


IP Storage Concepts



NAS & iSCSI Fundamentals



IBM TotalStorage™

Agenda


- What Business Problem Needs to Be Solved
- Storage Directions - Marketplace Trends
- Fundamental comparisons ... NAS & SAN
- Network Attached Storage (NAS) Fundamentals
 - Benefits
 - Solutions
- NAS Gateway Fundamentals
 - Benefits
 - Solutions
- iSCSI Fundamentals
 - Benefits
 - Solutions
- Summary

Business Problem



➤ Each server or client has its own dedicated storage


- Storage investment is tied to specific server
- Once server is replaced, storage investment is lost
- Migration of data to new server & new storage is required
- Underutilized storage can not be shared by other servers or clients
- Data in storage can not be accessed by other servers or clients
- Storage must be managed as individual entities
- No centralized management or data protection



Business Solution

► Pooled Storage on a network!

- Investment Protection
- Storage Consolidation
- Data Sharing and Access
- High Utilization of storage
- High Scalability
- Centralized Management
- Data Protection
- Disaster Tolerance



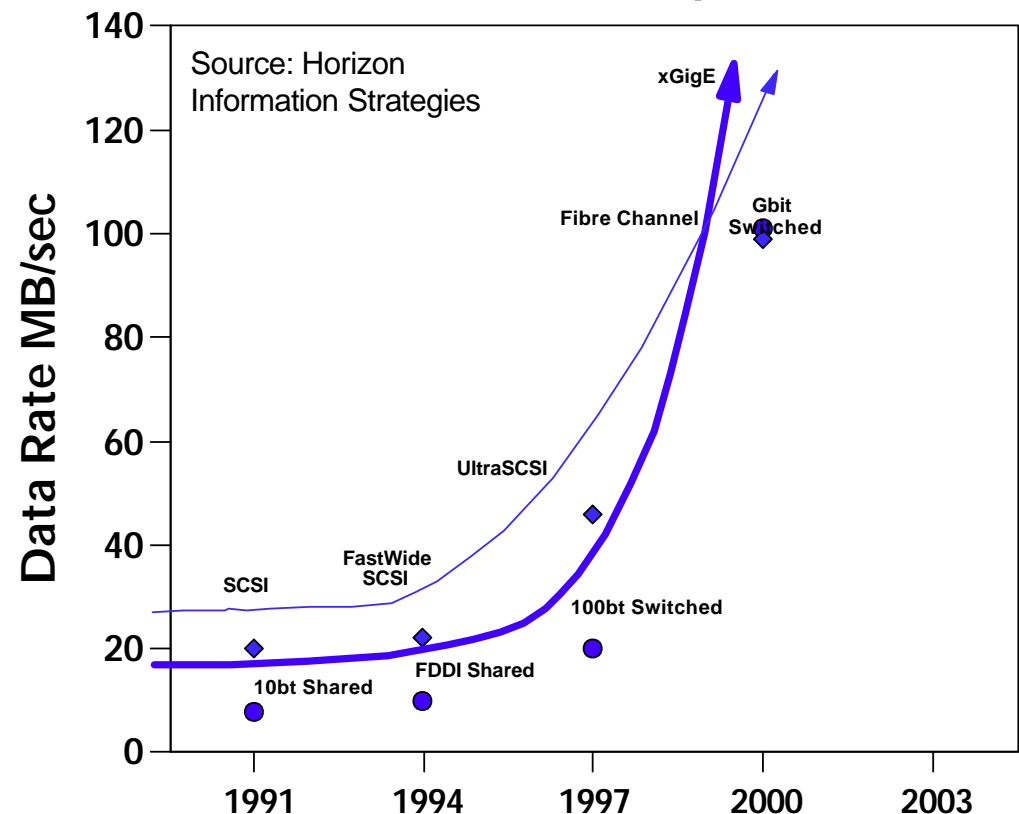
Provides Ultimate Agility & Efficiency

IP Storage Networks

For all these reasons, IP networks may someday accommodate a large portion of the world's storage networking needs


- **Speeds of Ethernet networks are eclipsing Fibre Channel**
 - 10 Gbit ethernet on the horizon
- **Much more R&D is pouring into Ethernet than Fibre**
- **Giga Group forecasts that IP based SANs will dominate (vs.. Fibre Channel) in 5 to 7 years.**
- **Ethernet technology is less expensive than Fibre Channel**
 - Switches, Hubs, infrastructure, IT skills
- **IP is more prevalent than FC**
- **FC may be viable in large enterprise environments**
 - But needs substantial investment
 - Not viable in most other business environments...small, midrange, departmental, workgroup, startup, emerging businesses
- **TCP/IP inherently adds processing latency and requires higher CPU cycles**
 - But higher engine processor speeds accommodate the demand for increased CPU cycles

Network vs. I/O Speed





Growth in Network Storage



Source IDC 9/00

Two Types of Storage Networks


1. Storage connected directly to IP Networks
2. Storage connected directly to Fibre Channel Networks

"Versus" may be giving way to "Convergence"
One picks up where the other's weakness leaves off.

Storage Network Comparison

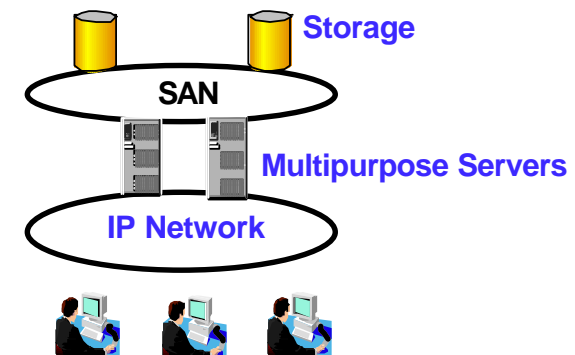
1. Network Attached Storage (NAS)

- Task-optimized, high performance storage appliance directly attached to IP networks, working independently of multipurpose servers providing "File Serving" services to clients and servers in a heterogeneous environment
- Stand-alone Appliance
- Quick to install
- Requires less IT skills
- Optimum Usage for File I/O operations
- Ethernet connections - TCP/IP protocols
- Typically less expensive (\$/MB)
- Performance & Scalability Considerations



2. Storage Area Network (SAN)

- High performance, dedicated storage network tailored to an environment, and combining Multipurpose servers, storage products, networking products, tape products, software & services.
- Composite (Fabric) of several entities
- Longer design & implementation time
- Requires higher levels of IT skills
- Designed for Intensive Block I/O
- Fibre Channel - SCSI protocols
- Typically more expensive (\$/MB)
- High Performance & Massive Scalability



File I/O vs.. Block I/O

■ File I/O

- Sometimes referred to as **File-System protocols (NFS, CIFS..)** >>>
- Identifies data by file name and byte offset
- Transfers file data or file metadata (files's owner, permissions, creation data, etc.) via TCP/IP protocols
- Handles security, user authentication, file locking

■ Block I/O

- Sometimes referred to as raw disk or raw data
- Identifies data by disk block number
- Transfers raw disk blocks via encapsulated **SCSI protocols** >>>

File-System vs.. SCSI Protocols

■ File System

- **File-System protocols (also called "file I/O") are used for accessing and sharing data. The protocols are device independent. A file-system command might just request reading the first 80 characters from file "ABC", without knowing the location of the data on the device.**
 - Network File System (NFS)
 - ▶ A file-system protocol
 - ▶ NFS has its origins in the UNIX world.
 - Common Internet File System (CIFS)
 - ▶ Most often pronounced "siffs"
 - ▶ A file-system protocol
 - ▶ CIFS has its origins in the Microsoft NT world

■ Small Computer Systems Interface (SCSI)

- **SCSI is often called a "Block level" protocol or "Block I/O" because SCSI commands specify particular block (sector) locations on a specific disk.**
- **SCSI I/O commands tell disk devices to store and retrieve data from a specific location on a disk drive**
- **Protocol most prevalent in Storage Area Networks**

File Applications vs.. Block Applications



- ▶ **NAS offers "File" application optimization**
 - Designed to work with File System applications
- ▶ **SAN's offers "Block" application optimization**
 - These applications are typically high performing and read and write large blocks of data directly to disk

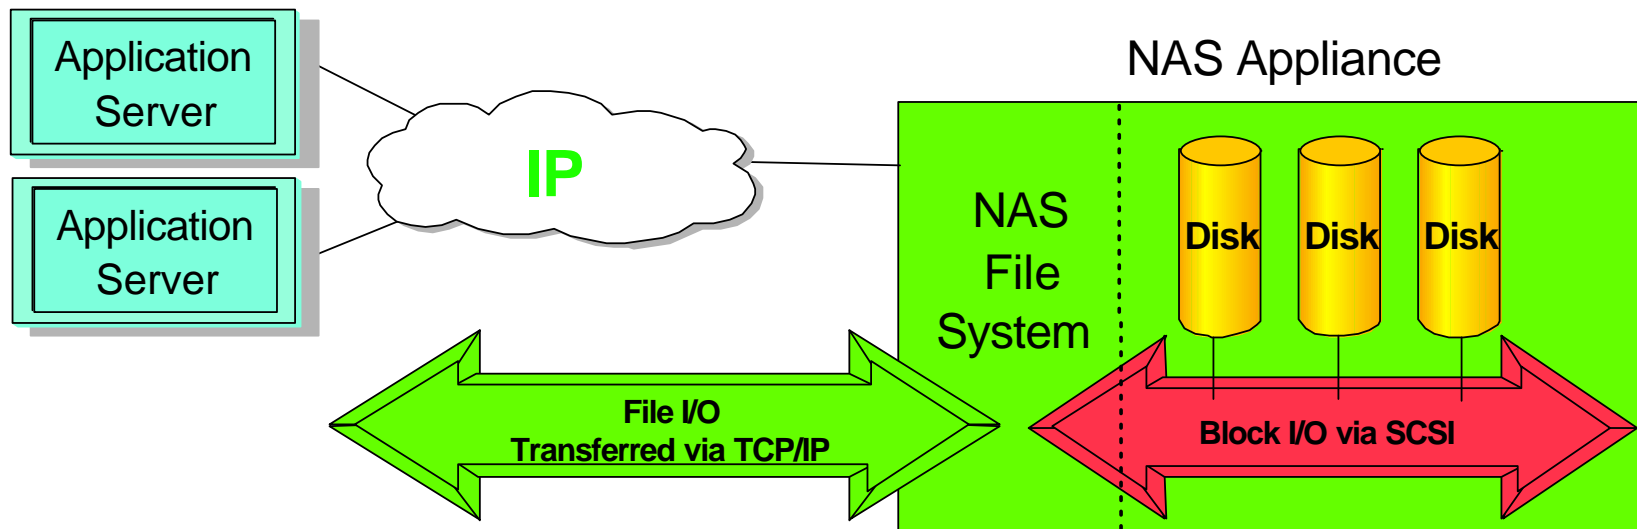

Typical FILE I/O Applications

- Freelance
- LOTUS 123
- Power Point
- Word Pro
- LOTUS NOTES
- Lotus Domino - Server
- Lotus Approach
- Video

Typical BLOCK I/O Applications

- Microsoft exchange
- DB2
- Informix
- Oracle
- Video Streaming
- ERP Applications
 - Supply Chain
 - JD Edwards
 - Lawson Software

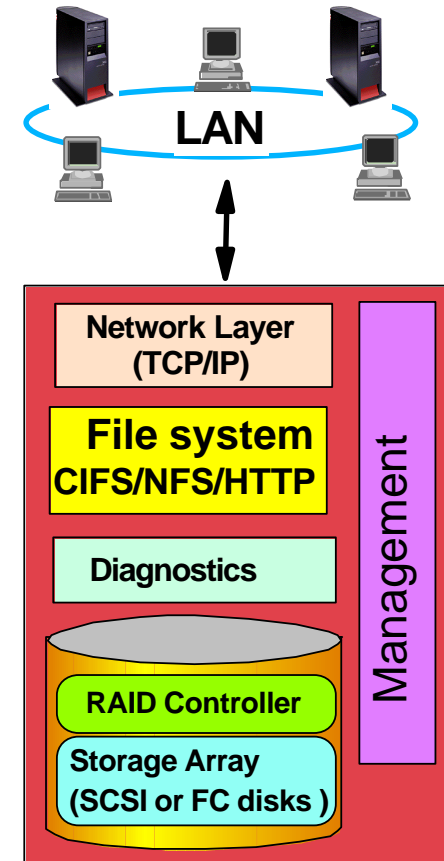
Everything is Stored on Disk in Block I/O via SCSI



Network Attached Storage Concepts

NAS Appliance Characteristics

- **Preloaded file system (O/S) that provides file sharing**
 - Windows (CIFS), Unix (NFS), Web (HTTP), Novell....
- **Storage Disk Arrays**
 - Scales from GBs to TBs
- **Installation software**
- **Management software**
 - Quota & Security
- **Data Protection Technology**
 - RAID Implementation
 - Data Protection on Disk
 - Data Backup to Tape
- **Diagnostic software**
- **Fault Tolerant Features**
 - Dual, Redundant, Hot Swap Components



Self Contained, Preinstalled, Preconfigured, Pretested at the Factory

NAS... Main Operational Characteristics

- NAS is an optimized "File Server"
 - Not a multipurpose or general-purpose server
 - Does not replace multi/general purpose servers
 - Cannot install multi/general purpose applications
 - ▶ Freelance, LOTUS 123, Power Point, Lotus Domino, ERP, DB2, Oracle, Microsoft Exchange.....
- NAS typically serves files faster than multi/general purpose servers
 - It has a finely tuned operating system (O/S) for file serving
 - It has been freed from the latency and cycles of running multiple applications for other reasons than "file serving"
- NAS directly connects to an IP Network
 - All data enters and exits a NAS appliance via an IP Network in "File System Formats (File I/O)"
 - NAS O/S internally converts the "File I/O" into "Block I/O Formats" and stores all data onto the integrated disk arrays via SCSI commands
 - ALL data, whether traversing NAS or Storage Area Networks (SAN) or Direct Attached Storage (DAS), is stored onto disks in "Block I/O Formats" via SCSI commands!
- NAS provides access to pooled storage with unlimited distance consideration
 - SAN pooled storage is limited to fibre channel distance restrictions (100 km)

NAS... Main Operational Characteristics

- NAS can talk several different "File System" languages
 - Most all NAS vendors supports:
 - **CIFS** - Common Internet File System
 - ▶ Developed by Microsoft for Windows Applications
 - **NFS** - Network File System
 - ▶ Developed by Sun for UNIX Applications
 - Other protocols supported in NAS environments
 - Netware (Novell), HTTP, FTP, MAC (Apple) ...

- NAS allows multiple clients and servers in a mixed "File System" environment to store and retrieve data on the same storage device

- NAS allows multiple clients and servers in a mixed "File System" to share each others data
 - This is called "heterogeneous file sharing"

- NAS Appliances scale anywhere from GB's to TB's of storage capacity
 - Current NAS designs have built-in performance & scalability limits whereas SANs generally do not

Network Attached Storage Solution Scenarios



Envision Media

Envision Media, a 14 person advertising agency in Soquel CA, install the NAS devices last June to help solve several file access and storage issues.

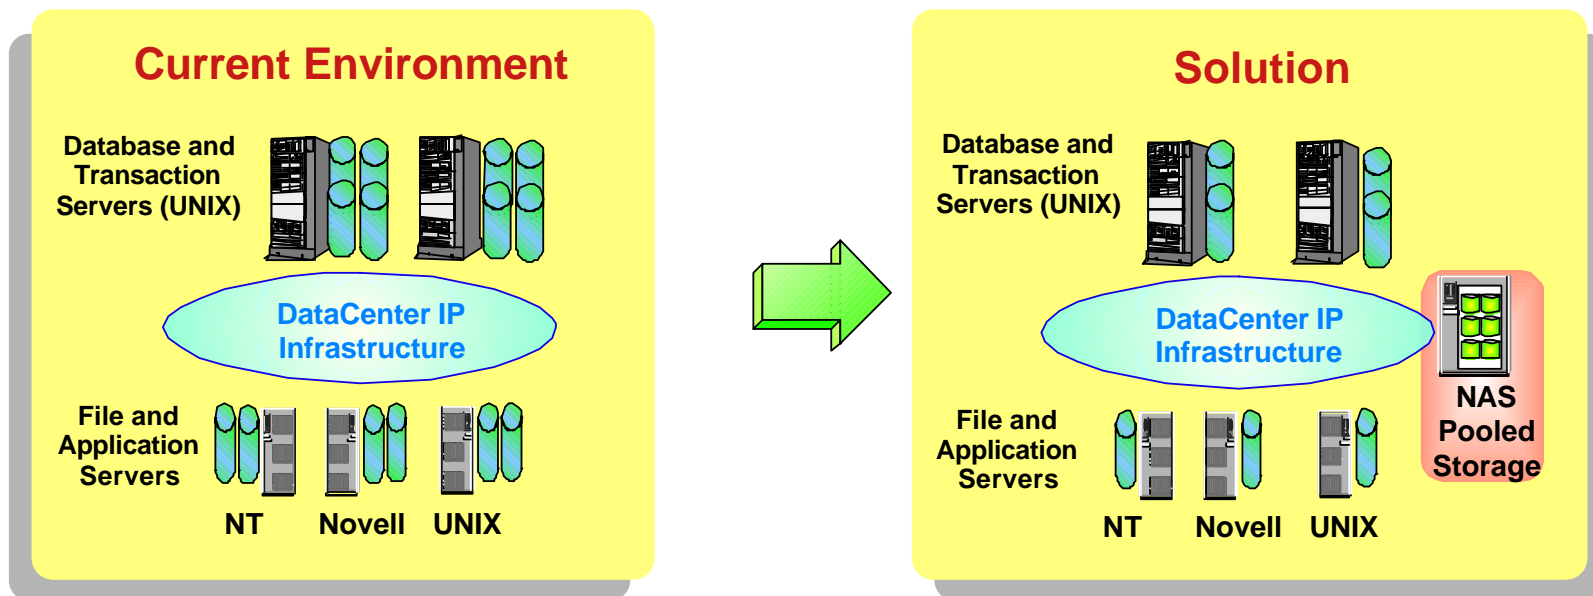
As part of its client services, Envision creates Web sites and portals. In addition, the agency runs a mixed shop of UNIX and Macintosh workstations for engineering, with business operations running applications stored on NT servers. The agency needed to share mixed files and currently with other clients.

Envision's chief operating officer, recalls the agency was increasingly concerned about failures within their general-purpose servers and the shortage of available disk space.

With no formal IT department, they looked for a simple solution capable of communicating with a variety of desktop platforms. They also needed a device that could serve as a temporary backup device or as an intermediate repository to store or access data, should a server go down.

They chose a NAS solution.

NAS Solution ... Pooled Server Storage & File Sharing



Benefits:

- Higher availability and Backup (Disk Backup, Tape backup and RAID support)
- Pooled / Centralized Storage
- Easier and quicker to administer vs.. other NT servers
- Centralized storage management
- File sharing and File synchronization



University of Georgia

The Information Technology Outreach Services (ITOS) arm of the University of Georgia in Athens, assists a variety of state and local agencies in creating, monitoring, and storing information for a wide range of activities.

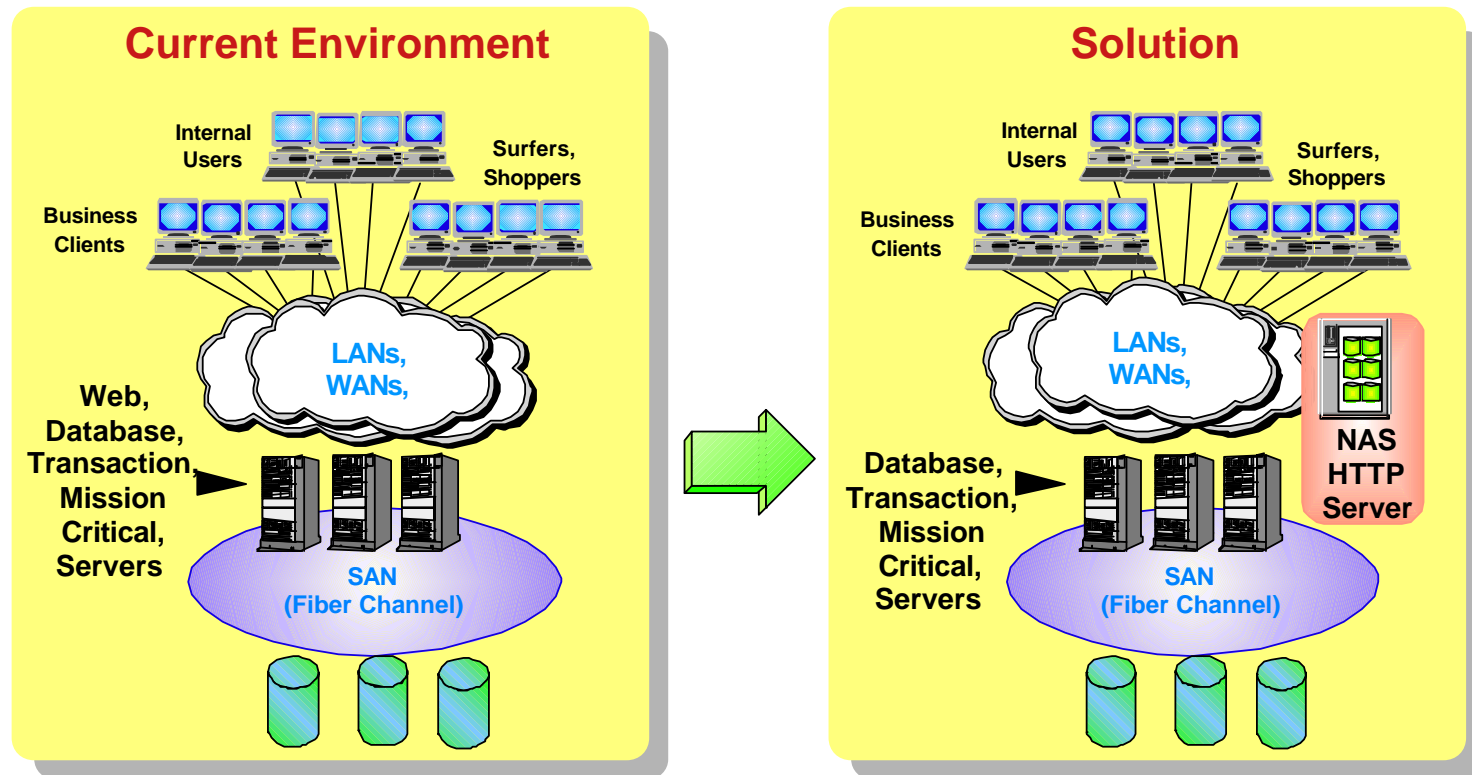
To help gage the effectiveness of the states education system, for example, ITOS maintains a 35 GB database that tracks the educational progress of every Georgia student from kindergarten through the workforce. ITOS also helps create in-store hurricane tracking computer models, and produces and maintains a repository for bandwidth hogging geographic information systems (GIS).

To serve about 100 users working on 60 workstations, ITOS uses about 10 UNIX and Windows NT servers, the servers are connected to a 12 port 3Com gigabit switch which is in turn connected to the NAS server.

Many ITOS employees using older UNIX workstations running software RAID experienced up to 35 second-screen refreshes when working on the GIS maps. Offloading that data onto a NAS server and away from the rest of the ITOS is server traffic caused refreshed time to drop to 5 seconds.

Since it can communicate to workstations and other servers via NFS and CIFS protocols, the NAS file server bridged older and newer technologies.

NAS Solution ...HTTP File Server



Benefits:

- Offloaded SAN cycles
- Increased performance throughput (service level) to end users
- Minimized investment in additional Servers
- Provided storage pooling
- Provided heterogeneous file serving
- Used existing infrastructure / tools / processes
- Isolated Web clients away from "Mission Critical" data.... SECURITY

Nellis Air Force Base

As a senior systems engineer at Nellis Air Force Base in North Las Vegas, the IT manager of a 5000 persons civil engineering (CE) squadron is responsible for designing, building, and maintaining everything on the base from sewer pipes to aircraft hangers.

After building a small server farm following a migration away from a minicomputer, the IT manager and his IT staff were able to consolidate and migrate to newer technologies, because the staff was spending too much time managing and adding capacity to older servers.

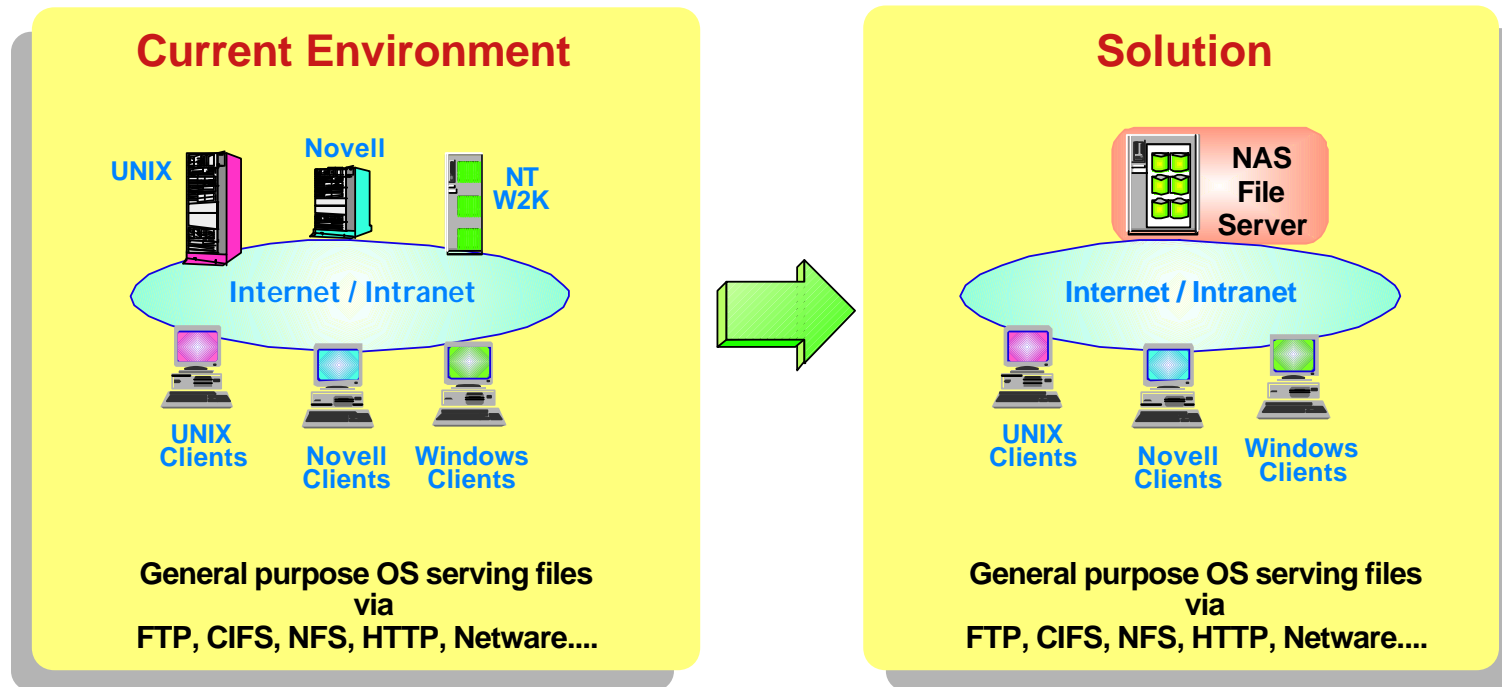
"I'm trying to simplify my job in an industry that has grown leaps and bounds beyond whatever we imagined, and the hardest thing for us to do is to keep up with the technology" (IT Manager).

This manager choose a NAS solution which quickly scaled to serve files among divergent groups of engineers and builders. CAD/CAM data stored on NAS devices provides detailed maps illustrating all electrical, water, sewer, and telecommunications lines on the base. Accessing files directly over the network was important because different user groups with different platforms required access to the same information.

For example, if an engineer uses CAD/CAM program to produce and store roadway plans on the NAS device, workers preparing to dig can access the plans without reconfiguring their client machines to access a particular server. As a result, this IT manager shaved a considerable amount of time off each project.

"A tangible NAS benefit is that it gives the IT person all the functionality of a standard server environment, but without the overhead of administration". According to this IT manager, "scaling capacity was one of the most time-consuming administrative duties and now if I need more capacity, I don't buy another server, a buy more capacity and plug into my network. It only takes about 10 to 15 minutes of configuration."

NAS Solution ...File Server Consolidation



Benefits:

- Provided continuous availability to files
- Heterogeneous file sharing
- Reduced cost for additional OS dependent servers
- Added storage capacity non-disruptively
- Consolidated storage management
- Lowered Total Cost of Ownership



Isdaner & Co.

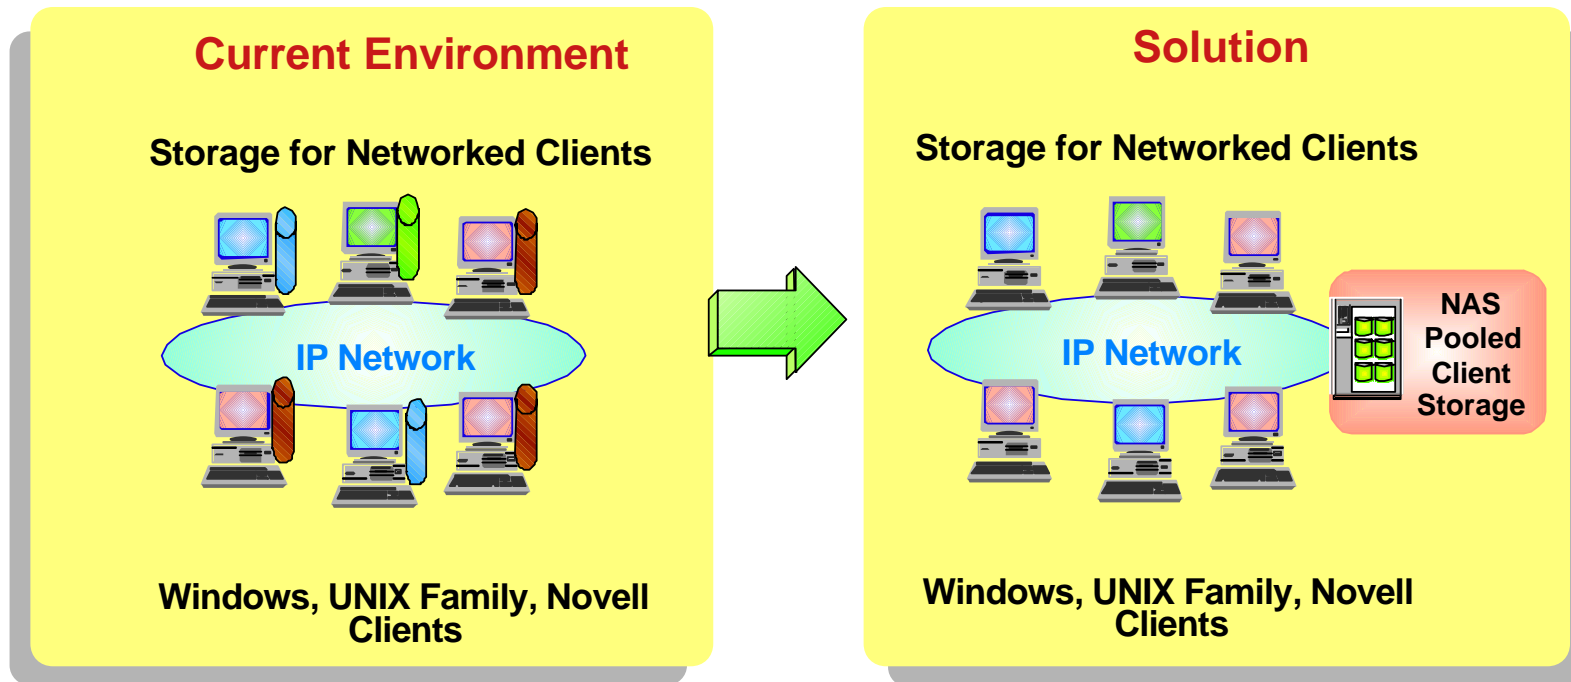
With 50 workstations and a flea of remote laptop users running NetWare and Windows NT, the accounting firm of Isdaner & CO. LLC, in Bala-Cynwyd, PA, was running out of space on its Windows NT server.

Already taxed with a slew of projects, the firm's three-person IT staff was concerned about having enough time to install and configure another server, and users, and manage users licenses. Isdaner chose a NAS solution in large part because of low administration requirements.

"Administration is a big issue because there's literally a hundred other projects that our IS staff are involved in. Layering in more administration while attending to other high's needs of a growing office was not something we wanted to encounter" (IT Manager).

An Isdaner & Co. policy mandates that laptop users take daily snapshot copies of their data while at a client site. Their NAS device proved useful as a data repository as accountants downloaded their data. Back in the office, employees now directly accessed remotely gathered information on the network without considerable time spent configuring menus and custom controls.

NAS Solution ...Storage Consolidation & Mission Critical Data Protection



Benefits:

- Pooled / centralized storage
- Easier & quicker to administer vs.. adding storage to each workstation
- Protected mission critical data with backup features (disk & tape)
- Higher availability
- File sharing / file synchronization

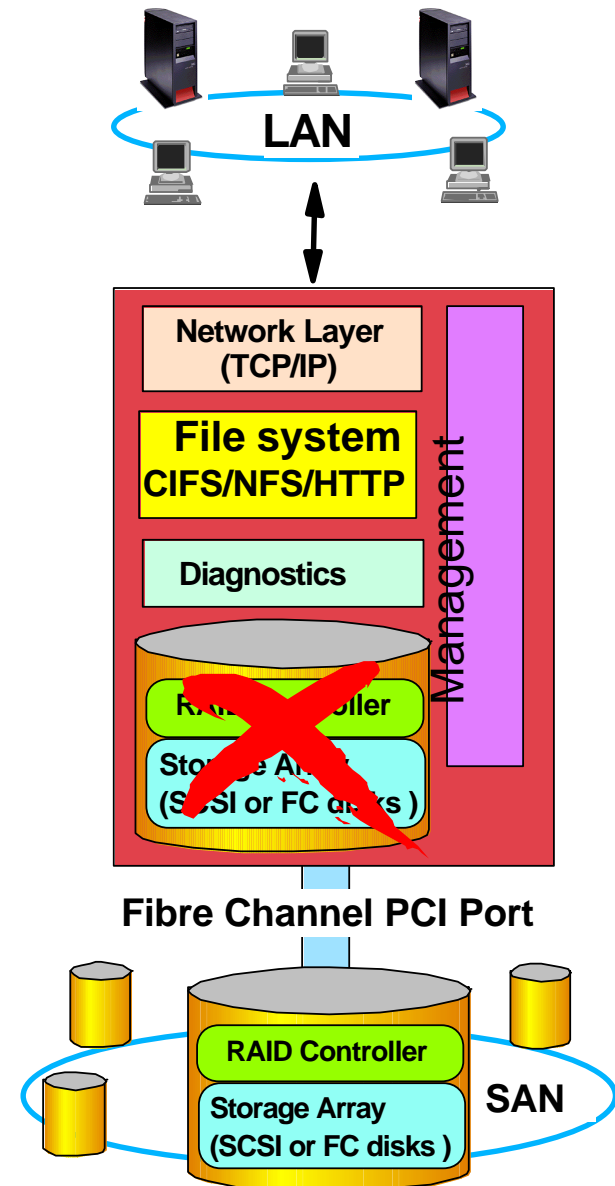
Network Attached Storage Summary

- NAS is a high powered File Server
- NAS offers pooled storage to clients and servers on an IP network
- NAS allows multiple mixed "File System" clients and servers to access, retrieve and store data on the same storage device
 - Windows, UNIX, Novell, HTTP, FTP...
- NAS allows heterogeneous file sharing to mixed clients and servers
 - i.e.: Windows client can access a UNIX file and visa versa
- NAS DOES NOT replace multi/general purpose servers
- NAS offers a low Total Cost of Ownership proposition
 - Operational, Administration, IT Skills, Storage (cost per GBs to TBs)

NAS Gateway Concepts

NAS Gateway Architectural Characteristics

- SAN vendors are starting to introduce NAS Gateways
- NAS Gateways offer the same benefits & characteristics of NAS
 - Connects to IP networks
 - Performs as a File Server
 - Heterogeneous file sharing
 - Data Protection
 - Diagnostic Capabilities
 - Clustering & failover features
- NAS Gateway is a NAS Appliance with one exception
 - NAS Gateways supports direct attachment to FC storage or connection to a storage devices across a SAN - NAS Gateways do not have any integrated disks for data storage
- Some NAS Gateway designs offer multiple connectivity options
 - FC switches, hubs, directors, RAID controllers and disk arrays



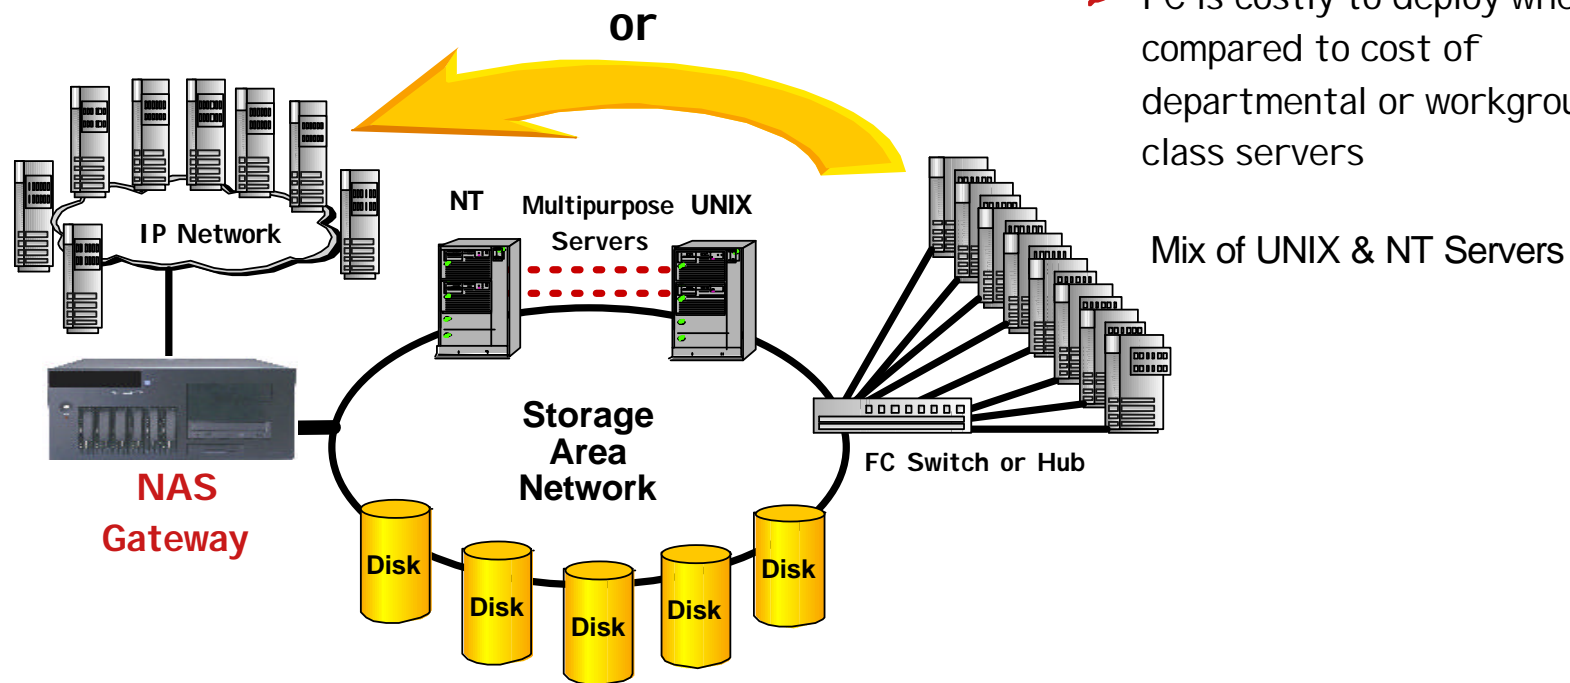
- Questions to ask when connecting remote servers to a SAN:
 - What are the distance considerations?
 - FC connectivity has distance limitations (100 km or less)
 - IP connectivity offers unlimited distance support
 - Is the cost of a FC PCI Adapter (per Server) and FC Switch/Hub cost effective?
 - Cost of a FC PCI Adapter = \$2000 (approx.)
 - Cost of a FC port on switch/hub/director = \$3000 (approx.)
 - Redundancy paths double the connection costs
 - Many times, FC connection costs exceed the value of the Server
 - Ethernet connection costs are significantly less
 - Do you need Heterogeneous file sharing/serving?
 - Does the environment have a mix of UNIX, NT, Novell Clients and Servers?
 - NAS devices support heterogeneous file sharing

NAS Gateway Solution Scenario



- 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



- Gives you the combined benefits of NAS and SAN
 - SAN scalability and performance
 - NAS flexibility and ease of use

- Increases the reach of SAN infrastructure
 - Extends beyond topology limitations of fibre channel
 - Provides IP device access to SAN storage

- Leverages the value of SAN investment
 - Reduces access costs
 - Allows access to underutilized SAN storage
 - Enables heterogeneous file serving

iSCSI Concepts

▶ **iSCSI (SCSI over IP)**

- **New and emerging technology**
- **Provides routing of SCSI protocol over IP networks**
- **Standards Based**
 - IETF - Internet Engineering Task Force
www.ietf.org/internet-drafts/drafts-ietf-ips-iscsi-06.txt
- **Prior to iSCSI there was no standardized method to transport "Block I/O Formats" over IP networks**
- **Clients and Servers wanting to store data on LAN connected storage devices (NAS) were required to transport the data in File I/O formats (CIFS, NFS...)**
 - Added additional layers of latency

▶ **iSCSI can enhance and compliment IP Storage solutions**

- **Remember, all data, whether traversing NAS or SAN or DAS, is stored onto disk drives in "Block I/O Formats" via SCSI commands**
- **Only standardized method for storing data over IP networks is in "File System" formats**
 - Prior to iSCSI standards
- **iSCSI reduces a layer of latency**
 - Eliminates the need for "block I/O" conversion to "File I/O" before transmission over IP network

What's the Immediate Benefit of iSCSI?

- **Data Base Management Systems (DBMS) applications are designed to run in a block-oriented environment**
- **These high performance applications do not need universal access, from many environments, at the file level (like NAS). They care more about block level performance and control**
- **DBMS applications demand direct connection to storage via Block I/O SCSI formats.**
- **Consider that DBMSs are constrained with File I/O based access to storage (like NAS)**
 - Example: If you are using SQL with NAS, the "Block I/O" needs to be converted to "File I/O", transported over IP network, deconverted from "File I/O" into Block I/O and stored on disk
 - This process creates enormous latency.
- **In order for high performance applications to interact with remote storage (connected via a LAN), the need to reduce latency is critical**
- **iSCSI a step in the right direction.**

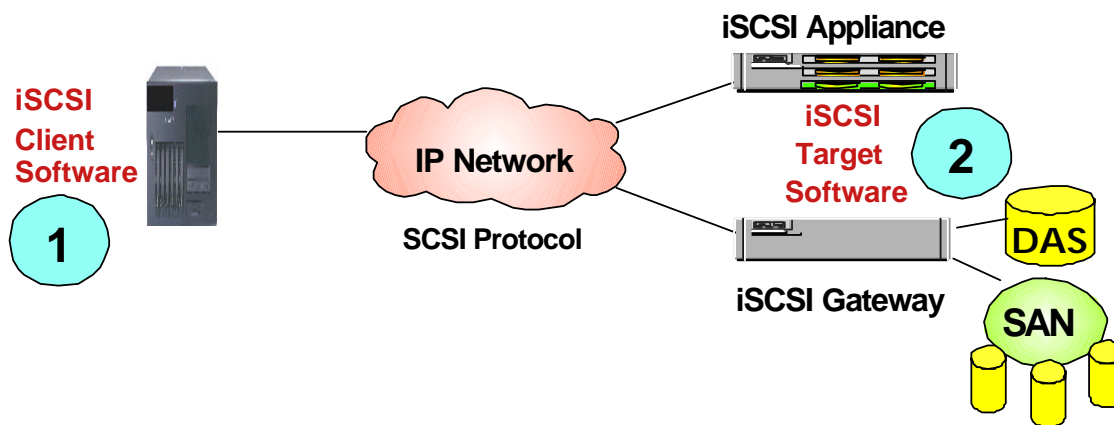
iSCSI solutions are currently available in two deliverables

1. iSCSI Appliance

- Totally integrated with Storage and Server combined into one entity
- Looks like a NAS device but does "block I/O" on IP network via iSCSI protocols vs.. "file I/O" on IP network via File System protocols

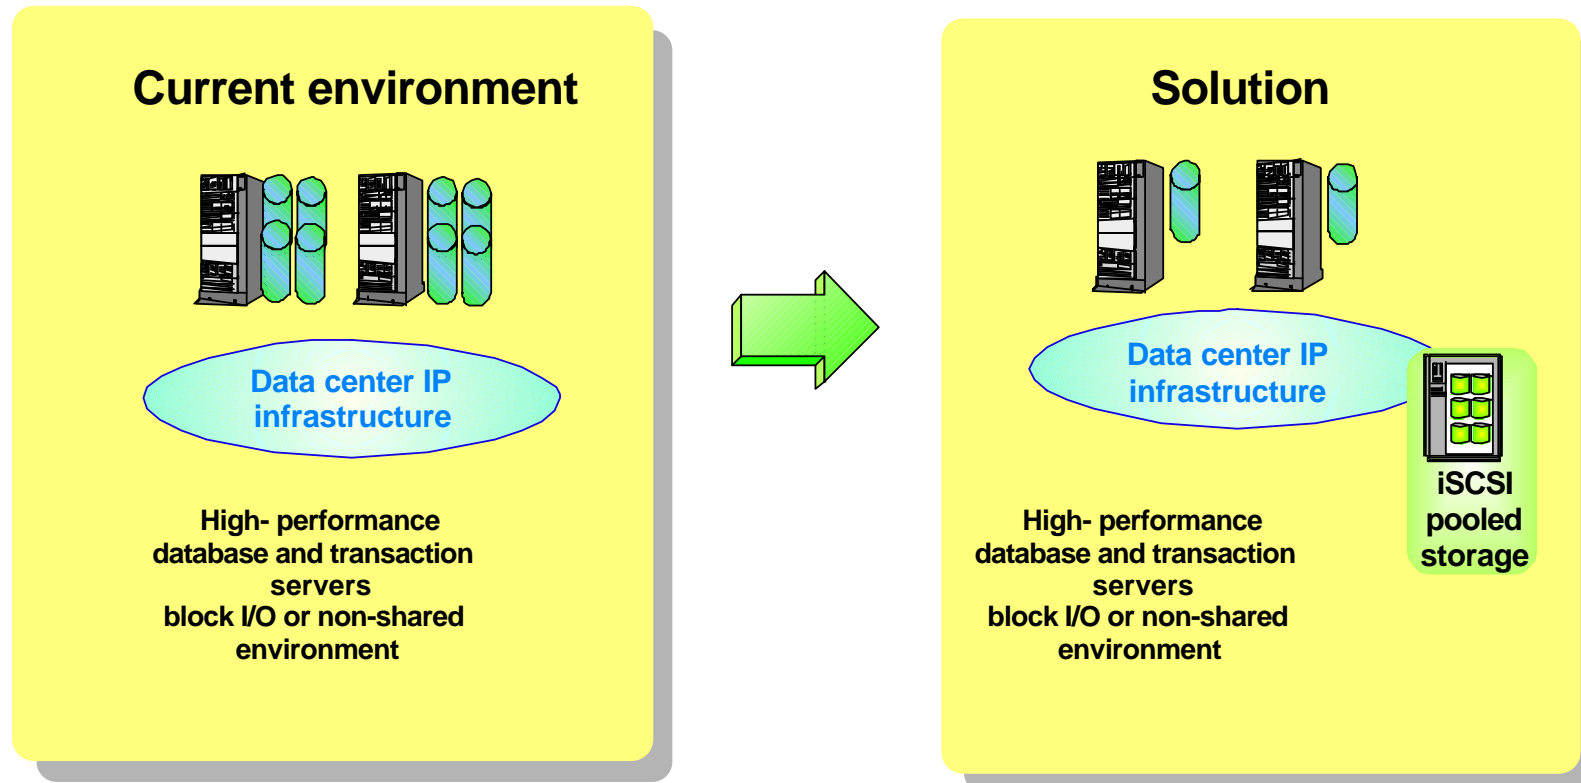
2. iSCSI Gateway

- Access to Storage Area Network or RAID Storage
- Looks like a NAS Gateway device but does "block I/O" on IP network via iSCSI protocols vs.. "file I/O" on IP network via File System protocols



1. iClient (initiator) code routes SCSI commands over IP network
2. iSCSI target code receives SCSI commands from IP network

iSCSI Pooled Storage Consolidation



Benefits:



- Pooled/centralized storage
- Non-disruptive growth
- Centralized storage management
- Utilized existing network/IP skills

Ethernet vs.. Fibre Channel Alternative

Current Option

Multipurpose Servers


Database applications
block I/O



Alternative Option

Multipurpose Servers

Database applications
block I/O



Benefits:

- Provides pooled storage to high performance database, applications
- Uses existing infrastructure/tools/processes
- Requires fewer IT skills
- Requires minimum investment

Positioning iSCSI and NAS



If Primary Benefactor = End Users And If Primary Need is Sharing Files		NAS
If Primary Benefactor = End Users And If No Need to Share Files		iSCSI
If Primary Benefactor = Server App.		iSCSI
If Primary Benefactor = Server App. And If Need to Share Files		NAS
If FC isn't realistically an option		Combination NAS & iSCSI Appliance

Credits, Acknowledgments & Disclaimers



© Copyright International Business Machines Corporation 2001
IBM Corporation
PO Box 12195
Research Triangle Park NC 27709
U.S.A. 06-01
All Rights Reserved

References in this presentation to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

All information being released represents IBM corporation's current intent, goals and/or objectives, and is subject to change or withdrawal without notice.

The performance data contained herein was obtained in a controlled environment based on the use of specific data. Actual results that may be obtained in other operating environments may vary significantly. These values do not constitute a guarantee of performance.

IBM, the IBM logo, Netfinity, Predictive Failure Analysis, and ServeRAID are trademarks of International Business Machines Corporation in the United States, other countries or both.

Tivoli is a trademark of Tivoli Systems, Inc. in the United States or other countries or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, or other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product and service names may be trademarks or service marks of other companies.