



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

Network Adapters for NFS at 10Gb and Beyond

Tom Hotchkiss

VP, Engineering

Emulex Corporation

Tom.Hotchkiss@emulex.com





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

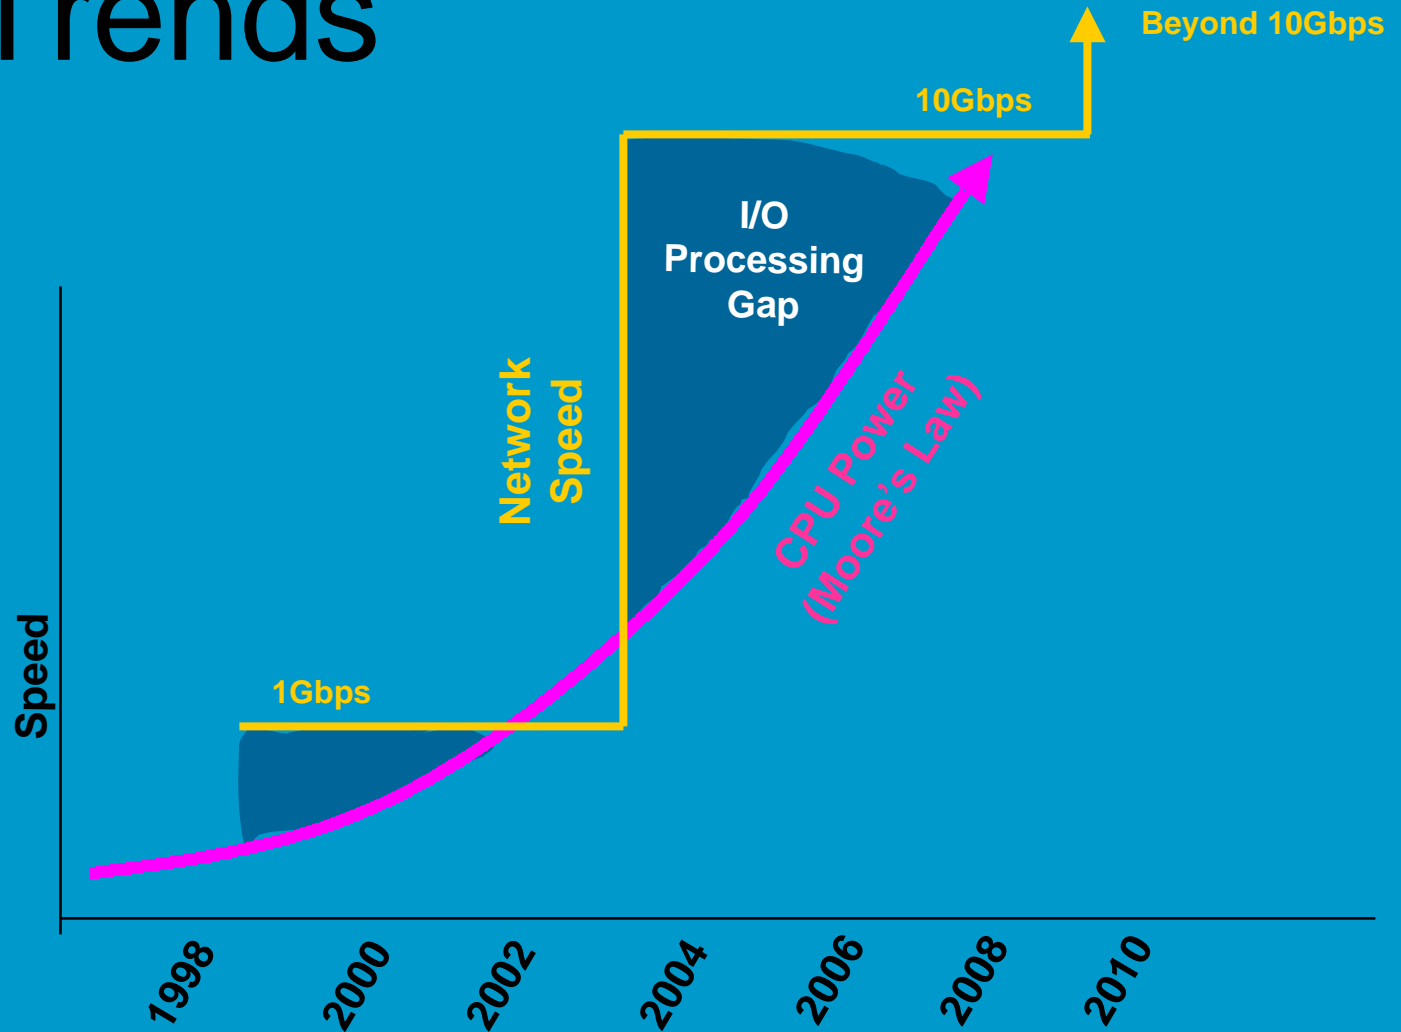
Overview

- **Networking Issues at 10Gb and Beyond**
 - Protocol Processing
 - Memory Bandwidth
- **The Role of the Network Adapter**
 - Standard NIC
 - TCP Offload Adapter
 - RDMA Adapter
 - NFS Acceleration Adapter
- **Multi-Protocol Network Adapters**
 - Flexible Solutions



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

CPU and Network Trends





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


10Gb Protocol Processing

- Requires a 10x Increase in I/O Processing
- Performance Limited by Host
 - Protocol Processing
 - Interrupts
 - Buffer Copies



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

Standard NIC





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


Standard NIC

- Will not Allow 10Gb Performance
 - TCP Processing on Host
 - Buffer Copy on Receive
- Overloaded Host CPU and Memory
 - Poor Application Performance



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

TCP Offload Adapter





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

TCP Offload Adapter

- TCP/IP Processing on Adapter
 - Not Just TCP Assist
- Eliminates Host TCP Processing Requirements
- But...
 - Still Requires Buffer Copy on Receive

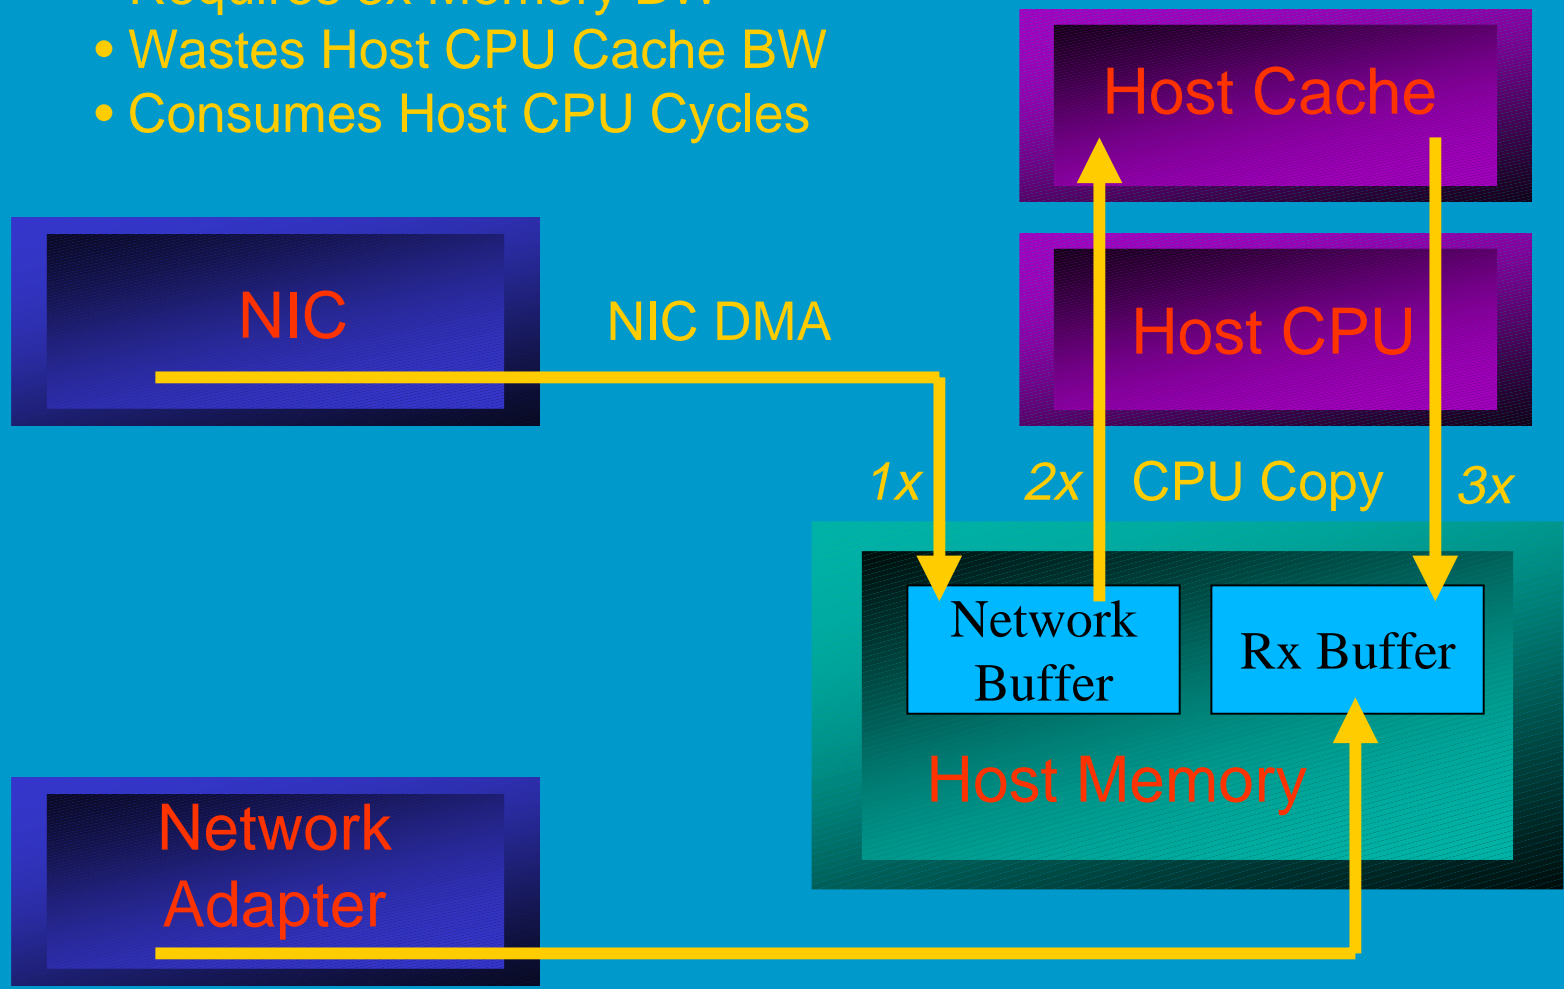


N I C
F N O
S D N
U S F
T R E
R Y N
C E

Direct Data Placement

Buffer Copy

- Requires 3x Memory BW
- Wastes Host CPU Cache BW
- Consumes Host CPU Cycles



Network Adapter Direct Data Placement (DDP)



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

Memory BW Requirements

Wire Speed	Network BW (MB/sec)	Memory BW (MB/sec)
1 Gb/sec	125	375
10 Gb/sec	1,250	3,750

**Memory BW = 3x Network BW due to Buffer Copy*



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

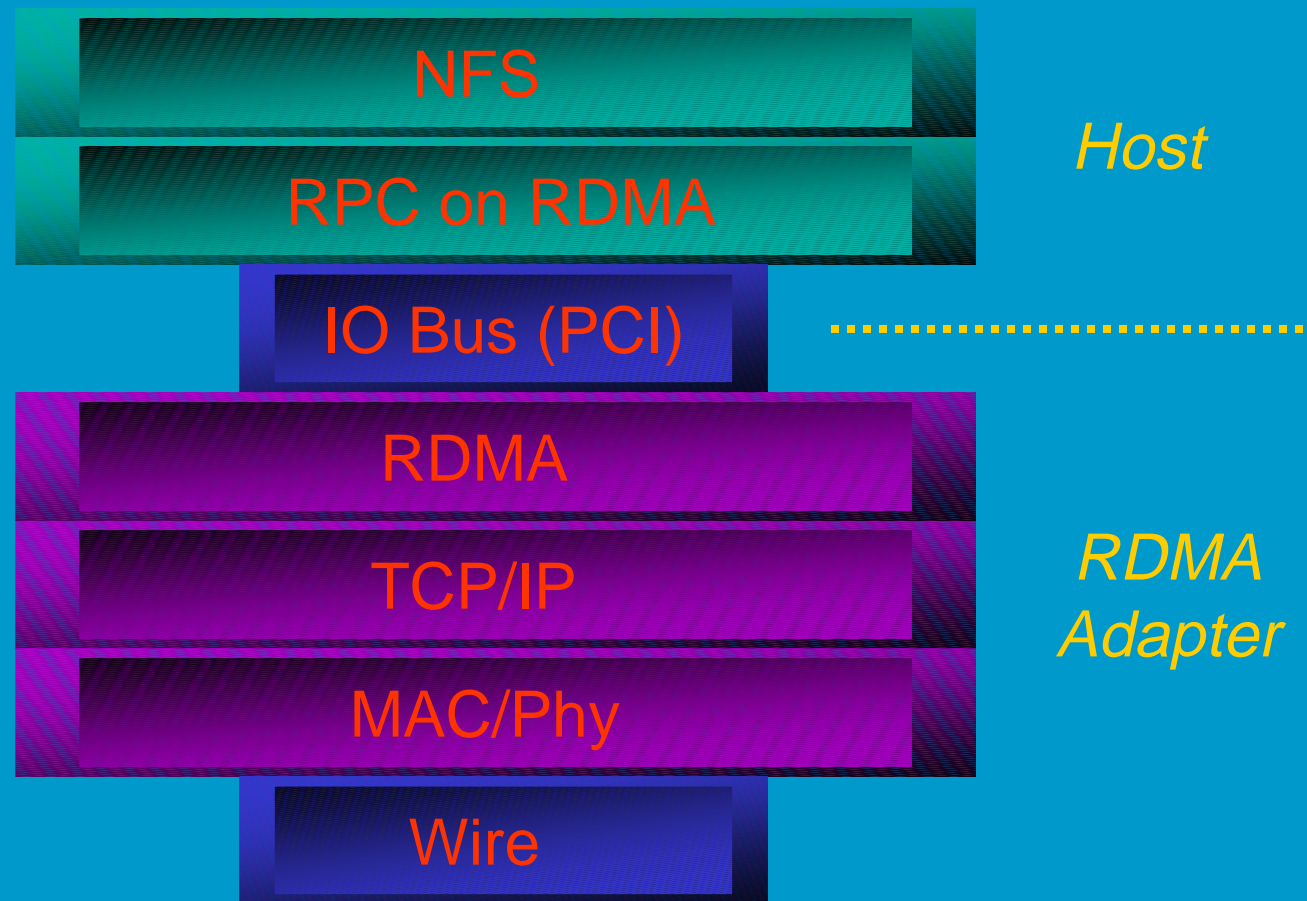
Zero Copy Receive with Remote DMA (RDMA)

- Direct Data Placement in Receive Data Buffer
- Requires 1/3 the Memory BW
- Eliminates Host CPU Copy Overhead



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

RDMA Adapter





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

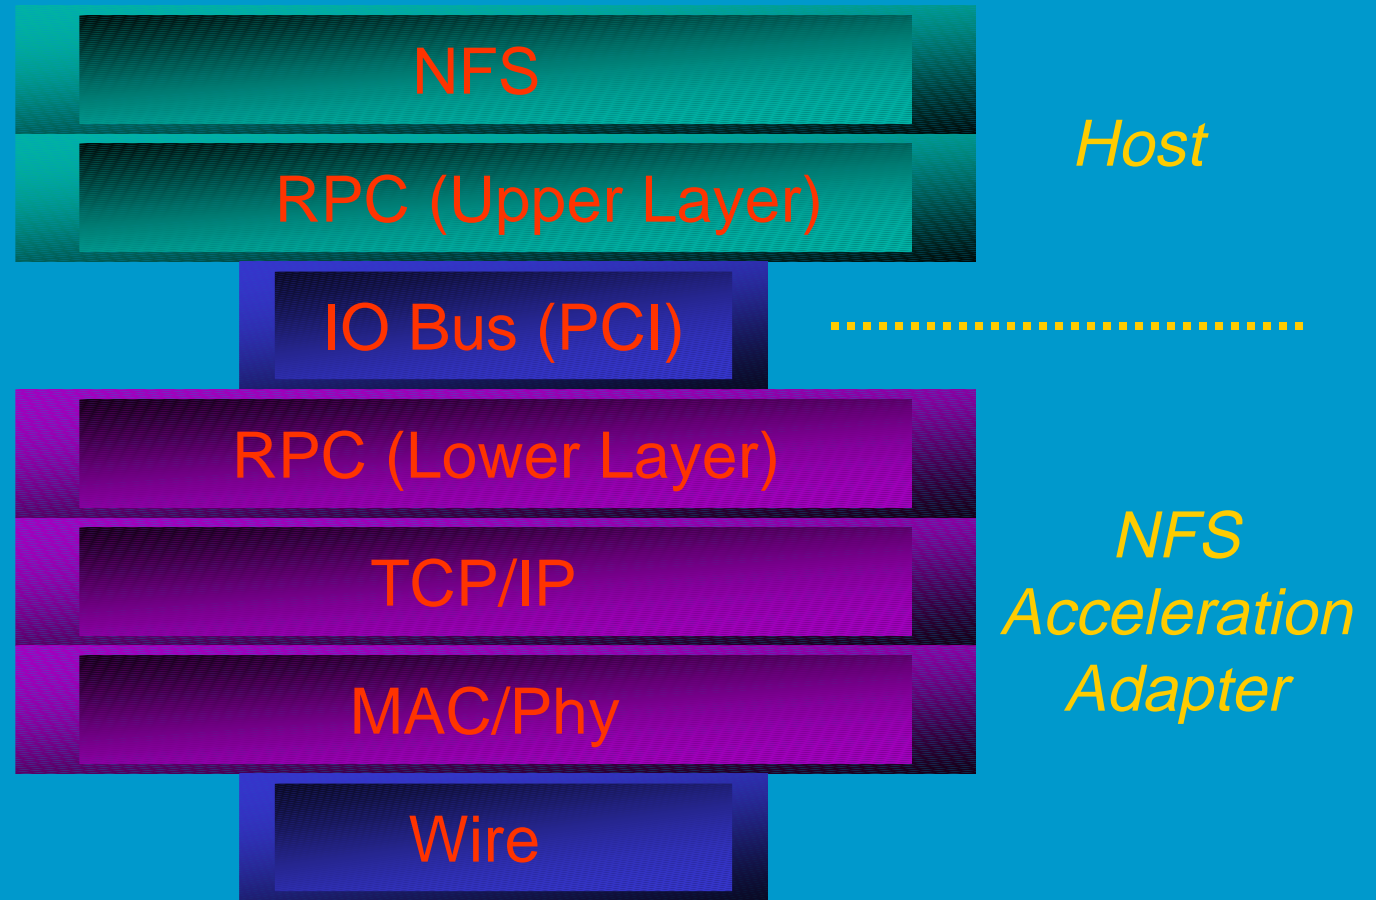
RDMA Adapter

- Modified RPC Layer Supports RDMA Protocol
- Transparent to Applications
- High Performance
 - Zero Copy Receives
 - TCP Protocol Offload
- Requires a New Transport Protocol
 - RDMA Protocol Supported by Server and Client



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

NFS Acceleration Adapter





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

NFS Acceleration Adapter

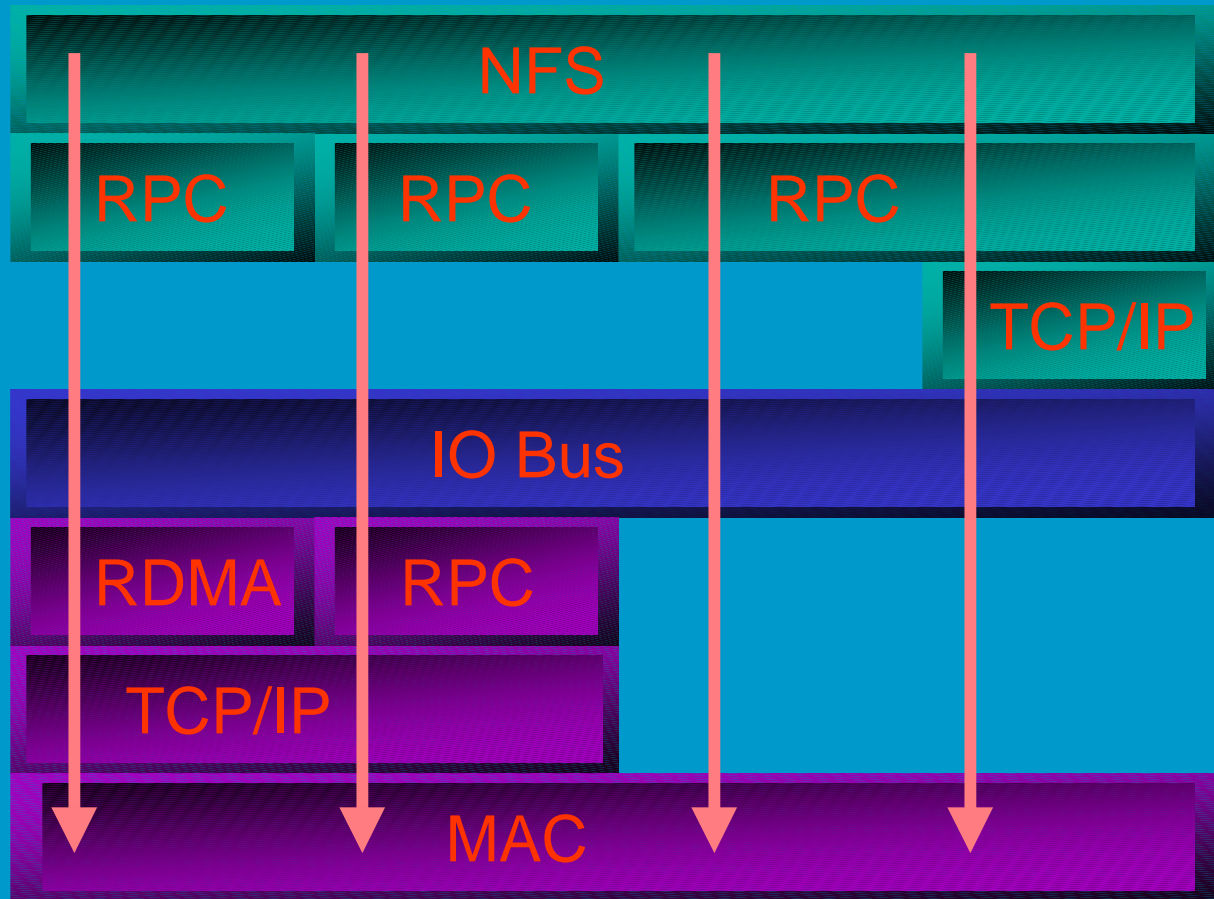
- Modified RPC Layer
- Direct Data Placement
- Transparent to Applications
- High Performance
 - Zero Copy Receives
 - TCP Protocol Offload
- No New Protocols
 - Single Ended: Client or Server can Implement Independently



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

Multi-protocol Adapter

RDMA Adapter NFS Adapter TOE Adapter Standard NIC



Host

Network Adapter

*TCP Offload
Zero Copy
RDMA*

*TCP Offload
Zero Copy
TCP/IP*

*TCP Offload
Buffer Copy
TCP/IP*

*Host TCP
Buffer Copy
TCP/IP*

*Green=Host
Purple=Adapter
Blue=Transport*



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

Multi-Protocol Adapter

- Intelligent Network Adapter Supports Multiple Protocols
 - Standard NIC Mode
 - TCP Offload
 - NFS Acceleration
 - RDMA



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

Summary

- Standard NICs will not Support 10Gb Performance
- Intelligent Network Adapters Needed
 - TCP Protocol Offload
 - Zero Copy Receive via Direct Data Placement
- Multi-Protocol Adapters Provide Flexible Solutions