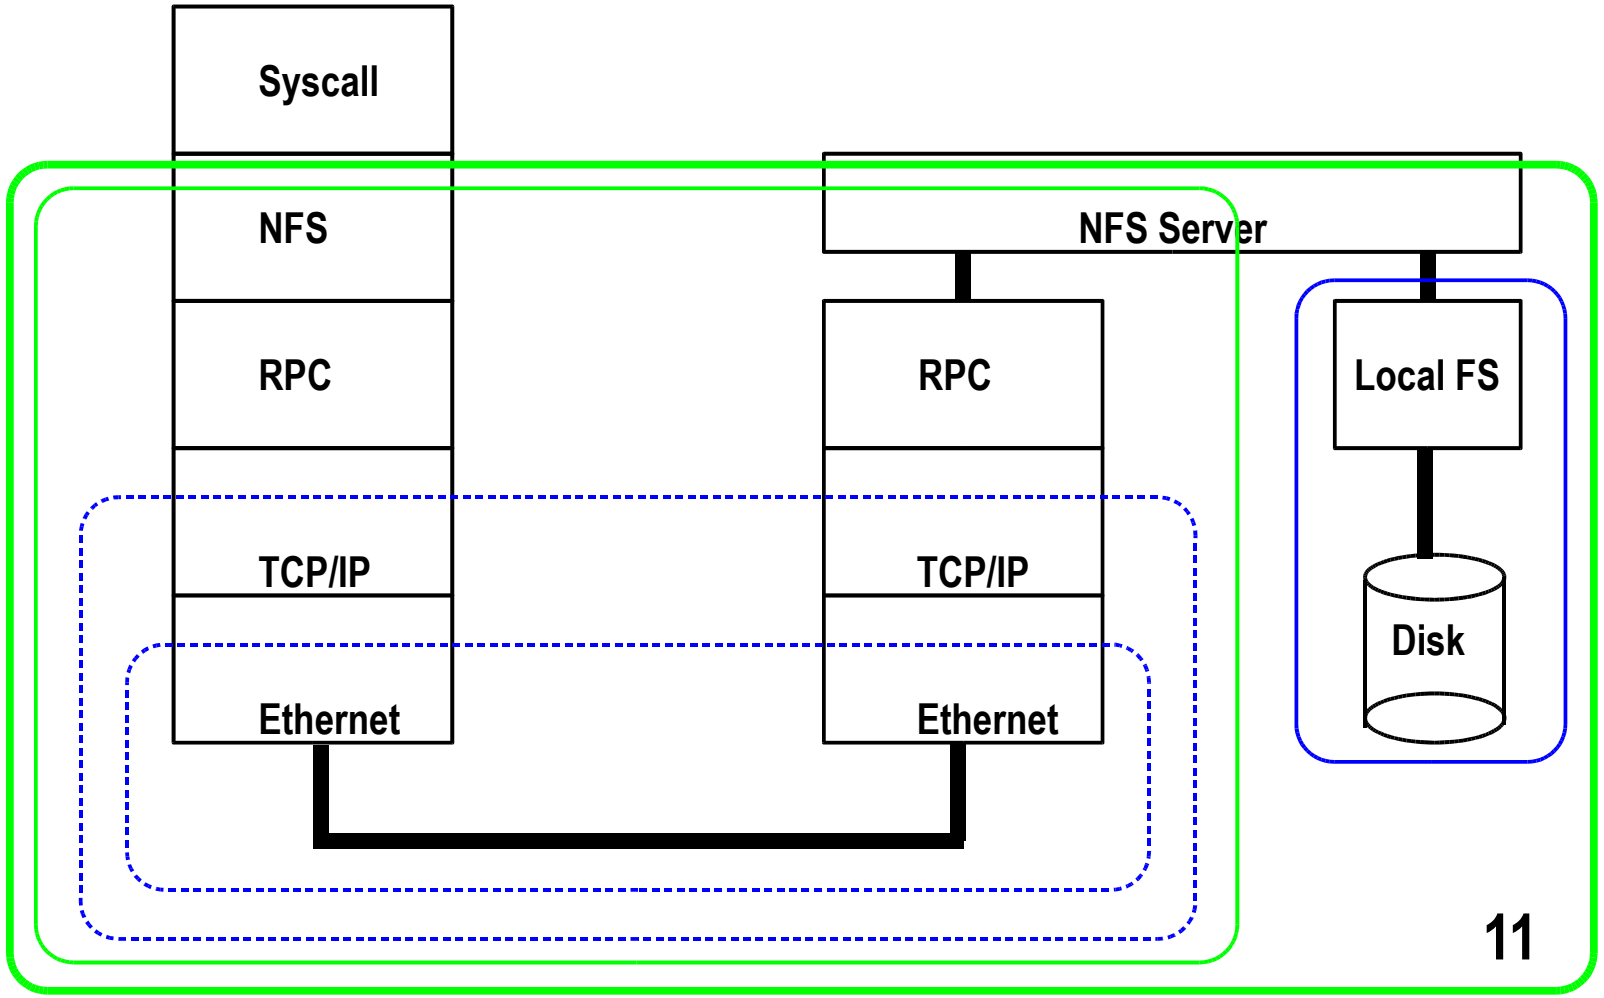

Requirements for NFSv4 Checksums

Protocol Perspective

- Protect against network corruption
- Extend the protection domain to the server storage media wherever possible. Server can optionally verify checksums before passing them to the underlying fs
- Define the “ends” clearly. Client memory to server storage media?
- Is “good enough” an acceptable solution?

Requirements for NFSv4 Checksums

Usability Perspective

- Easy to configure
- Provide choices between checksum algorithms
- Atleast one mandatory to implement checksum algorithm
- Relatively Inexpensive, performance impact should be minimal

Additional Information

- I-D on Checksum Extensions for v4 -

  draft-aggarwal-nfsv4-cksum-00.txt
- Draft related discussion on NFSv4 WG alias
- If still in doubt that TCP checksum is weak, read -

  http://citeseer.ist.psu.edu/stone00when.html

Checksums for NFSv4?

Alok Aggarwal
alok.aggarwal@sun.com