



**N I C  
F I O  
S D N  
U S F  
S T E  
T R R  
R E  
Y N  
C  
E**

# NFS Performance

Peter Staubach

Senior Staff Engineer

Sun Microsystems, Inc.

[peter.staubach@Sun.COM](mailto:peter.staubach@Sun.COM)



**N I C  
F N D  
S U S  
T R E  
R E N  
C E**

# Agenda

- Performance versus Consistency
- Read/Write Throughput
- Metadata Updates
- Data Copying
- Client Metadata Caching



**N I C  
F N O  
S D N  
U S F  
T R E  
R Y R  
E N  
C E**

# Performance versus Consistency

- "noac" / "actimeo=0"
- Locking
  - File/record
  - Share
- Non-cached I/O
- Delegation



**N I C  
F N D  
S U S  
T R E  
R E N  
C E**

# Read/Write Throughput

- Read ahead values
- Write behind values
- Block sizes
  - TCP
  - UDP
  - RDMA
- Non-cached I/O



**N I C  
F N O  
S D S  
I U F  
N D E  
D T R  
R E  
Y N  
C  
E**

# Metadata Updates

- Local file system
- NVRAM



**N I C  
F N O  
S D S  
I U I  
N D R  
D U E  
T S T  
R R E  
Y N C  
E**

# Data Copying

- XDR overlays
- RDMA
  - RDDP
  - Infiniband
  - VI
- Inference Engines



**N I C  
F N O  
S D U  
I N D U S T R Y  
C O N F E R E N C E**

# Client Metadata Caching

- Directory Name Lookup
- File access privileges
- Readdir response