



# **NAS: New Frontiers Beyond The Filer**

Michael Davis

NAS Strategy and Planning

Data Management Group

Sun Microsystems, Inc.

# NAS

## NAS (naz, nas)

1. Optimized server for file serving/sharing, including integrated storage, attaching to IP LAN.
2. Intelligent device for single, well-defined function (appliance).

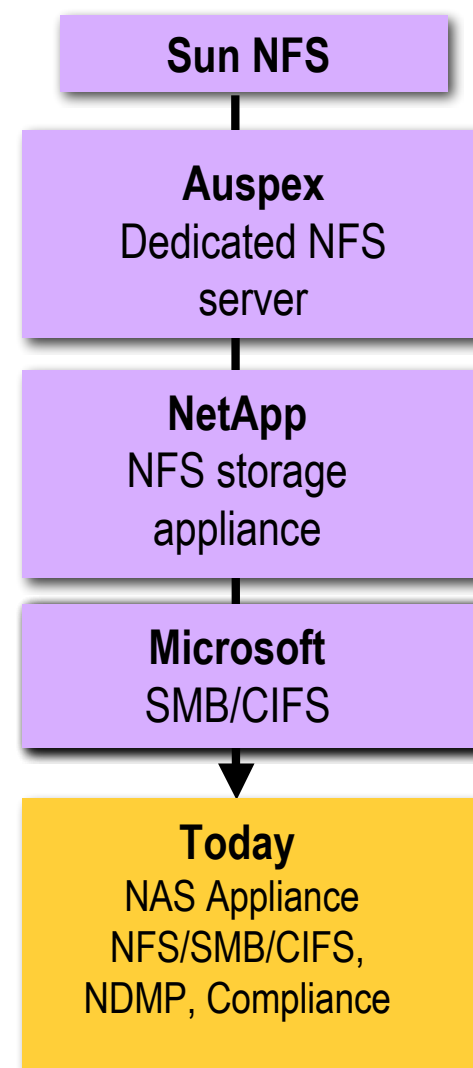
**NAS = file access**

Good for sharing data

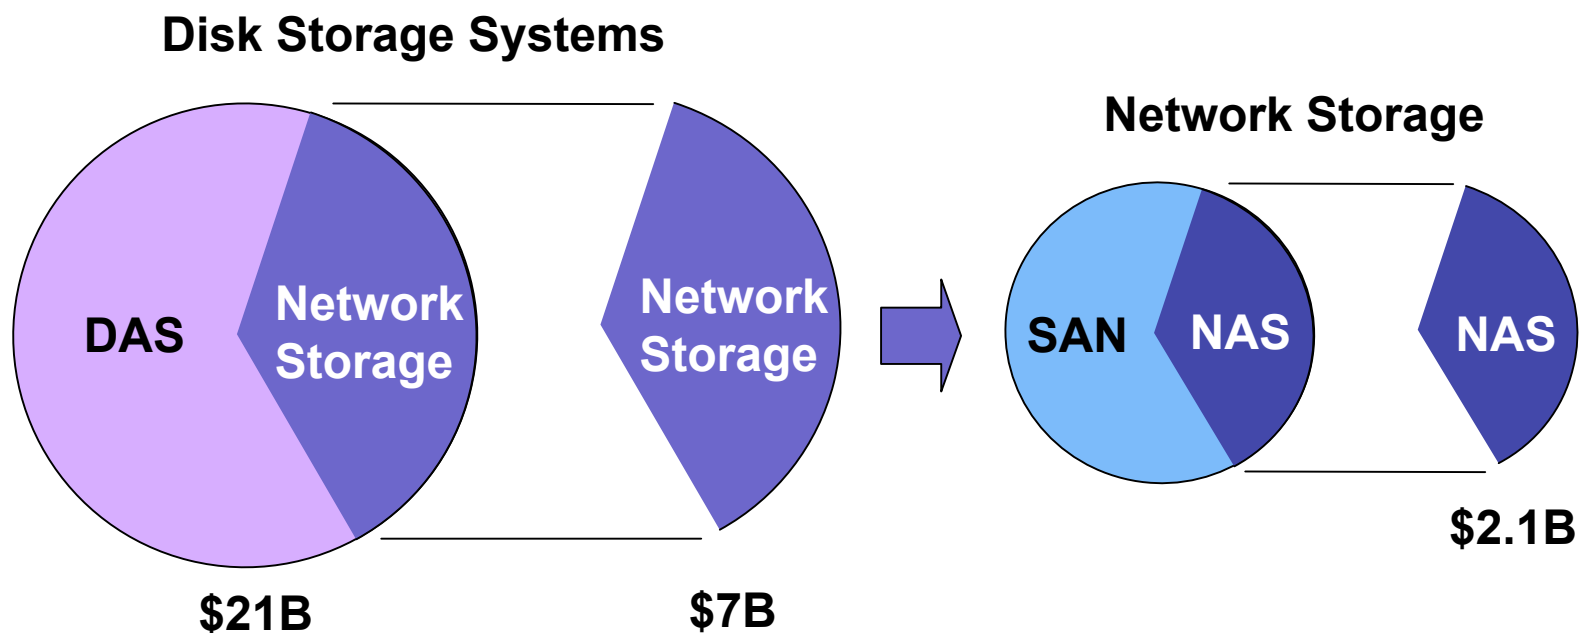
**SAN = block access**

Good for sharing resources

## NAS Evolution



# Consolidation is the Primary NAS Opportunity



# Many NAS Benefits

## Reduce management

- Simplify admin, protection
- Reduce repetitive tasks

## Reduce capacity needs

- Eliminate DAS redundancies
- Share/eliminate duplicate files

## Increase reliability

- Monitored platform and alerts
- RAID protection
- Clustering

## Improve performance

- Dedicated, optimized server

## Improve Disaster Recovery

- Backup from single source
- Eliminate redundant file backup
- Snapshot/replicate transparently

## Improve scalability

- Simplify configuration changes
- Reduce time/effort to add storage

## Improve security

- Set policies for file access
- Set policies for groups/users

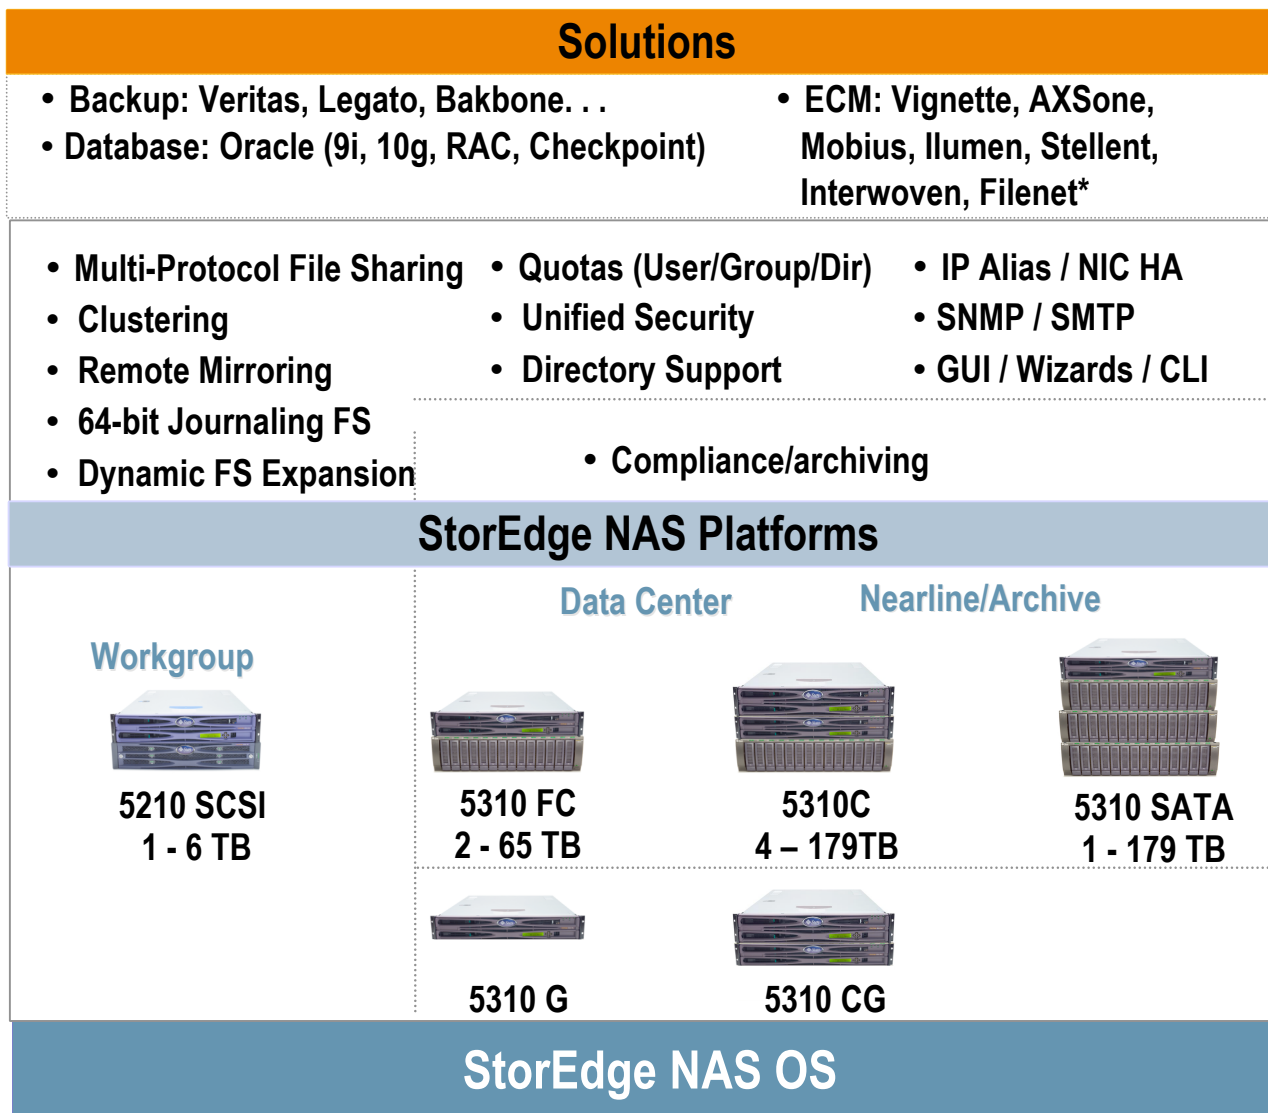
# Workgroup to Datacenter Solutions

3rd-party solutions

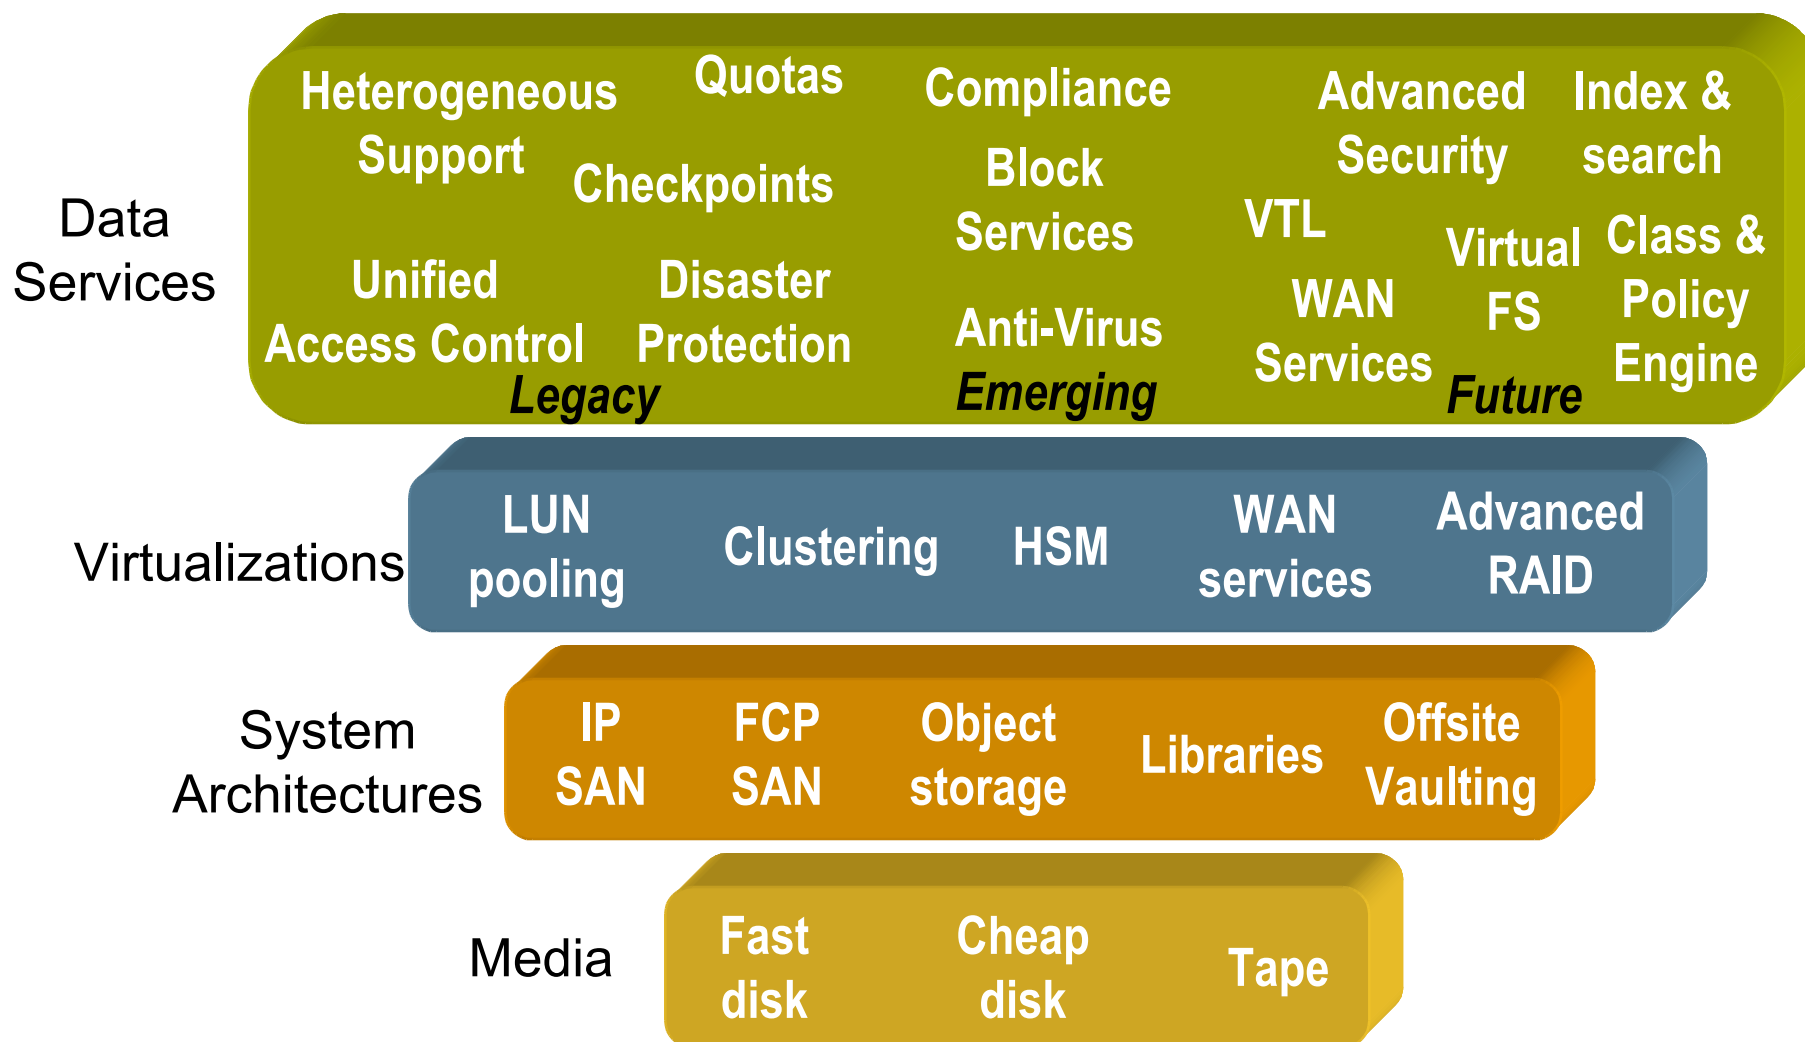
Solutions for file sharing, data protection, security and management

Industry-standard hardware platforms

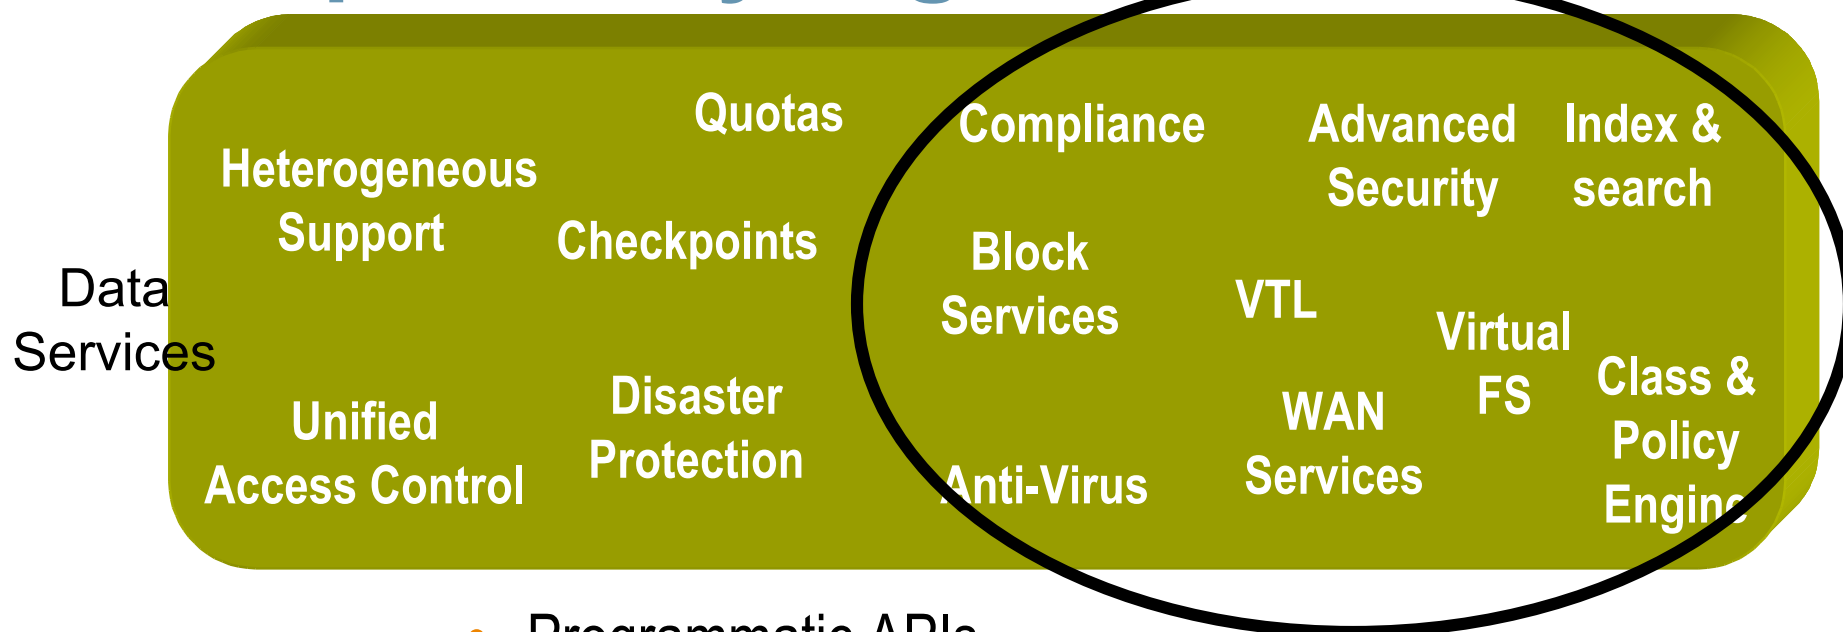
Storage appliance OS



# NAS is Not A Clear Definition



# Increased Intelligence = Worsening Interoperability Nightmare?



- Programmatic APIs
- Pseudo-APIs (eg CHMOD 444 for WORM)
- NFS Extensions
- Administrative: SMIS, CLI, GUI
- Custom agents and shims (ISV role in 'standards')

# What Problems Are We Solving?

## SAN

Fast

Low Overhead

Device Specific

Geometry Specific

Transport Specific

Limited Distance

Limited Data Sharing

Limited Hetero

“Blocks”

## NAS

Slow(ish)

High Overhead

Device Independent

N/A

Transport Independent

Any Distance

Excellent Data Sharing

Completely Hetero

“Files”



# This Is Not the End Point

**SAN (Block)  
FC Network  
FCP Protocol**

**NAS (File)  
Gig-E Network  
IP Protocol**



This is a filer and a disk controller in the same box  
Sun is doing this. This is good, but it's not the end  
of the road

# Is There Room For an Object Interface?



I want to write 26,763 bytes to  
you  
Handle = Get\_Rich\_Quick  
QoS = High  
Metadata = 10/19/05 NASConf  
Retention = 100 years  
Access Key = \*&^%#  
Data Payload...

Storage Node



“Devices”



## Big Message:


Applications (and humans) are not concerned with where the data is located, only that it is there & secure

## Storage Node Functions

- Outboard Space Management
- Security/Access Control
- QoS Implementation
- Cache, Backup, RAID
- Policy Migration, Retention

# Modernizing the Access Semantic

A New “Layer” to Talk To



# Problem: Data Outlives Infrastructure!

- Magnetic media life = ~5 years, Papyrus = ~3000 years
- Retention requirement = 1 day to “Forever”
  - > Life Insurance Policy & Claim data: 99 years after actuarial death of the insured
  - > Financial Data: 5-7 years and increasing
  - > Healthcare: 7 years, 20+ years, some “forever”
  - > Deeds and real property maps: Forever
- Solutions:
  - > Transparent technology refresh (media, CPUs, vendors?)
  - > Multi-generational support
  - > Policy-based de-commissioning

# Problems: Assured Delete

- A looming post-retention problem
- How do I delete data when:
  - > I may not know where it is?
  - > It might be on off-line tape?
  - > How many copies are there? Where are they?
  - > Forensic disk recovery can retrieve “erased” data
- Ask the right question: Delete or “make unavailable”?
  - > Encrypted data and robust key management  
[http://research.sun.com/spotlight/2005-03-01\\_Ephemerizer.html](http://research.sun.com/spotlight/2005-03-01_Ephemerizer.html)

# Problem: We are Losing Value

- In large-scale applications, file systems have lost meaning to customers
  - > Finding the needle requires a rich application schema
  - > Query is the everyday use model
  - > FS limitations: scaling, insuring uniqueness, static view
- Rich metadata = content awareness
- Content awareness adds the value back
  - > Query
  - > Virtual dynamic hierarchical file systems
  - > Smarter policy engines

# Technologies Influencing Trends

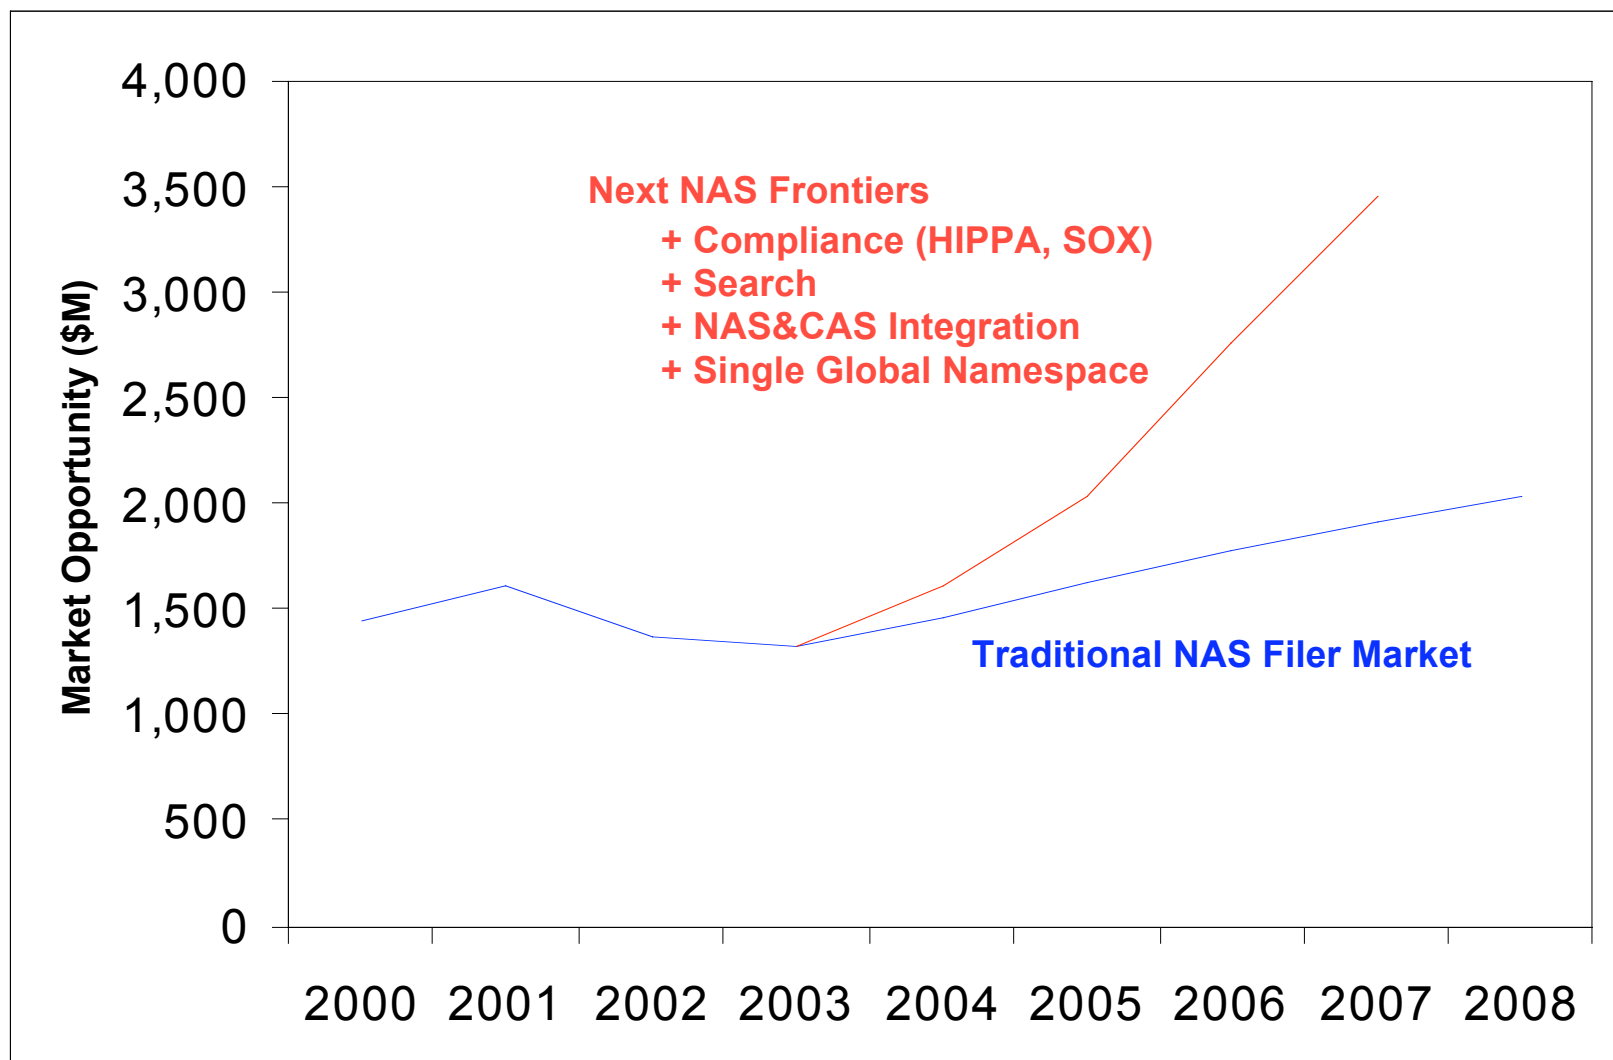
- NFS v4
  - > NT style ACLs
  - > File delegations
  - > Statefull connection
  - > WAN improvements (TCP, compound, security, locks)
  - > Migration & replication (although limited)
- PNFS extensions - horizontally scaleable performance
- 10G Ethernet transition
- iSCSI - Consolidating block and file services

# Where Can NFS Help Further?

- Open Standards support for new data services, to reduce business risk to partner community
  - > Capturing extended attributes
  - > Executing a query
  - > Invoking compliance functionality
- Heterogeneous interoperability
- Integration with identity management
- Namespace scaling
- Stronger WAN capabilities



# We Can Grow the Market



# **NAS: New Frontiers Beyond The Filer**

- **Michael Davis**
- NAS Strategy and Planning
- Data Management Group
  - > Sun Microsystems, Inc.