Checksums for NFSv4?

Alok Aggarwal

Member of Technical Staff

Sun Microsystems

alok.aggarwal@sun.com

Agenda

- Data Integrity Domains in a System
- NFS Data Integrity
- Checksums for NFSv4?
- Requirements and Possible Implementation Choice
- Q & A

Data Integrity Domains

System-wide

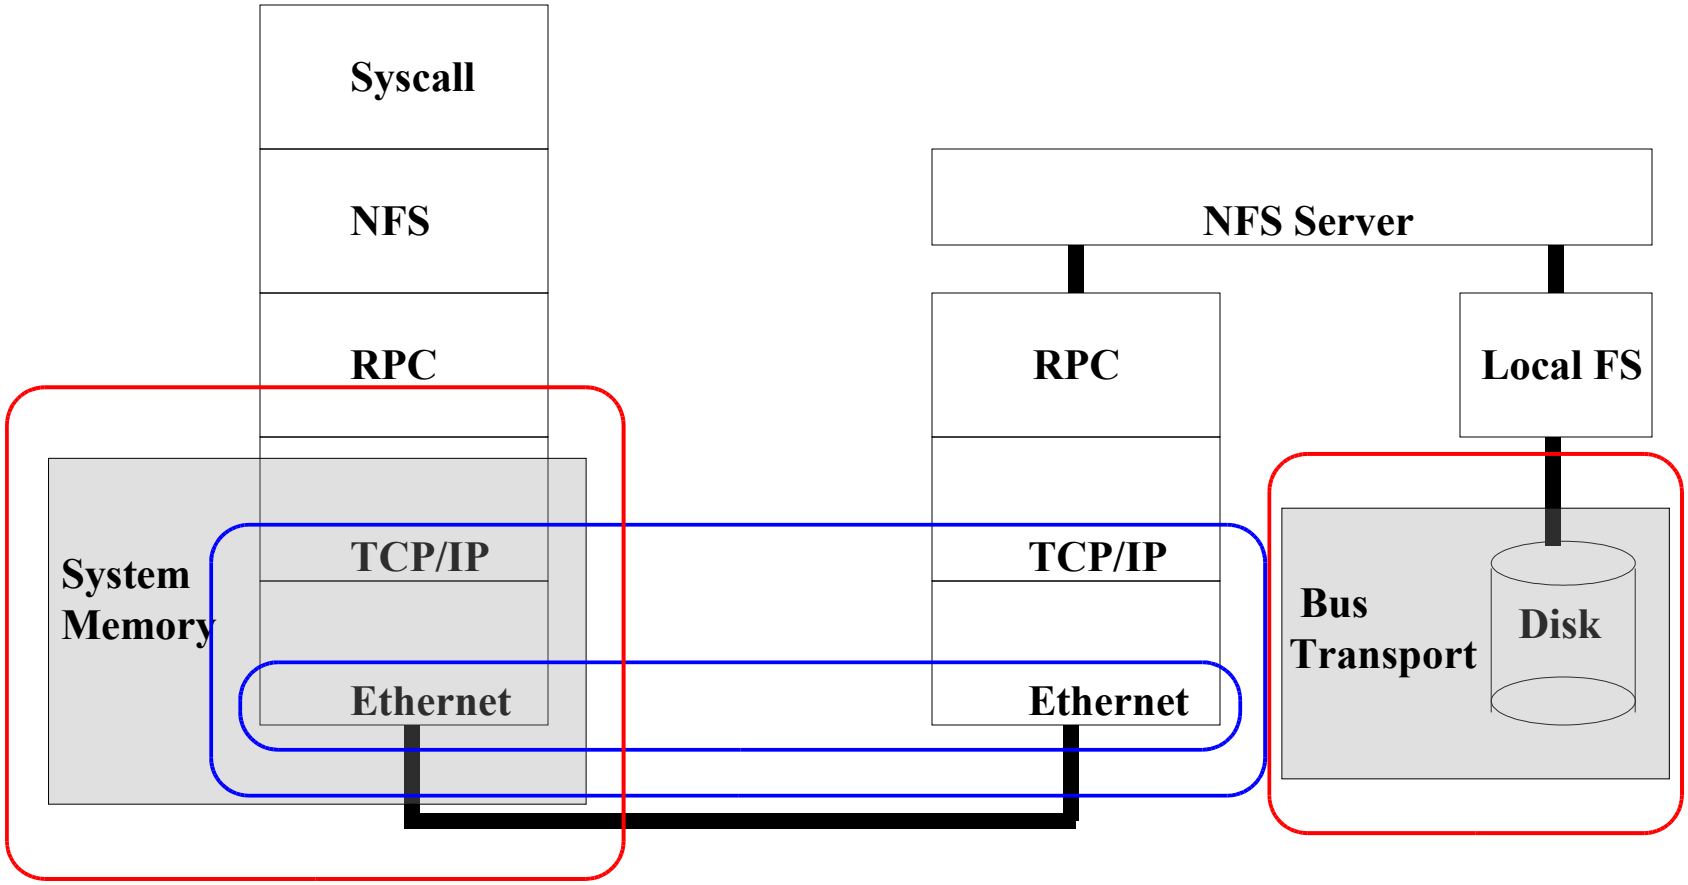

Data Integrity Domains

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 krb5i

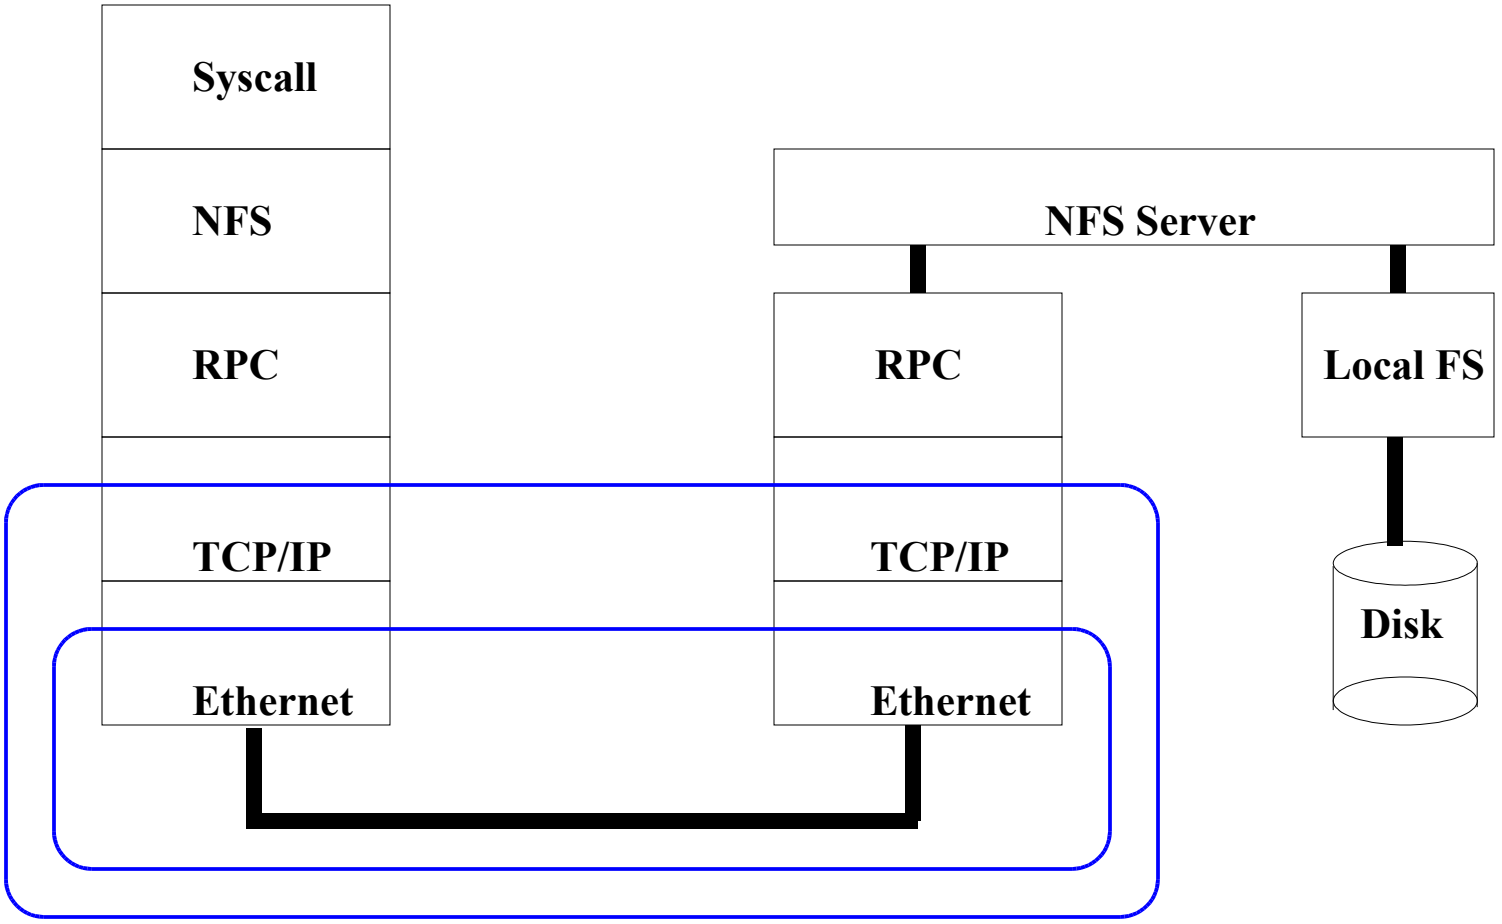

NFS Data Integrity

Without krb5i

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

NFS Data Integrity

Without krb5i

- Reliance on TCP checksums
  - Alternate checksum algorithm optional
    - Not widely implemented
- 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 are we left with?

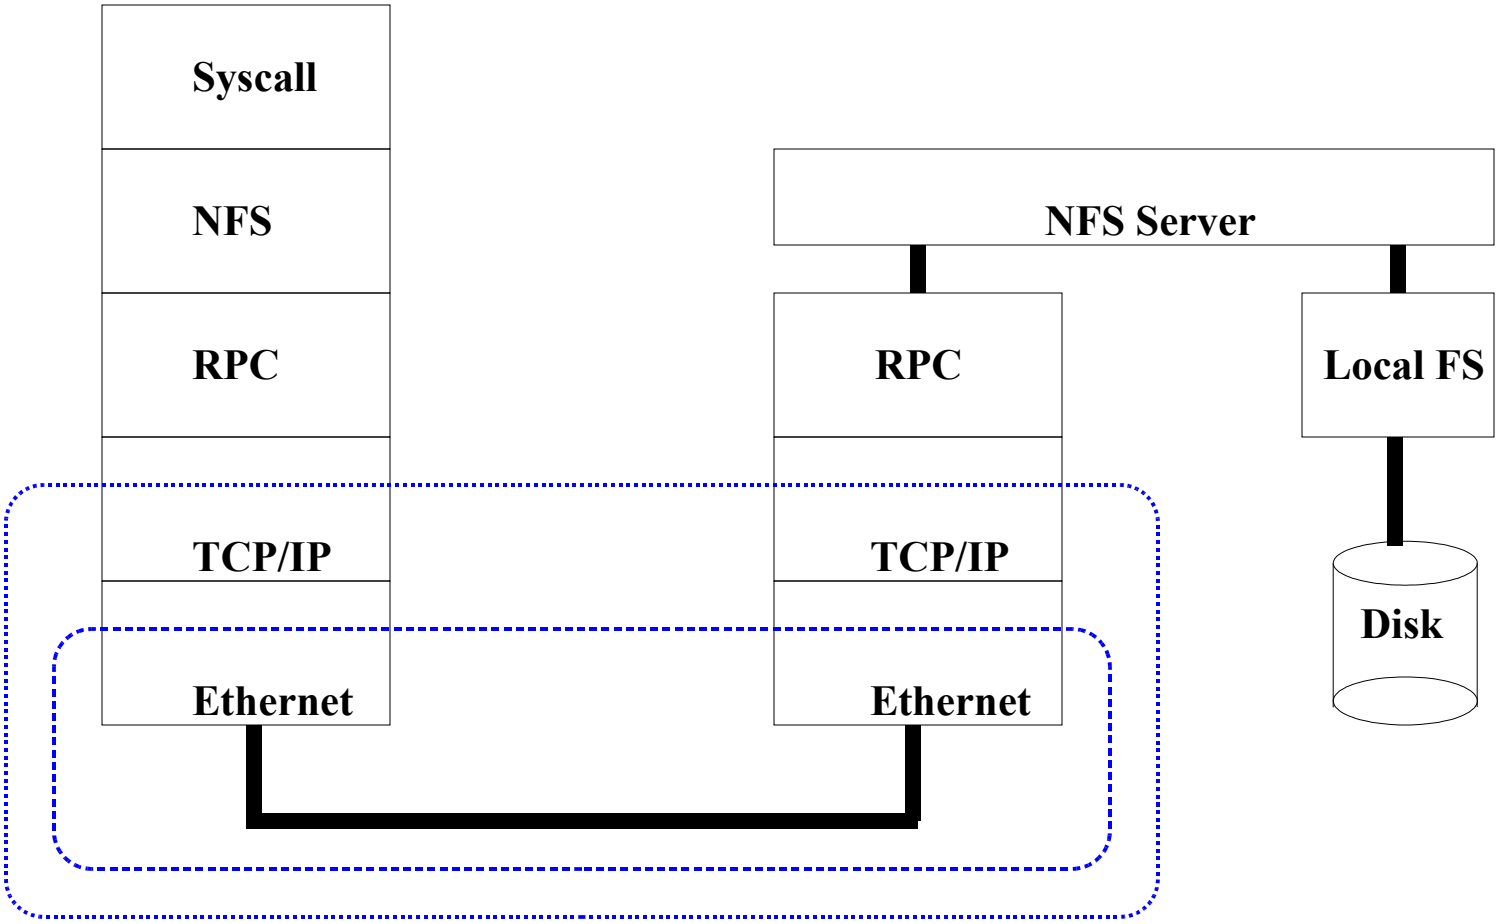

NFS Data Integrity

With krb5i

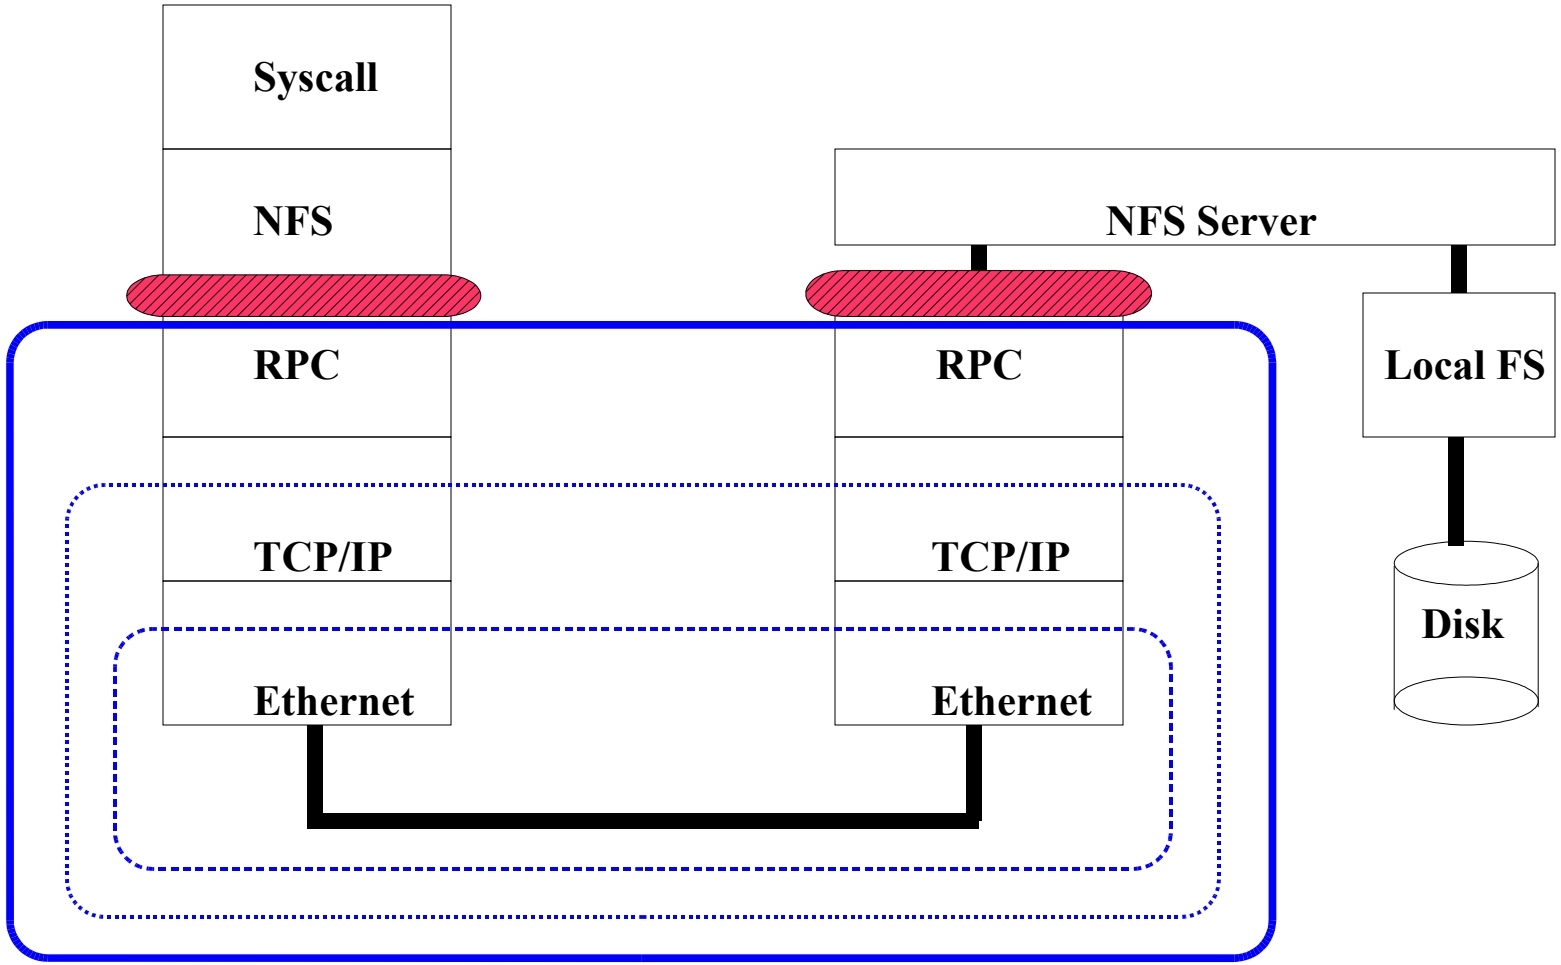

NFS Data Integrity

With krb5i

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

Checksums for NFSv4?

End-to-End Integrity

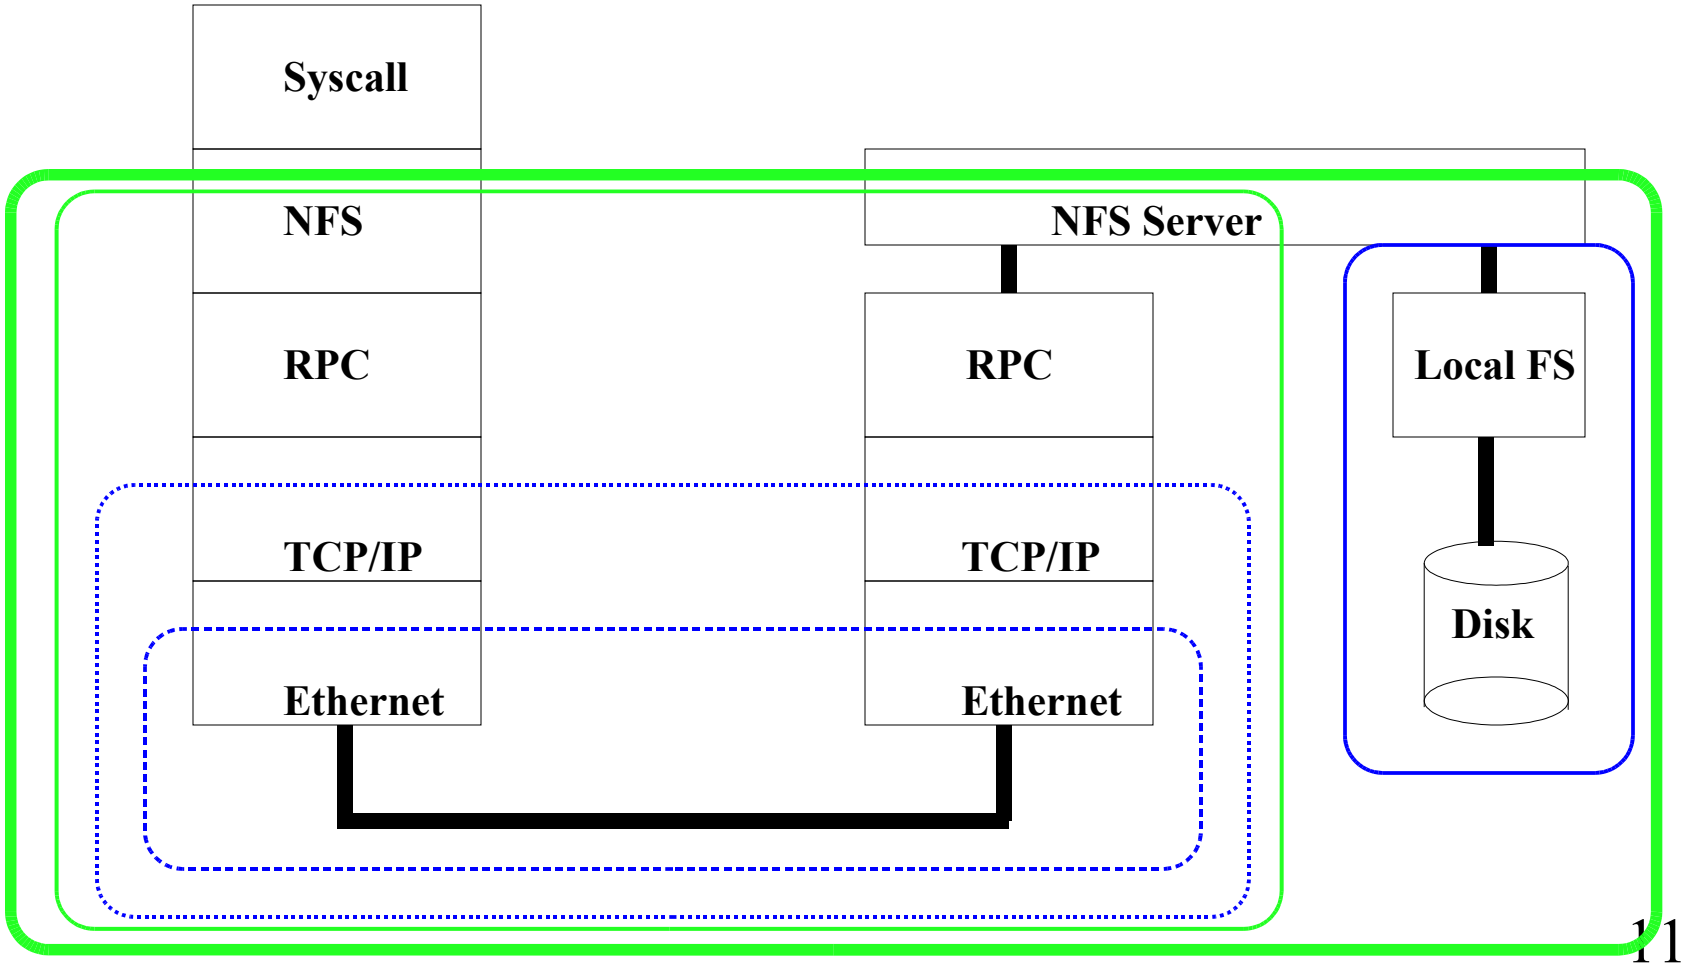

Requirements for NFSv4 Checksums

- Easy to configure
- Provide choices between checksum algorithms
- Relatively Inexpensive, performance impact should be minimal

Implementation Choice

- Additional checksum operations
  - OP_CKSUM to carry actual checksum information, tagged with OP_READ/OP_WRITE
  - OP_CKINFO to negotiate checksum algorithms
- Implemented via NFSv4 minor versions

Questions?