



Education

NAS and iSCSI Technology Overview

Wolfgang Singer, IBM Austria

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

- ◆ **Marketplace Directions**
- ◆ **General Differences of DAS, SAN, NAS, iSCSI**
- ◆ **Uses of NAS**
- ◆ **NAS Gateway/Head Architecture and Solutions**
- ◆ **iSCSI Status and Future**

Reference to Other Tutorials



Check out SNIA Tutorial:

Server and Storage Consolidation with iSCSI Arrays

Large Scale File and Block Serving for Clouds

Case Study: iSpy iSCSI

- **CIO/CTO**
- **Consultants**
- **Systems Architects and Technologists**
- **Systems Administrators**
- **Technical Specialists**

- **Trends**


- ◆ Storage growth 20-150% per year (depending on the industry)
- ◆ Lack of skilled personnel

- **Industry Directions**


- ◆ Pool storage for flexibility
- ◆ Intelligent storage to reduce administrative costs
 - > Simpler (less skills required)
 - > Automation (less effort)
- ◆ Continuous availability
- ◆ Evaluation of Cloud Storage offerings

Connectivity for Storing Data

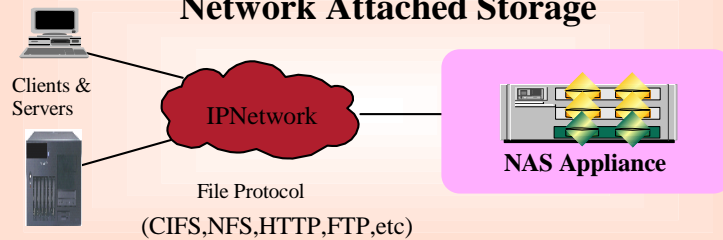
Direct Attached




SAN Attached




Network Attached Storage



NAS Gateway




iSCSI Attached



Uses of NAS

"Storage"

1. Server Storage



Shared Pooled Intelligent Storage

Business Problem:
 Cost, skills, downtime for
 adding storage to servers

2. Client Storage



Shared Pooled Intelligent Storage

Business Problem:
 Backup/Recovery, adding
 storage, sharing files

"File Server"

3. File Server Appliance



CIFS, NFS, HTTP, HTML, XML,
 RFC 1852, Multi-Media

Business Problem:
 Availability, performance, skills
 for serving files


Unified Storage



Multi-Protocol

- NAS
- IP SAN
- FC SAN
- FCoE (near future)

Traditional Arrays
Only Offer
FibreChannel

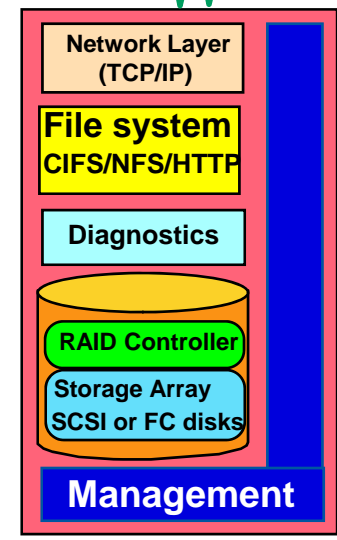



What is Network Attached Storage?

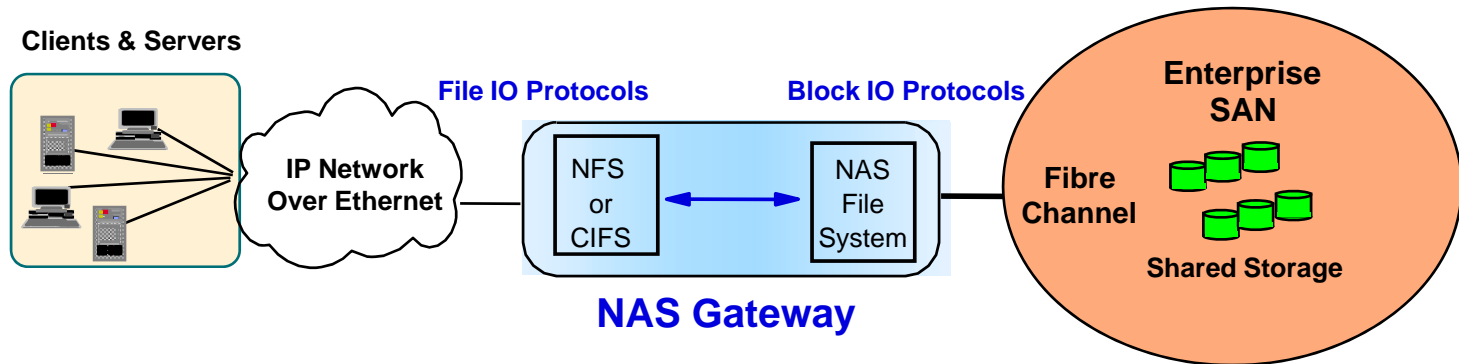
Task-optimized, high performance **storage appliance directly attached to IP networks**, providing “File Serving” to clients and servers **in a heterogeneous environment**

- **Preloaded file system that provides heterogeneous file sharing**
 - Windows (CIFS), UNIX (NFS), Web (HTTP), Novell, FTP, Apple FP
- **Installation/Configuration of software for Clients/Appliance**
 - Requires minimum IT skills to maintain / install
- **Scales from GBs to TBs**
- **Connects to IP network, mainly running over Ethernet**
- **Management software**
 - Manage & Setup from remote location
- **Diagnostic software**
 - Predictive Failure Analysis, Alerts
- **Fault Tolerant Features**
 - Dual, Redundant, Hot Swap Components
- **Data Protection Technology**
 - Data Protection with RAID, & Backup to Disk & Tape

- **Preloaded**
- **Preconfigured**
- **Pre-Tuned**
- **Pre-Tested**



NAS Gateway/Head Configuration



- **Gives the combined benefits of NAS and SAN**
 - ➔ NAS flexibility and ease of use
 - ➔ SAN scalability on the IP network
- **Increases the reach of Fibre Channel storage devices**
 - ➔ Extends beyond topology limitations of Fibre channel
 - ➔ Allows FC devices to be used on the IP network
 - ➔ Connectivity to switches, directors, RAID controllers and disk arrays
- **Leverages the value of Fibre Channel investment**
 - ➔ Reduces access costs to Fibre devices
 - ➔ Allows access to underutilized SAN storage
 - ➔ Enables heterogeneous file serving on SAN storage devices

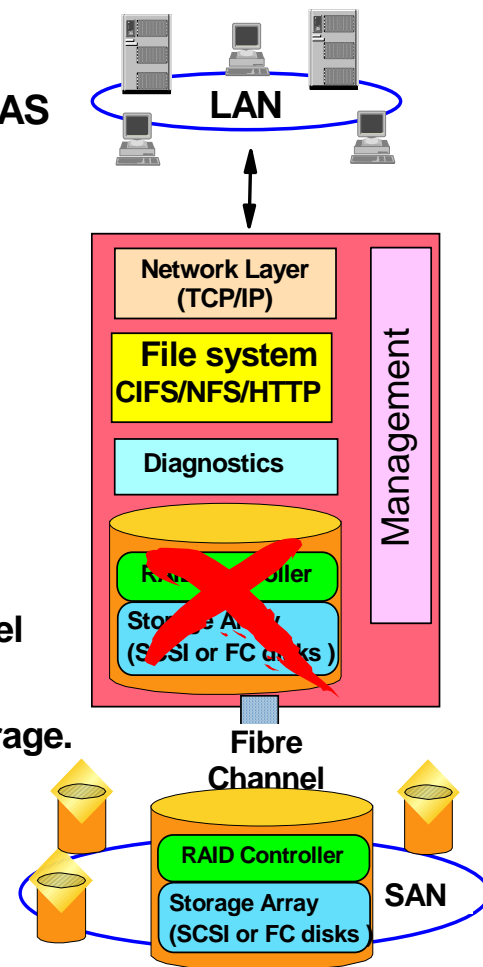
NAS Gateway/Head Architecture

- **NAS Gateway offers benefits and characteristics of NAS**

- Connects to IP networks
- Performs as a file server
- Heterogeneous file sharing
- Data protection
- Clustering and failover features

- **NAS Gateway is a NAS Appliance with one exception**

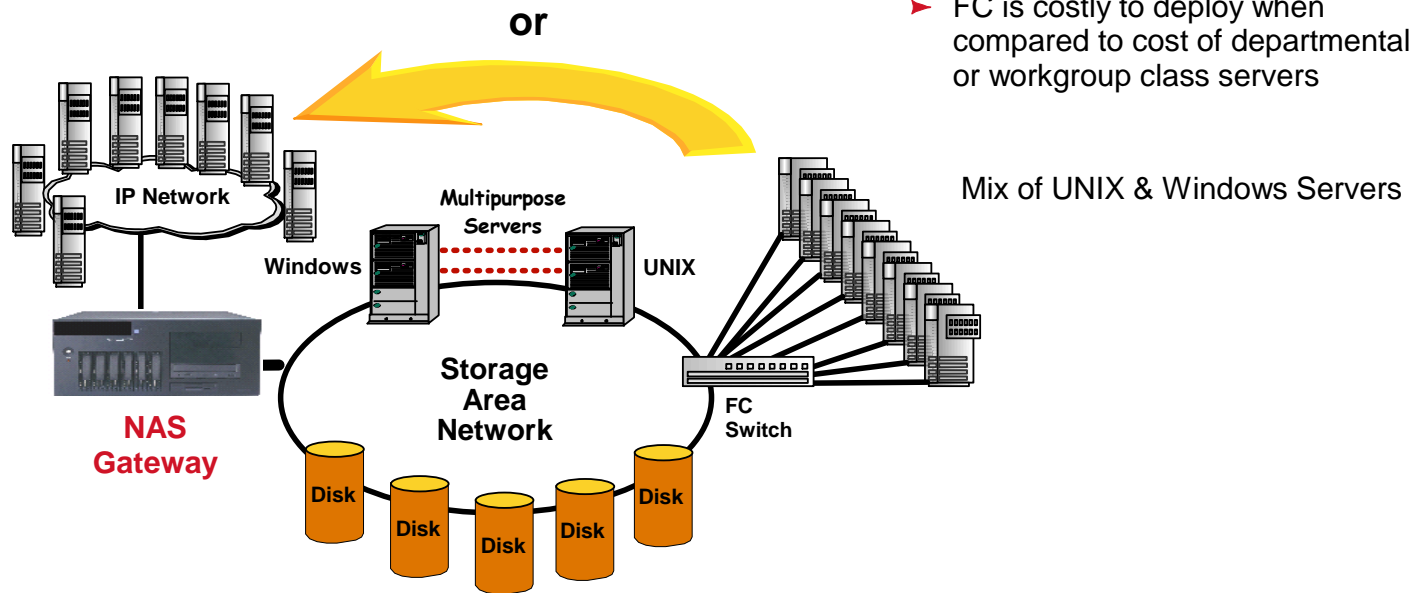
- NAS Gateway supports direct attachment to Fibre Channel storage or connection to a storage device across a SAN.
- NAS Gateways do not have integrated disks for data storage.



NAS Gateway/Head Solution

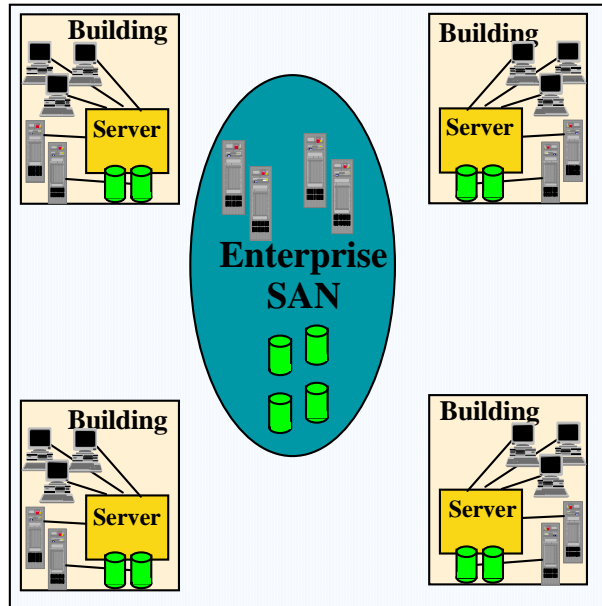
- Why purchase additional integrated NAS storage when you already have SAN storage
- Capitalize on your storage investment and purchase NAS functionality without the cost of additional NAS storage

- FC has distance limitations
- FC is costly to deploy when compared to cost of departmental or workgroup class servers

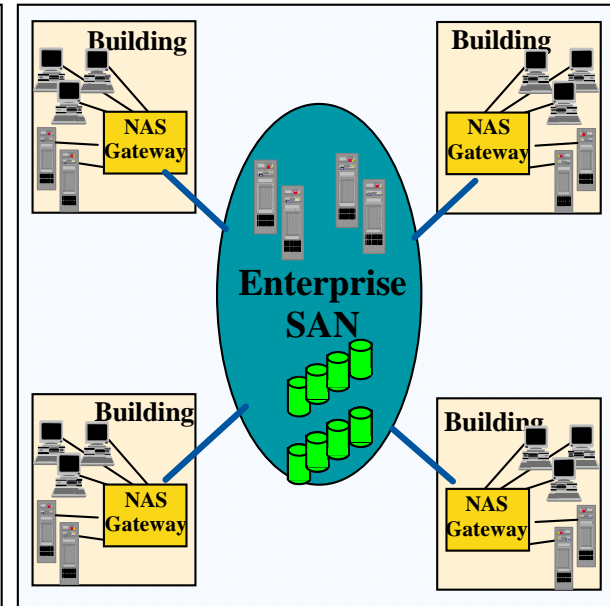


Customer Example

Islands of Storage



SAN/NAS Integration



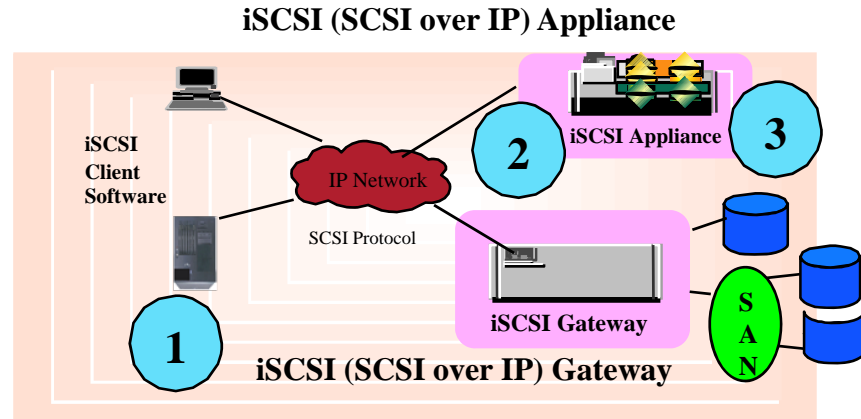
Benefits

- **Storage Consolidation**
- **Cost effective connection to SAN capacity**

iSCSI - An Alternative Technology

SCSI over IP Networks "SAN" with IP fabric

Two industry Approaches:



iSCSI Appliance (Native iSCSI w/ embedded storage)

iSCSI Gateway (Implemented on Fibrechannel Switch or Standalone Appliance, w/o embedded storage)

1

iClient (initiator) code reroutes SCSI commands over IP network

2

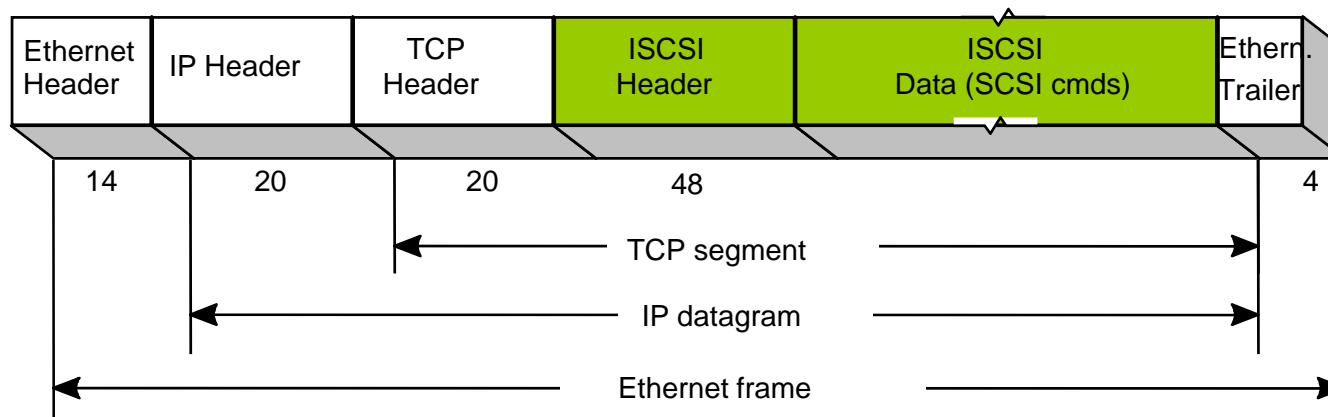
iSCSI target code receives SCSI commands from IP network.

3

SCSI commands then either routed directly to embedded storage (iSCSI Appliance) or routed to FC SAN (iSCSI Gateway)

What is iSCSI ?

- Enables the transport of **Block I/O** data over **IP Networks**
- Operates on top of **TCP** through encapsulation of **SCSI** commands in a **TCP/IP** data stream
- Transport of **iSCSI** mainly over **Ethernet (LAN/Metro)**; **WAN Protocols (PPP, Frame Relay..)** possible as well



Remark: not all TCP segments carry an iSCSI header


What is iSCSI ? - cont.

- **Transport for SCSI Commands**
- **End to End protocol (between Initiator and Target)**
- **Can be implemented on Desktops, Laptops and Servers**
- **Can be implemented with current TCP/IP Stacks**
- **Can be implemented completely in a HBA**
- **Can use existing routers/switches without changes**
- **Transport includes Security as a base concept**
 - Initiator and Target (RADIUS) authentication
 - Uses CHAP, SRP, Kerberos, SPKM
 - Enabled for IPSec Encryption, Digests and anti-Replay
- **Defines Discovery as a basic element**
- **Defines process for Remote Boot, as a basic element**
- **Excellent SAN solution for servers with less thruput demand today**

Benefits of Ethernet in Storage

- **Ethernet and TCP/IP are widely deployed and dominant**
 - Not just the Fortune 1000 (as is Fibre Channel)
 - Well understood technology
 - Low acquisition cost
 - Unlimited distance (with TCP/IP)
 - Companies do not have to retrain for TCP/IP networks
- **Ethernet is a scalable technology, with 10/100/ Mbps, 1/10 Gbps**
 - 40 and 100 Gbps are on the roadmap
- **Allows the creation of a single physical network using familiar standards**
- **VLAN's maybe used for separating storage traffic from intranet traffic**
- **Brings Interoperability & Ethernet economics to storage**
- **Enables remote data replication and disaster recovery**
- **Faster deployment than with FC**

iSCSI Connectivity



- Initiators and targets can be implemented in H/W or S/W
- Session between initiator and target
 - ◆ One or more TCP connections per session
 - ◆ Login phase begins each connection
- Services (e.g., authentication, security) negotiated during login
- TCP Protocol provides
 - ◆ Delivery of SCSI commands in order
 - ◆ Recovery from lost connections (and lost packets)

TCP Overhead

❖ TCP Processing

- ◆ Every TCP connection that is part of an iSCSI session has processing overhead potential
 - › Connection setup / teardown
 - › TCP state machine:
 - Acknowledge, Timeout, Retransmission
 - Window management
 - Congestion Control
 - › TCP segmentation
 - › IP fragmentation
 - › Checksum calculations
- ◆ TOEs help here very much (as do faster processors)

❖ **Gigabit Ethernet and TOE not mandatory requirements for iSCSI;**

Many servers are still 100 Mbps connected

- ❖ **Therefore today, iSCSI mainly used in entry level to mid size servers**
- ❖ **High End Servers may require 10 Gigabit Ethernet and TOE**

- **TCP/IP Off-load Engines (TOEs), help at Gigabit wire speed NICs**


- ◆ Required to be competitive with Fibre Channel
- ◆ 1 Gbps links will NOT require full integrated ASIC
 - › Increase in system processor speeds might be sufficient
- ◆ Different Implementations: TCP/IP or TCP/IP and iSCSI offload;
- ◆ Defacto standard software TOE support with “TCP Chimney”

Several NAS's already implemented TOEs

- ◆ 1 Gbps iSCSI cards available
 - › Some with ASIC Chips that includes a TOE and MAC
 - › Some with ASIC TOE Chip that include iSCSI and MAC
 - These chips can replace FC chips in Storage Controllers
 - › Others just use MIPS or PowerPC processors with SW TCP/IP Stacks
- ◆ 10 Gbps cards available
 - › Full integrated ASIC Chips required here

Use Jumbo Frames → 20-50% performance increase


iSCSI & TOE Adapters




iSCSI Deployments

Same HW Configurations as NAS Workgroup, Departmental, & Enterprise (Appliances and Gateways)


Independent
iSCSI
Deployment



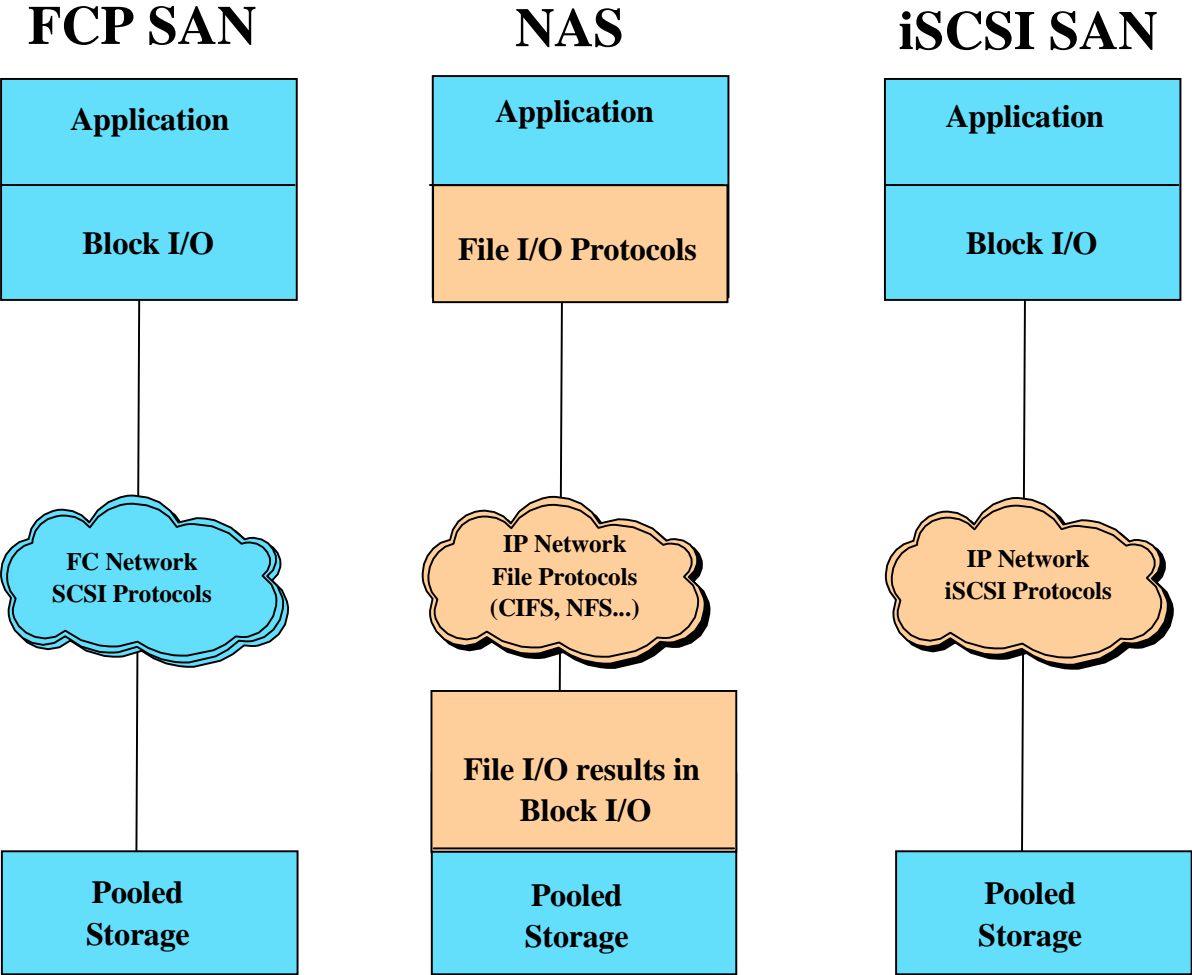
Extending
the SAN



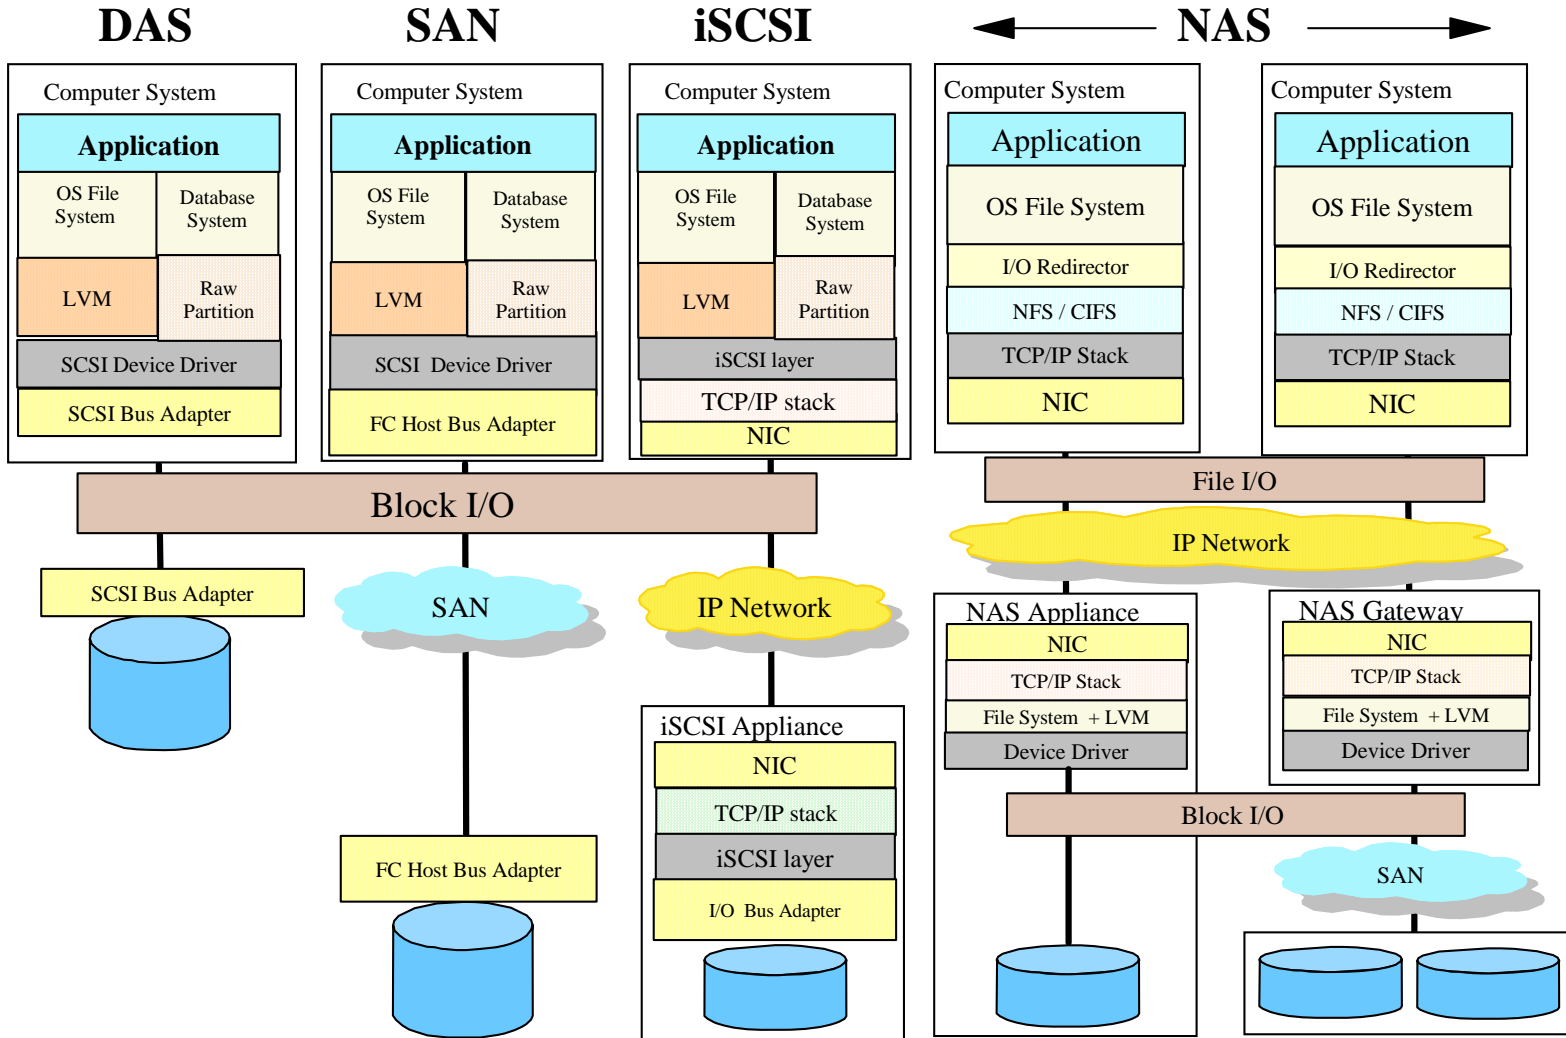
In
Combination
with NAS/FC



Application Protocol Support



Transporting Application Data



- **Key Enhancements**
 - ◆ **Additional Ease of Use**
 - > Automatic Discovery
 - > Automatic Configuration
 - ◆ **National Language Support**
 - ◆ **Enhanced Hardware**
 - > iSCSI Native Host Bus Adapter
 - > iSCSI “features” on high end storage hardware
 - > “Dual Dialect” - NAS and iSCSI support on one Box
 - > 10 Gb Ethernet for iSCSI Appliance/Gateway
 - > iSER – iSCSI Extension to RDMA

- **First iSCSI Plugfest in July 2001**
- **Promontory Summit iSCSI demo in September 2001**
 - ◆ iSCSI data transfer between east- and westcoast
 - › pioneered by 8 companies
- **iSCSI ratified by IETF in February 2003 (RFC 3720)**
- **Strong increase of iSCSI implementations in 2009**
 - ◆ Lots of new products (including Unified Storage), lots of vendors
 - ◆ Many locations will begin to install in many areas
 - ◆ 10 Gigabit products introduced
 - ◆ Also 10 Gigabit Ethernet volume Shipments for Campus Backbones and Host NIC's

- Please send any questions or comments on this presentation to SNIA: tracknetworking@snia.org

Many thanks to the following individuals
for their contributions to this tutorial.

SNIA Education Committee

Wolfgang Singer
David Dale
Ahmad Zamer
John Hufferd
Walter Dey
Elaine Silber
Paul Massiglia
Joe White