Network Attached Storage

Randy Kerns

Evaluator Group, Inc.
7720 E. Belleview Avenue • Suite 210 • Greenwood Village, CO 80111
(303) 221-7867 • Fax: (303) 221-1615
www.evaluatorgroup.com

Topics to be discussed

- Fundamentals of file storage
  - What is NAS
  - Characteristics of NAS
  - Performance
  - Architectures
- NAS Management
- Status and Futures
- Summary

Fundamentals of file storage

- Storage media can be arranged for access in many ways:
  - Direct attached storage (DAS)
  - Storage area networks (SAN)
  - Network-attached storage (NAS)

Fundamentals of file storage

- Evolution of NAS
  - Demand for remote file-based storage
  - General-purpose servers led to specialized file servers
  - Now known as NAS - offers reliability, performance and ease of installation
    - Also known as “filers”

What is NAS?

- Special purpose device to provide remote file system to other servers on network
- Server that is dedicated to nothing more than file sharing
- SNIA defines NAS as "storage devices that connects to a network and provide file access services to computer systems”
- Attached to local area network (typically and Ethernet network) and assigned an IP address
- File requests mapped by main server to NAS file server

What is NAS?

- Data is accessed from a NAS over a network via a file “redirector”
- Utilizes IP for connection protocol and UDP or TCP
- Ethernet/Gigabit Ethernet for media access
- Usually a kernel or thin server that provides file-level access using standardized protocols
  - Network File System (NFS) for UNIX environments
  - Common Internet File System (CIFS) for Windows environments
  - Others

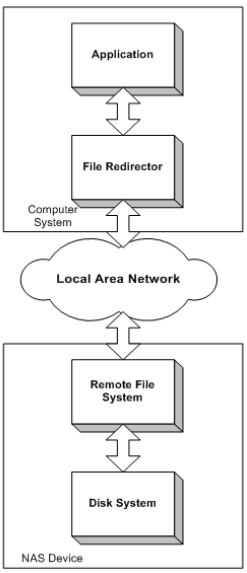

Hardware Elements in NAS

Component	Function
Network Interface	Ethernet, etc. network attachment
Processor	Customized operating system to process remote file system access
Device Adapter	Attachment for storage device
Storage	Storage device of some type (usually disk that is attached via Fibre Channel, SCSI, or ATA)
Service Interface	Ethernet and/or RS232 interface that is used to perform maintenance and administration. May be a separate processor and adapter for the service.

Software Elements in NAS

Component	Function
Network Layer	TCPIP protocol stack for network attachment
Remote File System	Implementation of a file system with remote access protocols. Today's implementations would typically be NFS and CIFS with others such as HTTP, NTP and AppleTalk available
Multi-protocol Handler	Special software to be able to simultaneously handle more than one Remote File System type of access (both NFS and CIFS for example)
Native File System	Implementation of the file system used in the NAS device
Driver for storage device	Device specific software to handle the type of device attached to the NAS controller
Administration	Software to perform configuration and management tasks

Protocols - NFS

- Network File System
  - Originally developed by Sun and released in 1985
  - Provides remote file system with access semantics across a network (local or wide)
  - File system is mounted for accessing by client
  - Open standard for file sharing
  - Security is controlled by authentication and ACLs (access control lists)
    - Administrator managed names of users and groups
  - Lock manager used to control access to files
    - Advisory level locking used in Unix
    - File locking (and cooperating software) required for cache consistency

Protocols - CIFS

- Common Internet File System
  - Developed by Microsoft as a remote file system access protocol for use over the Internet
    - Derived from Microsoft SMB in 1992
  - File sharing support with file locking mechanisms
  - Remote file systems don't have to be mounted but can be referred to by globally significant names
  - Supports Unicode file naming
  - Implemented on some Unix, Linux, and VMS
    - Usually via SAMBA
  - Locking is implemented as "Opportunistic Locks" to allow local buffering and exclusive access (hard locks)
  - Security is done with authentication and both share and user level access to file system, Active Directory

Protocols - Other

- NCP (Netware Communication Protocol): Netware file serving protocol that competes with NFS and CIFS
- AppleTalk: the Apple file serving protocol that competes with NFS and CIFS
- FTP (File Transfer Protocol): protocol between computers where file data is transferred with appropriate handshaking
- HTTP (Hyper Text Transport Protocol): method of accessing data over the web typically for display of information in web pages

NAS I/O Structure

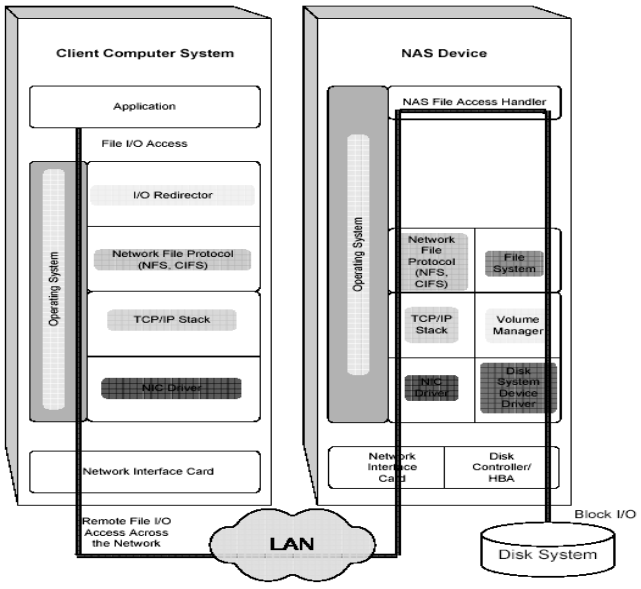

Block I/O vs. File I/O

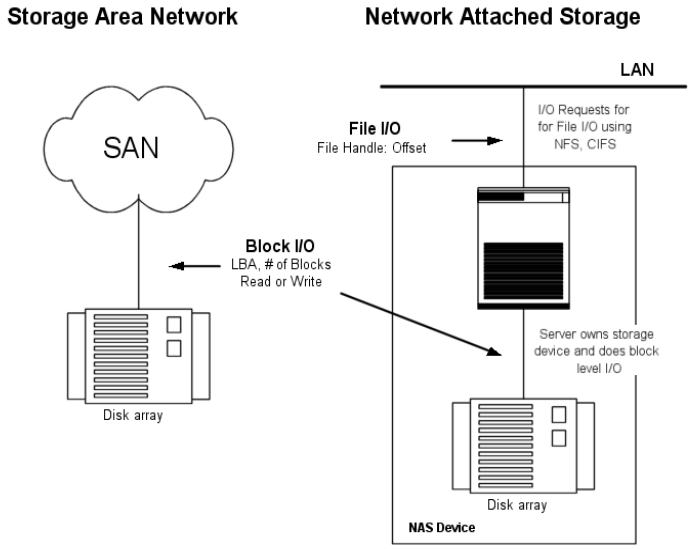

Characteristics of NAS

- Simplicity
  - Plug and play operation
- High-Availability
  - Built-in fault tolerant capabilities or clustering
  - Some utilize fault tolerant RAID storage systems
  - Fault tolerance in the NAS control function

Characteristics of NAS

- Scalability
  - Scalable in capacity and performance
  - Not linear improvements due to
    - LAN dependency
    - Application server handling the TCPIP connection
    - Other traffic on the connectionless topology
- Connectivity
  - Allows for multiple network connections for support of multiple networks and users
- Centralized storage management
  - Only for the specific NAS device

Characteristics of NAS

- Data Sharing
  - Via its implementation of a remote file system
  - Users on different client systems can have access to the same file on the NAS with the access serialization
  - Some NAS systems provide a translation mechanism to allow heterogeneous sharing for both NFS (UNIX) and CIFS (NT/2000) implementations

Characteristics of NAS

- Integrated Backup
  - Integrated or third-party backup capabilities
- Infrastructure
  - Leverages existing IP infrastructure and network administration skills
  - Drawbacks of TCP/IP
    - Packet sizes are typically small
    - Network congestion may result in performance variability
    - Dropped packets
    - Network transport layer has high overhead
- TCO - Lower than SAN or DAS

Issues with TCP/IP

- Processing overhead can be 50 – 100% of CPU cycles
  - congestion control algorithms
  - timer management
  - connection management
  - checksum calculation
  - transmitter data buffer handling
  - segmentation and reassembly
  - sliding window flow control mechanism
- Protocol overhead time usually exceeds data transfer time
  - If increasing network speed, then processors, operating system overhead, and NICs must also scale accordingly
- Large base implementation makes significant changes very difficult

TCP/IP Accelerators

- Improvement is to replace standard NIC with TCP/IP Accelerator (also called TCP/IP Offload Engine or TOE)
  - Does much of the TCP and IP function in new NIC card (host bus adapter)
    - Segmentation and reassembly
    - Buffering between processing layers
    - Checksum processing
    - Checking for dropped packets
  - Relieves processor and improves performance
- Requires change-out of NIC cards
- Requires new drivers
- Not commonplace yet – several vendors with slightly different implementations

Remote Direct Memory Access

- RDMA
  - Bypass internal data moves (Zero Copy) and operating system involvement for data movement over a network
  - Reduces CPU processing required
  - Improve performance by reducing overhead
    - Line rate can result in only 1/3 **real data** transfer rate
- Implementations
  - Virtual Interface Architecture
  - InfiniBand – VIA with IB as the transport
  - DAFS – file transfer using VIA
  - RDMA over IP – proposed standard over IP
    - Requires new NIC – RNIC
    - Initial proposal is for NFS only

RDMA

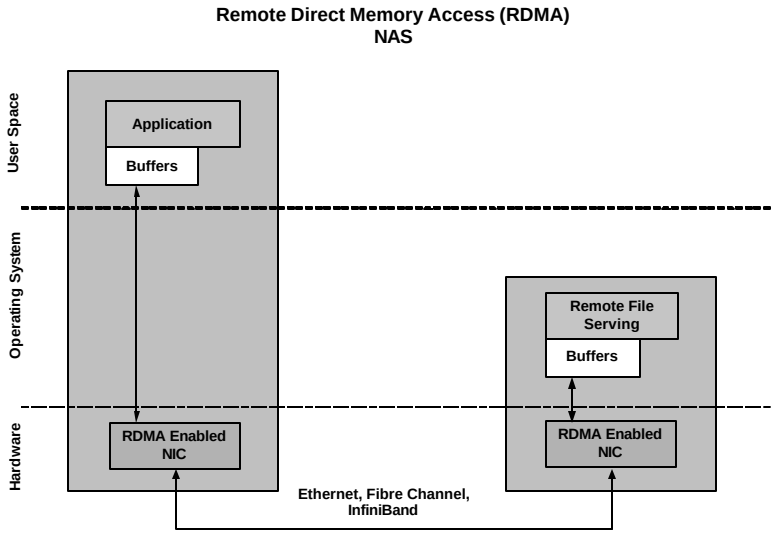

Architectures/Implementations

- Traditional NAS
  - +TCP offload engines (TOE) - move the processing of the TCP/IP stack out of the OS and onto a NIC
- Gateway
  - Filer that has been optimized to serve files connected to SAN storage on the back end
  - Fusion of NAS/SAN technologies
- NAS aggregation
  - Specialized NAS operating system that allows the file system information to be distributed across many NAS nodes
- Two-stage backup
  - Disk-based storage between application storage and tape libraries
- NAS on a blade

NAS Gateway

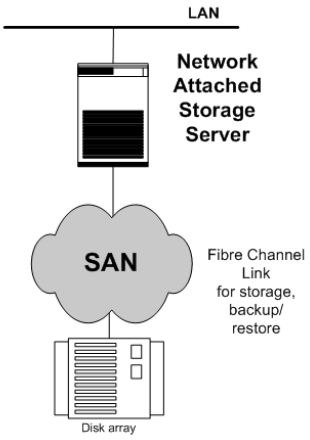

- NAS "head" with SAN storage
- The controller function of NAS with the storage decoupled
- Connected to a SAN – fibre channel typically
- May be separate product or a version of a standard NAS product
- Can utilize SAN capabilities
- Does introduce a different management scheme
  - Adds the SAN storage management as well as the NAS administration

NAS Gateway

- Typical usage profile
  - Enterprise data center environments with existing SAN
    - Requirement for file storage
    - Consolidation of existing, independent NAS devices
    - Centralized administration and purchasing
    - Storage system partitioned for NAS filesystems (LUNs)
  - Administration cost
    - Target is to reduce overall storage administration costs

NAS Management

- Evolution
  - NAS devices each individually managed
    - Installed and administered as unique devices
  - Management does not scale
    - Additive costs for additional NAS devices
  - Few sophisticated tools
    - Expectation for capacity planning, performance monitoring and reporting, and SLA adherence didn't exist
  - Changes are underway for NAS management

NAS Management

- Aggregation
  - Management software to manage multiple NAS devices
  - Two methods
    - Individual devices with information accumulated
      - Statistics, capacity information
      - Controls, status
    - Multiple devices managed as single entity
      - Dynamic allocation of capacity – adding of capacity (units)
      - Data protection dynamically spread among devices
      - Performance tuning by distributing data

NAS Aggregation Example

- Software to collect information about NAS appliances
  - Same vendor
- Unified view to monitor and control
  - Status and events
  - Capacity
  - Appliances, OS versions
  - Volumes
  - Disk devices
- Addresses administration costs in homogeneous environment

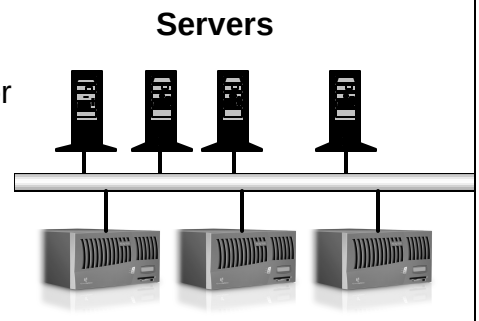

4 May 2004

27

NAS Aggregation Example

- Data accessed as if single NAS appliance
- For increased capacity, additional appliances added
- Data is redistributed across appliances – non-disruptively
- Clients see additional capacity with no required changes
- Data protected as RAID across appliances
- Management as a single entity

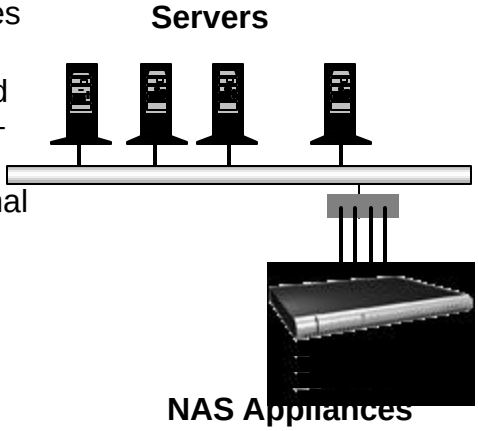

4 May 2004

28

NAS Management

- Summary
  - NAS Management is maturing
  - Economics of management / administration cost is now becoming important
    - Necessary with expanded use of NAS
    - Capacity growth accentuate the problem
  - Need to integrate with SRM tools
    - Must include NAS devices from different vendors
    - SMI-S Version 1.1 addresses NAS
    - Still a long way to go to have enterprise tools

NAS Status

- NAS functionality
  - Added-value functions as differentiators
    - Point-in-Time copy and Remote copy
    - Virus protection and Authentication controls
  - Capacities – scale to enormous amounts
  - Management
    - Aggregation in some instances
    - Tie-in with overall SRM software beginning
- NAS Performance
  - Quantum increase with TCP/IP Accelerators
- NAS Competition is iSCSI storage
  - Databases – primarily block oriented where performance or application limitations are an issue

Final Look - NAS Status

- NAS Offerings
  - > 40 companies
    - Around 20 are public
  - Analogy to RAID storage market 8 years ago
    - 200+ vendors
    - Most did not survive
    - Number of non-system vendors with > $100M in revenue now = 4 (HDS, EMC, STK, NetApp)
  - Low barrier to entry
  - Crowded market channels
    - Most sales are indirect
    - Resellers have many choices
    - Vendors with direct sales usually with specialized sales force

NAS Status

- Overall view of market – continued
  - Prices are declining
    - Tough to make livable margins
    - Some vendors moving up in food chain – higher-end NAS devices for the enterprise
    - Network Appliance has RDMA implementation
      - "Application-specific" NAS offering
      - Higher margin strategy
    - Expect other vendors to seek differentiation
      - NAS as a backup device is now being promoted
      - Target of massively-scalable NAS devices

NAS Futures

- Ease of installation and administration will continue to be main reason for deployment
- Performance increases will remove obstacle for many customers
  - TCP/IP acceleration will become common
  - OLTP performance acceptable for larger segment
  - RDMA will offer greater performance beyond narrow implementations
- Fewer providers
  - Margin pressures will be difficult
  - Many will be small players (<$100M)

Summary

- Market that was controlled by one vendor has expanded to many
- Major server vendors have all introduced products
- Management and security are beginning to become a differentiating factor
- Major battle ensuing with IP-Based solutions
- New developments
  - RDMA
    - High performance
    - Vendor offerings?
  - TCPIP accelerators
    - Improved performance