

Caching Requirements

Brent Callaghan

NFS v2 & v3 Caching

- **Caches - most clients use memory**
 - **Little use of disk caching**
 - **High bandwidth, low latency, little caching incentive**
- **Cached items: file data, attributes, symlinks, directories, ACCESS results, filehandles**
- **Close-to-Open consistency for file data & attrs.**
- **Probabilistic caching otherwise**
 - **3 - 30 seconds for data**
 - **30 - 60 seconds for directories**
 - **Varies depending on frequency of change**

Caching & Internet

- **High latency (100's of ms)**
 - **Impact on every file open**
 - **Consistency check is expensive**
- **Low bandwidth (order of magnitude)**
 - **Impact of cache miss much higher**
 - **Encryption reduces bandwidth further**
- **Scalability Issue**
 - **LAN WG server may have 10's - 100's of clients**
 - **Internet server may have 1000's - 10,000's of clients**
 - **Caching is Kind to busy servers**

NFS v4 Caching Design Considerations

- **Close-to-open still sufficient ?**
- **Callbacks to reduce consistency checks (AFS, CIFS)**
- **Leases (NQNFS)**
- **Easy server recovery**
- **Cachable objects: file data, directories, ...**
- **Robust protocol**
- **Proxy caching**
- **Transport assumptions: TCP or UDP. Firewalls**
- **Simple. Easy to implement**

