SAN y NAS

Acceso al almacenamiento

- El acceso al almacenamiento puede ser directo
  - *Direct Attached Storage (DAS)*
  - En ese caso cada servidor necesita su sistema de almacenamiento
  - Estos sistemas de almacenamiento externos están infrautilizados
  - Pueden ser por ejemplo para llevar a cabo *backups*
  - Un backup nocturno significa que el resto del día esos discos están inutilizados

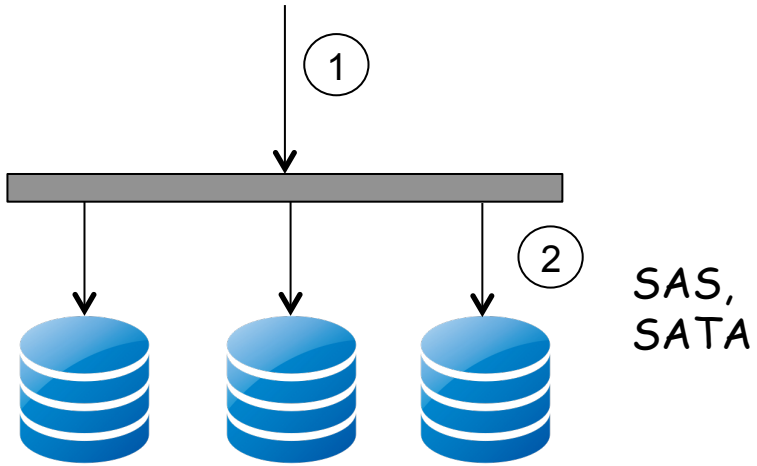

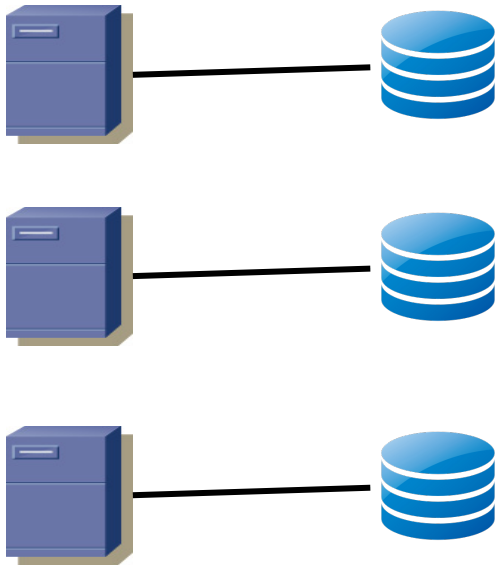

Acceso al almacenamiento

- El acceso al almacenamiento puede ser directo
- O puede ser a través de una red
  - Hablamos en ese caso de una *Storage Area Network (SAN)*
  - O de *Network Attached Storage (NAS)*
  - Estamos hablando del *front-end* de acceso a los discos
  - El *back-end* es común que sean discos SAS o SATA

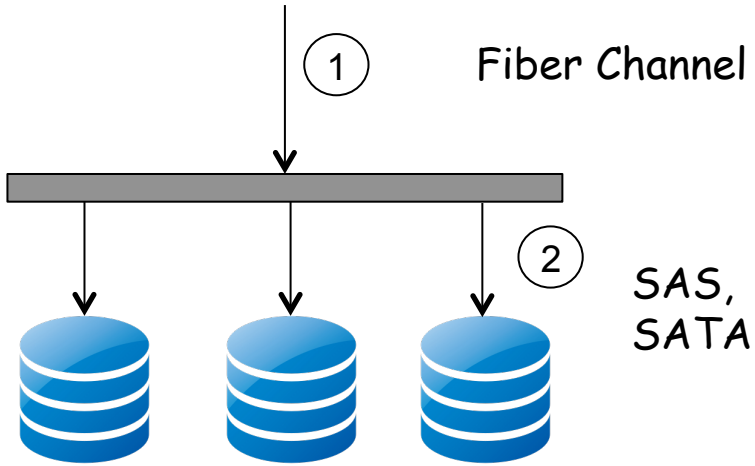

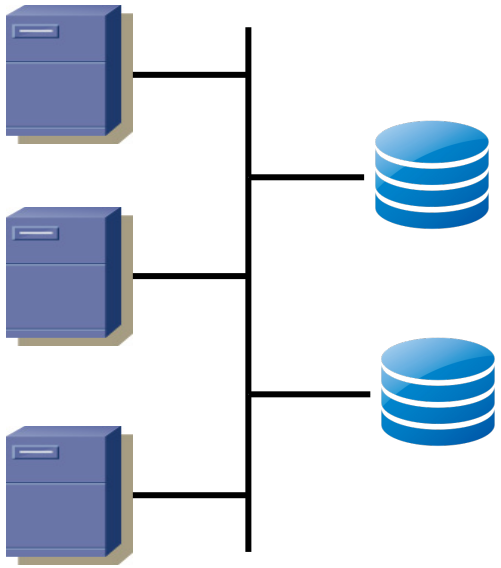

SAN ≠ NAS

En una SAN

- Se accede de forma serie a bloques de disco
- Normalmente mediante comandos SCSI-3
- Los protocolos están optimizados para baja latencia y nulas pérdidas
- La solución de transporte más habitual es Fibre Channel
- Acceso de varios servidores al mismo volumen requiere sistemas de ficheros especiales

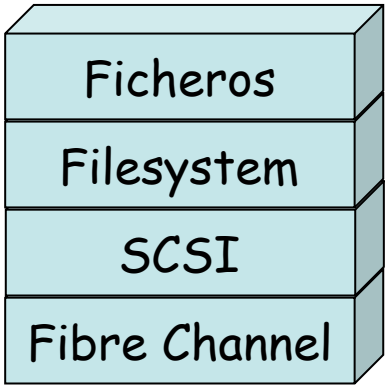

En una solución NAS

- Se accede a ficheros
- Se suele transportar sobre una tecnología LAN (o LAN + IP)
- Los protocolos no garantizan baja latencia ni nulas pérdidas (su recuperación aumenta la latencia)
- NFS, SMB/CIFS, AFP, etc

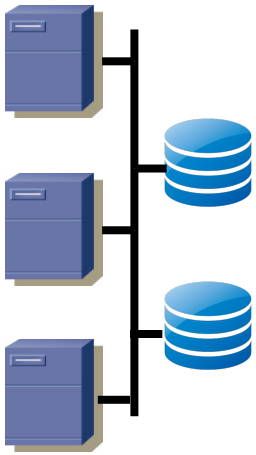

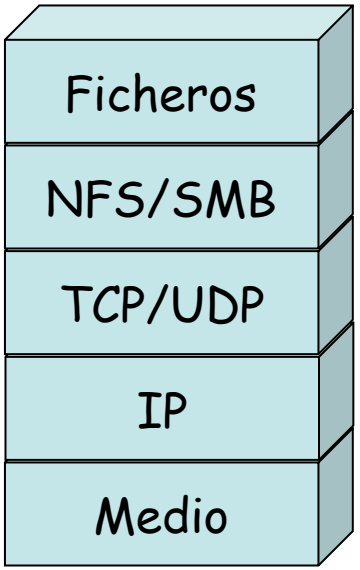

SAN vs NAS

- Rendimiento haciendo backups
  - Mayor en SANs, generalmente más rápida
  - En algunos casos se pueden mover datos del disco de un servidor a una cabina sin intervención de la CPU
  - En sistemas de ficheros con gran cantidad de ellos es más eficiente una copia del dispositivo *raw*
- Bases de datos
  - Algunas requieren un acceso al disco a nivel de bloques por rendimiento e integridad

¿ Dónde encaja la SAN ?

- Estamos hablando de una tecnología de red
- Cuenta con sus propios conmutadores
- Se dice que forman un *fabric*
- Cuenta con su propia pila de protocolos
- El interfaz en el host es el *Host Bus Adapter (HBA)* que sería el equivalente a la NIC en una LAN

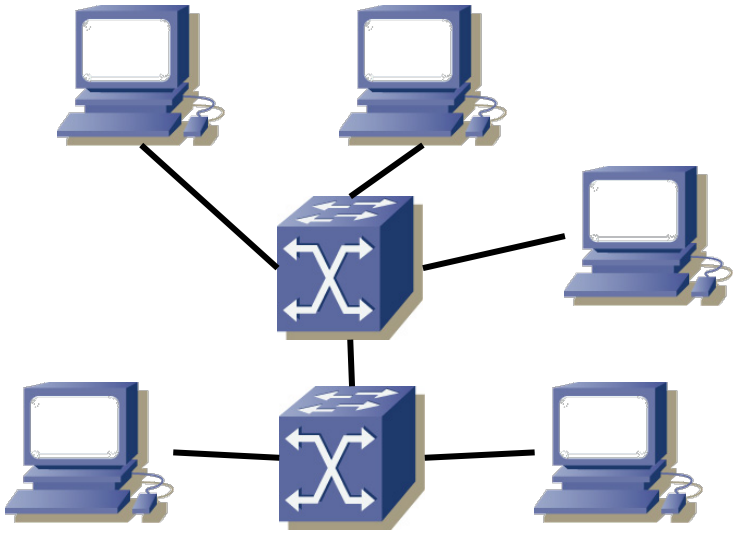

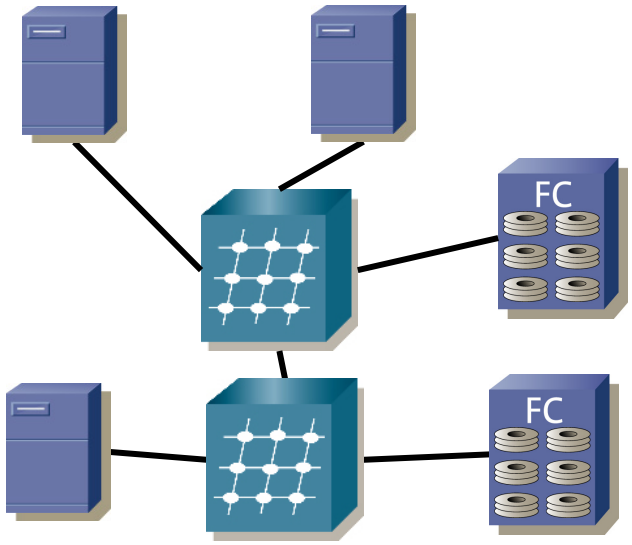

Es decir…

- Y recordemos que lo estamos poniendo en el back-end del servicio porque es donde suelen estar los datos
- Pero nada impide que los servidores de cualquier otra capa empleen almacenamiento SAN
- O incluso en la misma SAN (…)

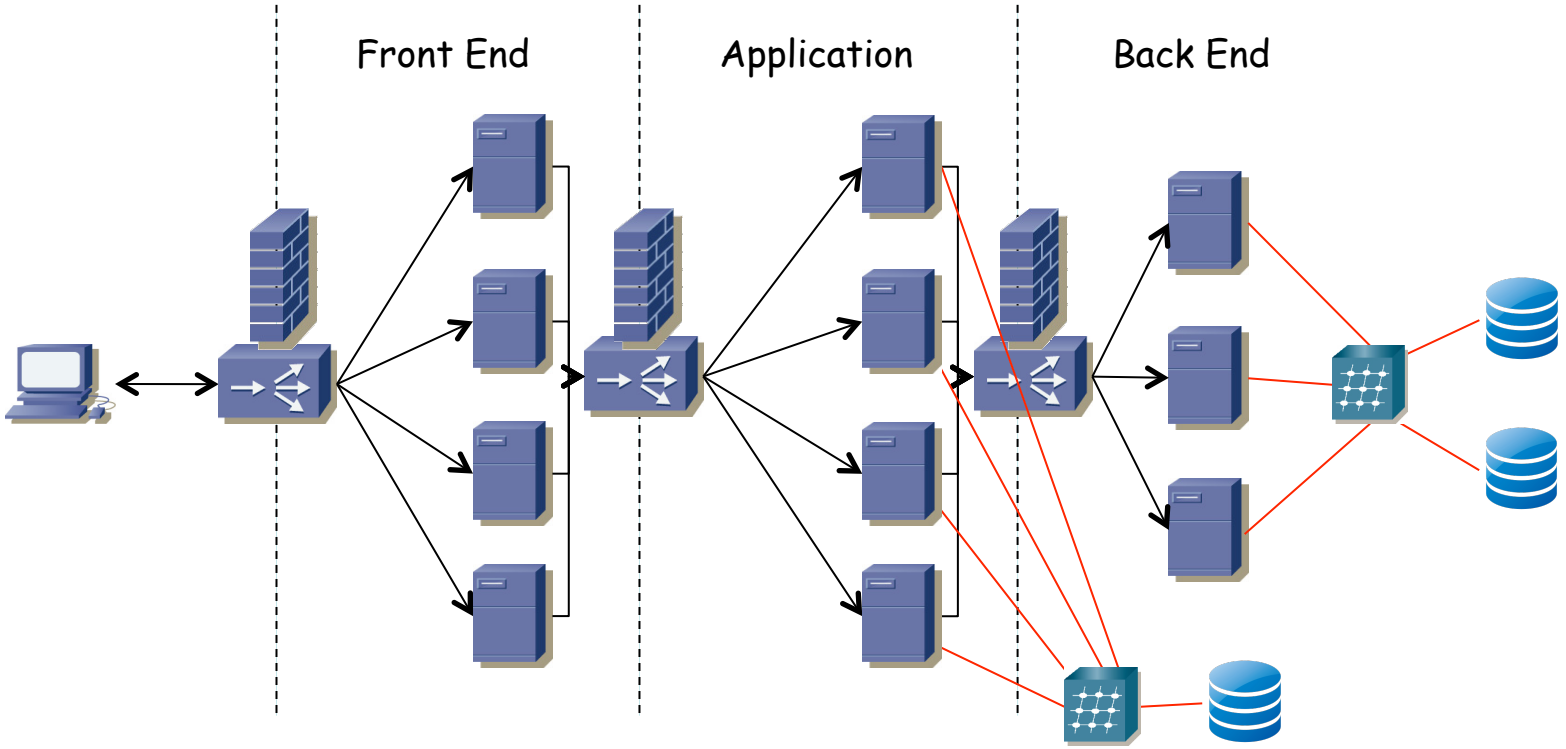

High Availability

- Y no hay que olvidarse de la redundancia
- (Retiro los interfaces de los servidores de la capa de aplicación por claridad)

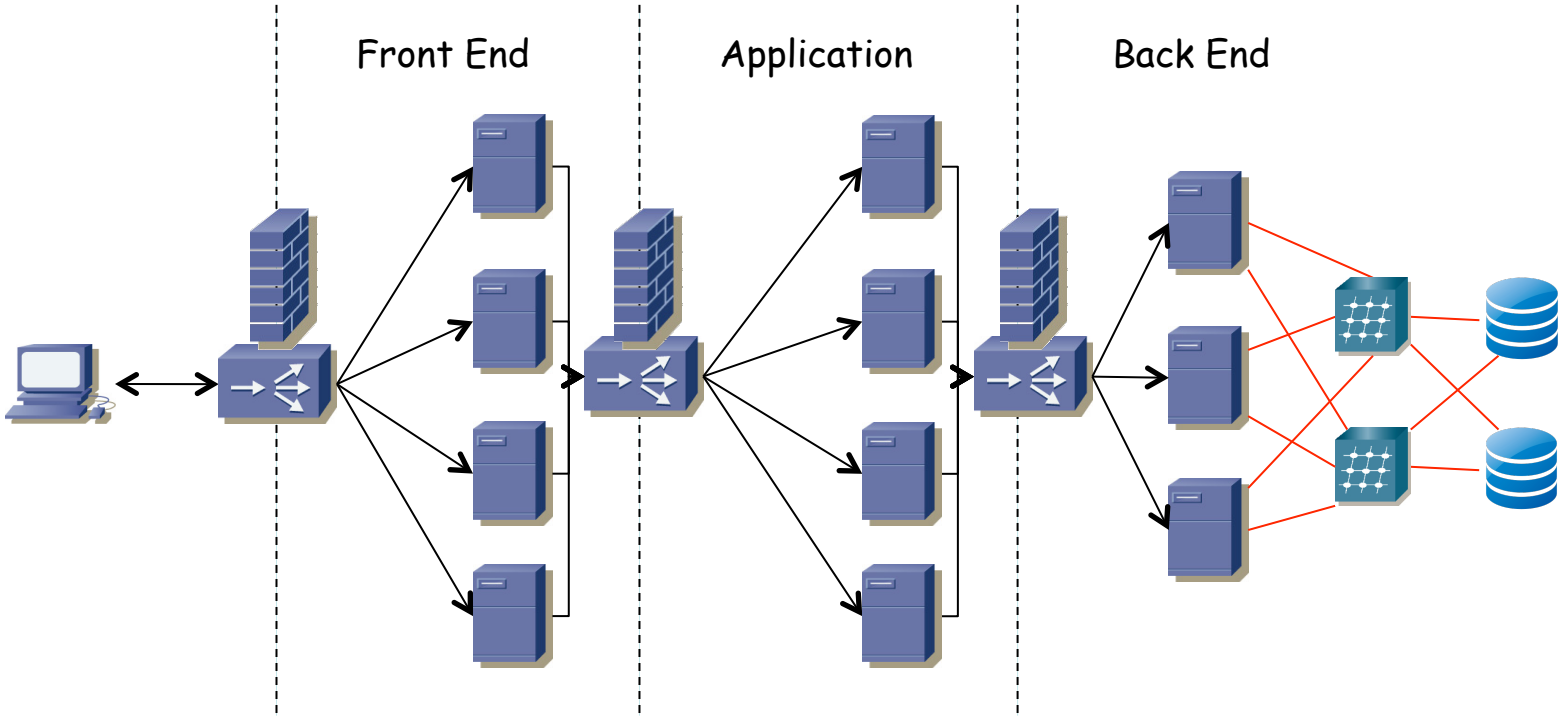

Diseño

- ¿Dónde encaja esa(s) SAN(s) en la red del data center?
- Es independiente, pero hablaremos más sobre esto

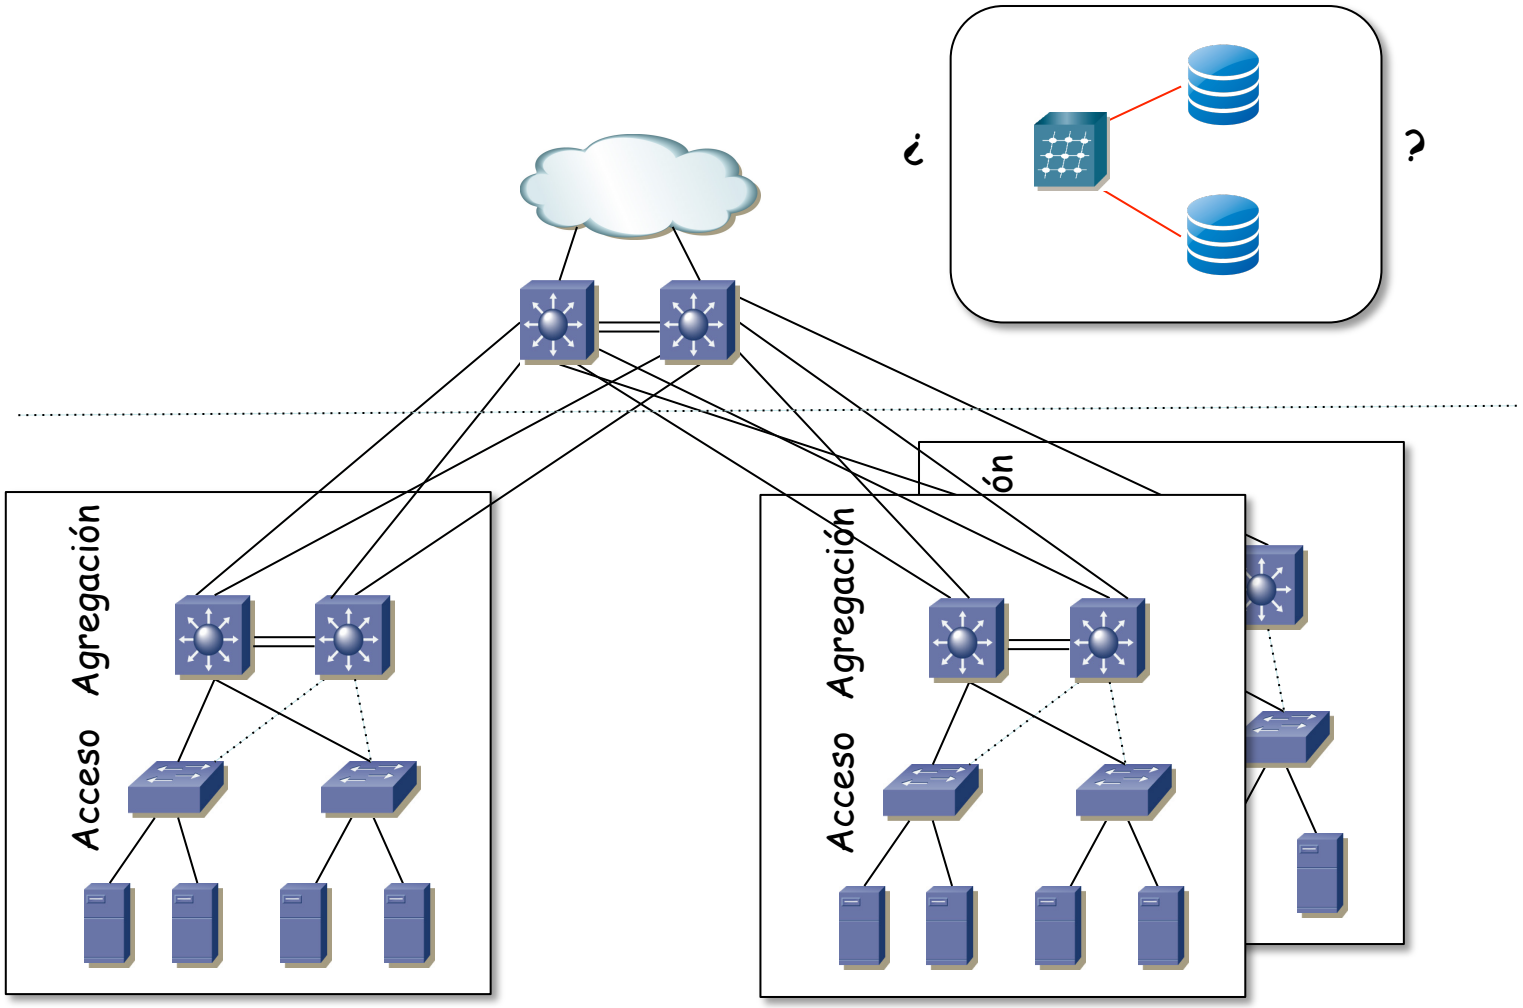

upna Universidad Pública de Navarra Nafarroako Unibertsitate Publikoa
Redes de Nueva Generación
Área de Ingeniería Telemática

FC: Arquitectura

Fibre Channel

- Desarrollo comenzado a finales de los 80s
- Lo normal es que sea sobre fibra pero puede ser sobre cobre
- "*Fibre*" es el *spelling* británico en lugar del americano para "fibra"
- Soluciona el transporte pero no fija lo que transporta
- Así, puede transportar comandos SCSI pero también IP o ATM
- Inicialmente una de sus ventajas era su alta velocidad en comparación con las tecnologías LAN de la época

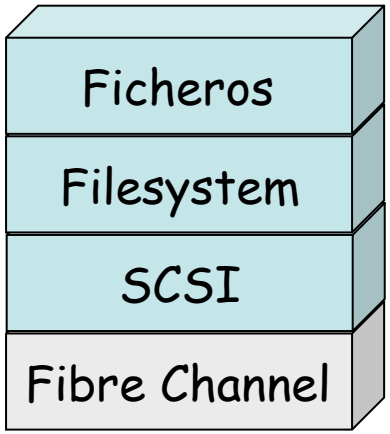

Fibre Channel

A fecha de 2014

- Para los nodos el comité T11 del INCITS tiene estandarizado: 1GFC, 2GFC, 4GFC, 8GFC, 16 GFC, 32GFC y 128GFCp (que son 128Gbps mediante 4 canales)
- Para la conexión entre conmutadores tiene 10GFC, 20GFC, 40GFCoE (FC over Ethernet), 100GFCoE y 128GFCp
- En el *roadmap* se habla de hasta 1TFC

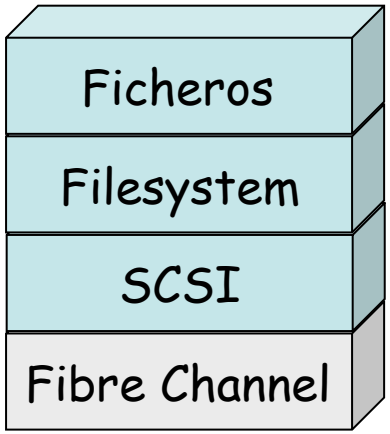

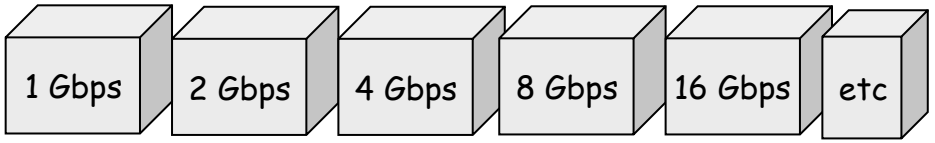

Fibre Channel

- FC-0: Capa física
- FC-1: Codificación, sincronización, control de errores
- FC-2: Formato de trama, señalización para gestión
- FC-3: Ofrece un conjunto único de servicios aunque por debajo haya varios puertos físicos (*name, login, address manager, alias server, fabric controller, management, key distribution, time*)
- FC-4: Capa de adaptación para protocolos superiores como puede ser SCSI o IP (ULP = Upper Level Protocol)

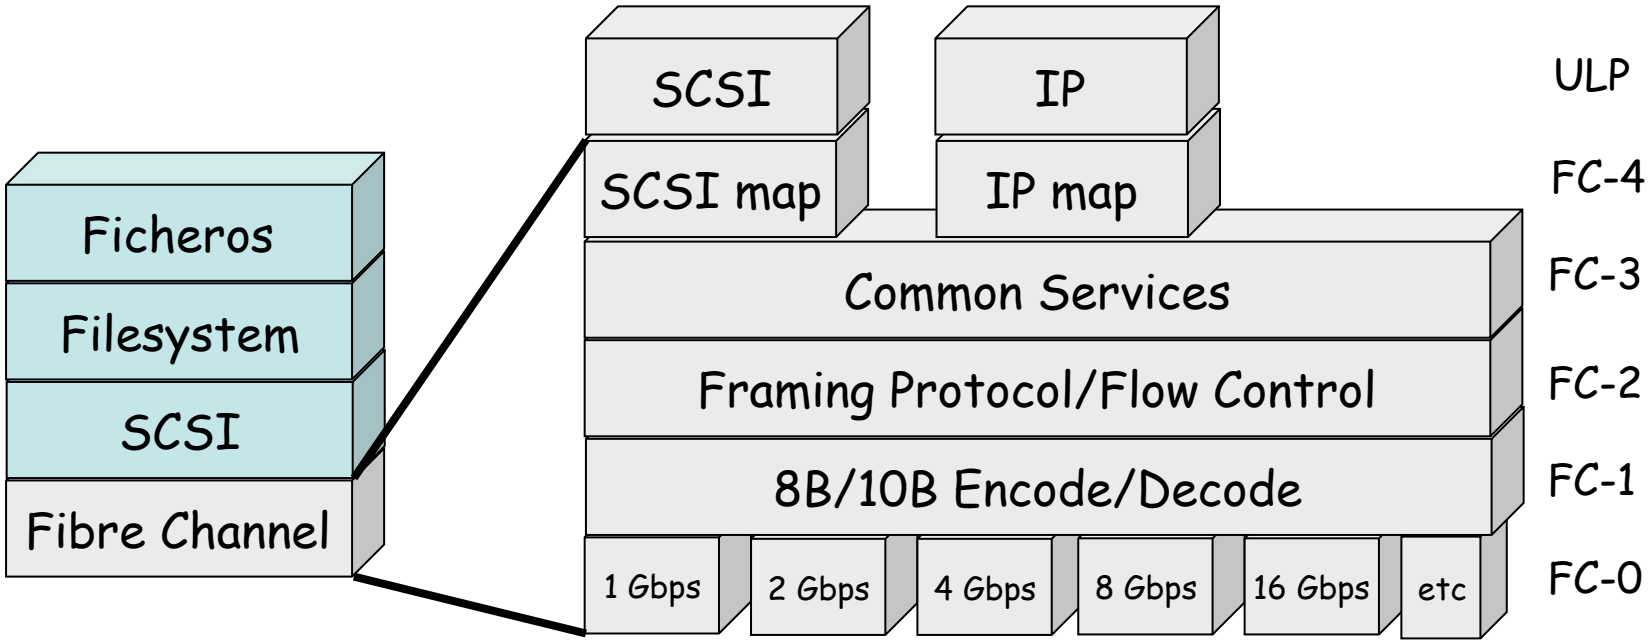

FC: Topologías

Topologías

Point-to-Point

- No hace falta un conmutador o crear “una red”
- Igual que en una Ethernet, podríamos tener simplemente un enlace punto-a-punto
- Ganamos algunas de sus características técnicas pero desde luego no la de compartir el uso de los discos
- Es una conexión directa entre los puertos de 2 nodos, que se vienen a llamar “N_Ports”

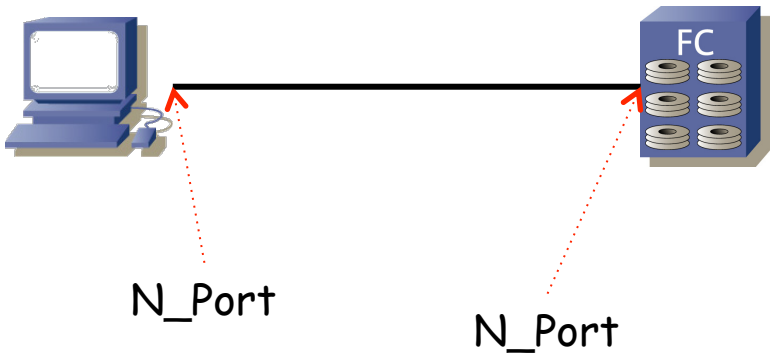

Topologías

Arbitrated loop (FC-AL)

- Anillo compartido
- 2-126 dispositivos pero muchos menos por rendimiento
- Los puertos se llaman “L_Ports” (NL_Port o FL_Port)
- Se negocia cuál de los nodos actúa como *master*
- Un nodo establece un circuito entre dos puertos que monopoliza el anillo para poder comunicarse
- No puede hacerlo de nuevo hasta haber dado turno a todos los demás
- (…)

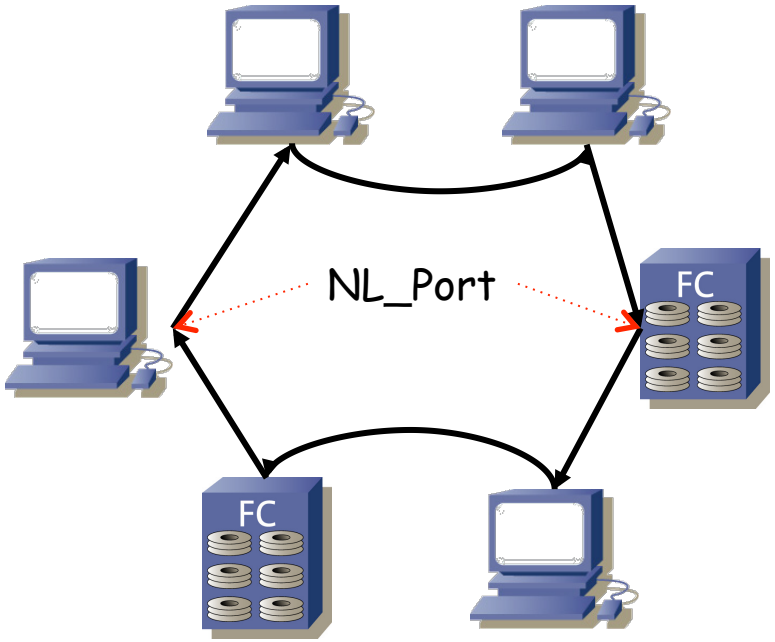

Topologías

Arbitrated loop (FC-AL)

- Anillo compartido
- 2-126 dispositivos pero muchos menos por rendimiento
- Los puertos se llaman “L_Ports” (NL_Port o FL_Port)
- Se negocia cuál de los nodos actúa como *master*
- Un nodo establece un circuito entre dos puertos que monopoliza el anillo para poder comunicarse
- No puede hacerlo de nuevo hasta haber dado turno a todos los demás
- También hub

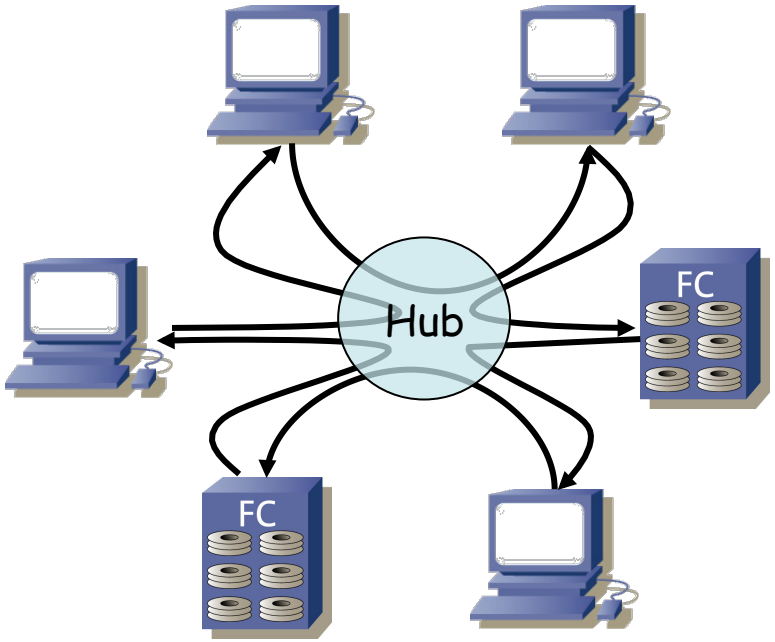

Topologías

Crosspoint switched (fabric)

- Uno o más conmutadores, interconectando múltiples nodos
- El direccionamiento permite en teoría hasta 2^{24} nodos
- "F_Port" (Fabric Port) a los nodos, "E_Port" (Expansion) entre switches
- ISL = *Inter-Switch Link*
- "FL_Port" para conectar a un FC-AL
- "G_Port" puerto genérico que se comporta según lo que se le conecte
- Y otros (propietarios)

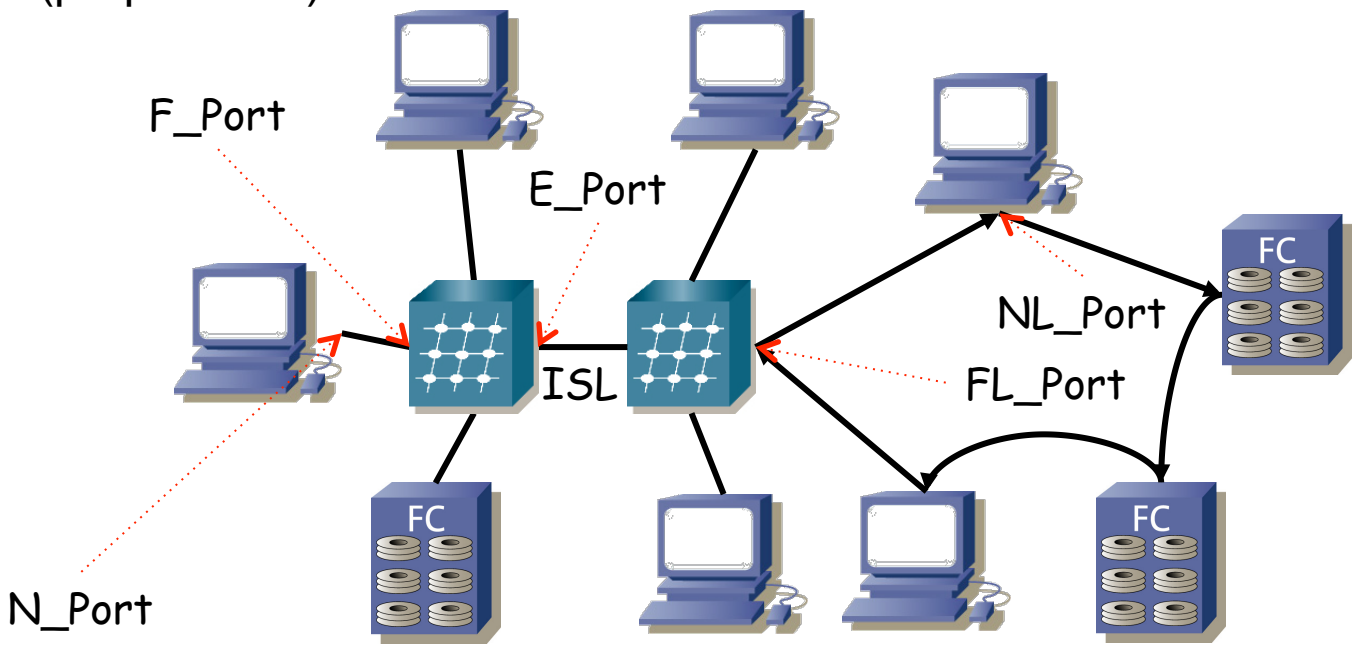

FC: Service Classes

Service Classes

- Diferentes “clases de servicio” (Class 1, Class 2, etc)
- Con circuitos virtuales (dedicados o compartidos) o datagramas
- Con ACKs y NACKs o sin ellos
- Con control de flujo en cada salto y/o extremo a extremo
- Con garantía de orden o no
- Con control de flujo
  - *Credit Based*
  - Básicamente una ventana deslizante
  - Mejora el rendimiento en situaciones de alta carga

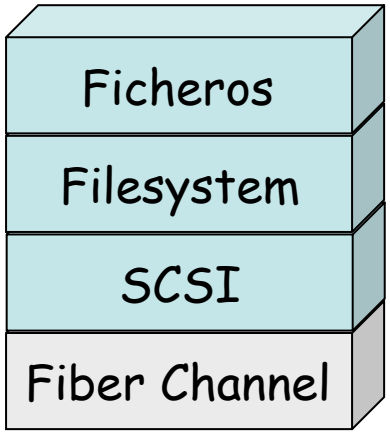

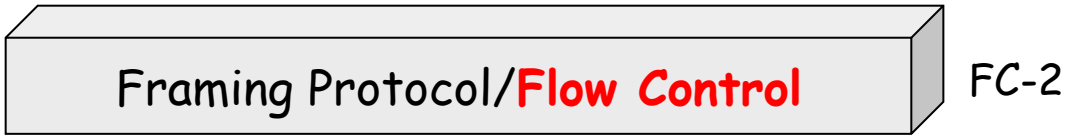

Service Classes

- Class 1 (*Dedicated connection)*
  - Se establecen un circuito dedicado entre los dos nodos
  - Durante el mismo los enlaces están dedicados así que no requiere control de flujo en cada salto (solo extremo a extremo)
  - Garantiza la máxima capacidad entre los Nx_Ports que se comunican
  - Garantiza entrega en orden de las tramas
  - El *fabric* no requiere buffer
  - Empleado para servicios continuos y sensibles al retardo (voz, vídeo)

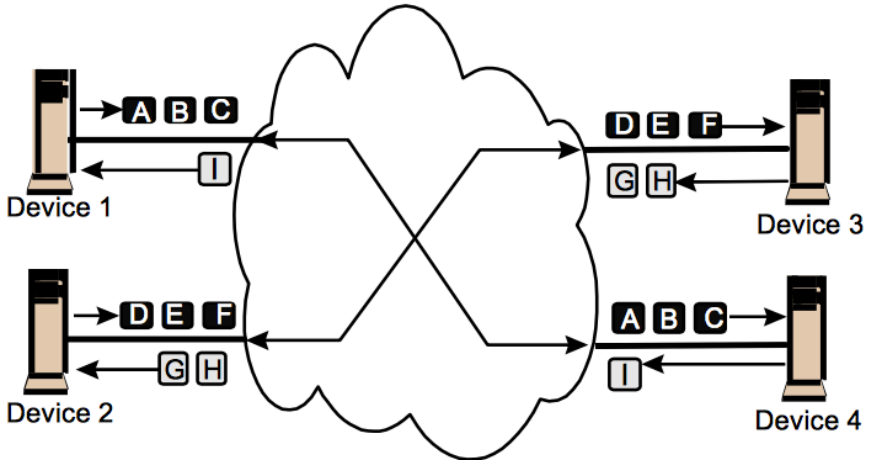

http://www.brocadechina.com/download/resource/ClassofService.pdf

Service Classes

- Class 2 (*Multiplex)*
  - Sin conexión, no garantiza entrega en orden
  - Emplea ACKs y NACKs
  - Se pueden estar recibiendo en un puerto tramas de varios nodos, así que comparten la capacidad (multiplexación)
  - Las velocidades de los puertos pueden ser diferentes así que requiere control de flujo en cada salto y extremo a extremo

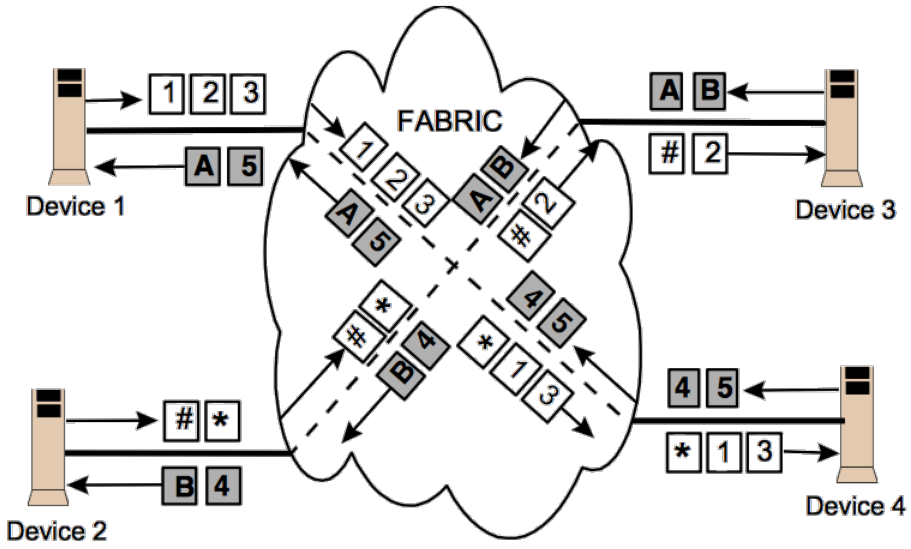

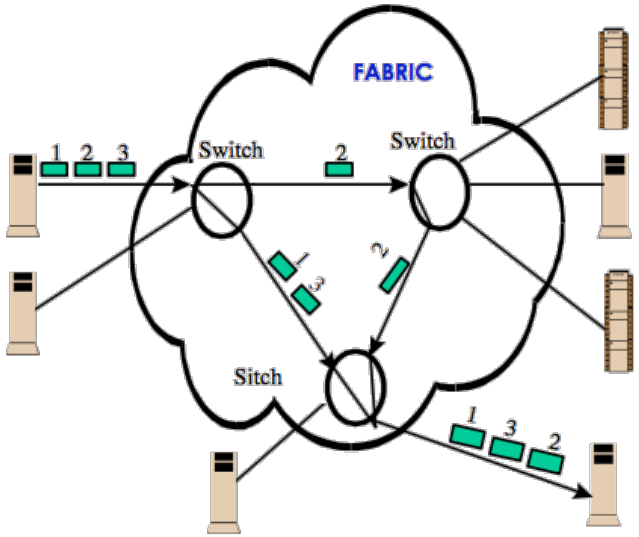

Service Classes

- Class 3
  - Sin conexión, no garantiza entrega en orden
  - Sin ACKs ni NACKs, lo debe hacer el ULP
  - Hace control de flujo salto a salto
  - Podría haber pérdidas con elevada congestión

- Class 4 (*Fractional bandwidth)*
  - Se establece un circuito virtual entre los dos nodos que empleará una fracción de la capacidad disponible
  - Un N_Port puede tener más de un circuito clase 4 mientras que solo podía tener uno de clase 1 (tomada toda la capacidad)
  - Garantiza entrega en orden
  - Emplea ACKs y NACKs
  - Hace control de flujo salto a salto

Service Classes

- Class 5 (¿sin definir?)
- Class 6 (*Multicast connection)*
  - Como un clase 1 pero con múltiples destinos
  - Unidireccional
  - Un servidor en el *fabric* se encarga de la replicación
- Class 9, Class F…
- Clases 2 y 3 las más habituales para sistemas de almacenamiento

FC: Otros servicios

Multipath

- Permite caminos redundantes
- Varios HBAs en host y varios interfaces en sistema de almacenamiento
- FC dispone del protocolo FSPF (*Fibre Channel Shortest Path First)*
- No bloquea caminos alternativos sino que puede repartir carga
- O pueden servir como caminos redundantes ante fallos
- El sistema operativo del host requiere software que le permita ver los dos accesos al volumen como un solo volumen

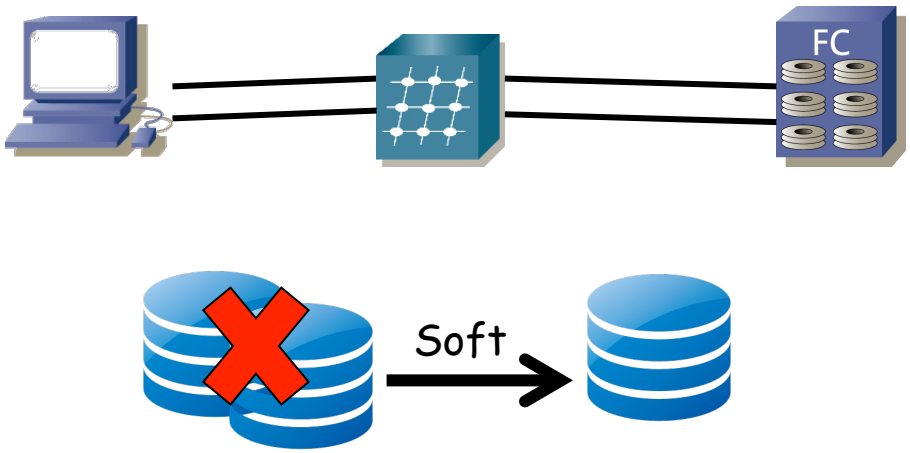

Masking/Zoning

- Tradicionalmente en un bus SCSI solo había un elemento que enviaba comandos al resto
- En realidad el protocolo permite a un dispositivo comunicarse con varios hosts
- Pero muchos sistemas operativos esperan tener el control exclusivo de cualquier dispositivo que ven en el bus

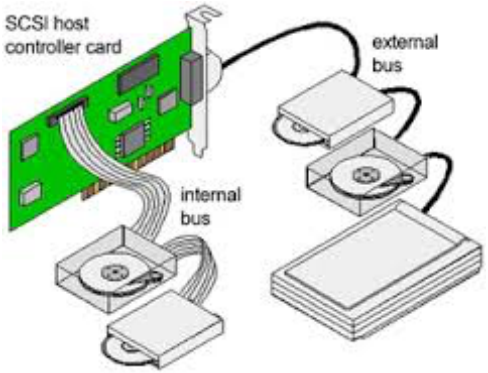

Masking/Zoning

- En una SAN múltiples hosts pueden acceder al mismo dispositivo de almacenamiento
- Se puede limitar mediante *zoning*
- Una zona es como una SAN privada virtual de forma que se limita la comunicación a los dispositivos en la misma
- En el caso *Hard zoning* se crea la zona con una lista de puertos físicos de conmutador
- En el caso *Soft zoning* se emplean las direcciones de los dispositivos (WWNs, *World Wide Names*) para limitar el acceso
- *LUN masking* funciona por encima y permite que aunque el host tenga acceso a la cabina no lo tenga a todas las LUNs
- Muchas soluciones propietarias

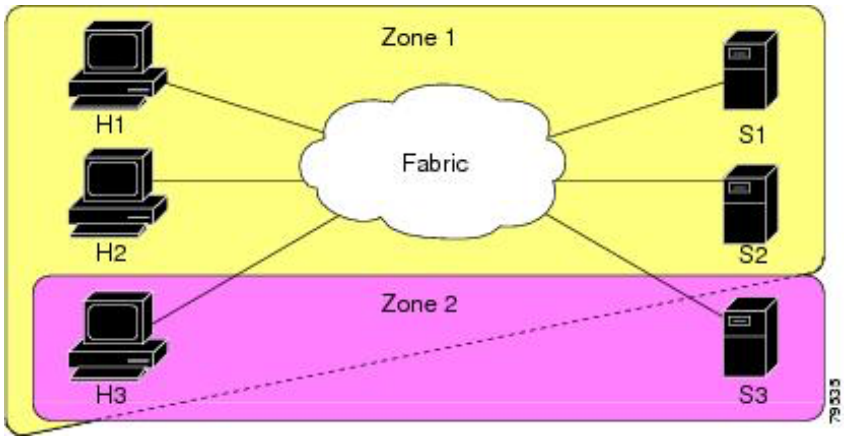

NAS

Acceso a ficheros

- Solución cliente-servidor, hoy en día sobre TCP/IP
- Integrado con el sistema operativo en el cliente
- Eso permite que las aplicaciones vean el sistema de ficheros remoto como si fuera local
- Ejemplos de protocolos: SMB, NFS, AFP
- Otras alternativas menos frecuentemente integradas en el cliente: FTP, HTTP (WebDav), SSH/SFTP
- El protocolo da acceso a ficheros y directorios, no a bloques

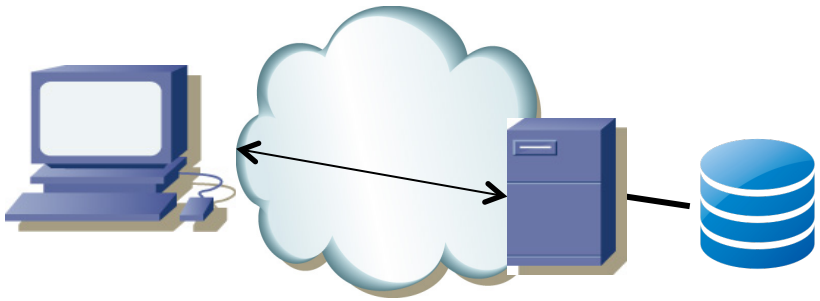

Acceso a ficheros

- Inicialmente el servidor es un hardware de propósito general que comparte parte de su disco
- Protocolos más habituales
  - SMB
  - NFS

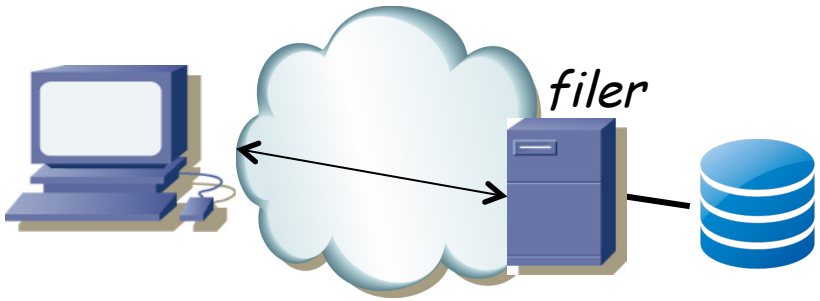

SMB

- *Server Message Block*
- Desarrollado por IBM y Microsoft
- Principalmente en sistemas Windows
- Empleado para resolución de nombre, navegar recursos compartidos, compartición de ficheros, acceso a impresoras, autentificación, etc
- Sobre TCP, UDP y otros
- SMB, CIFS (*Common Internet File System*), SMB2, SMB3, …

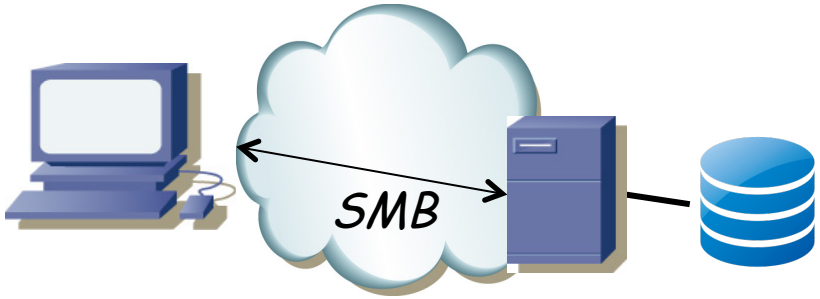

NFS

- *Network File System*
- Principalmente en sistemas UNIX (desarrollado por Sun Microsystems)
- Emplea los protocolos XDR (External Data Representation, RFC 1832) y RPC (Remote Procedure Call, RFC 1831)
- Muy común para centralizar los *home directories*
- mountd, nfsd
- NFSv2 (RFC 1094) sin estado (UDP), tamaños de 32 bits
- NFSv3 (RFC 1813) sin estado (UDP aunque también TCP), tamaños de 64 bits
- NFSv4 (RFC 3530) con estado (TCP), seguridad
- NFSv4.1 (RFC 5661) introduce *parallel NFS* (pNFS)

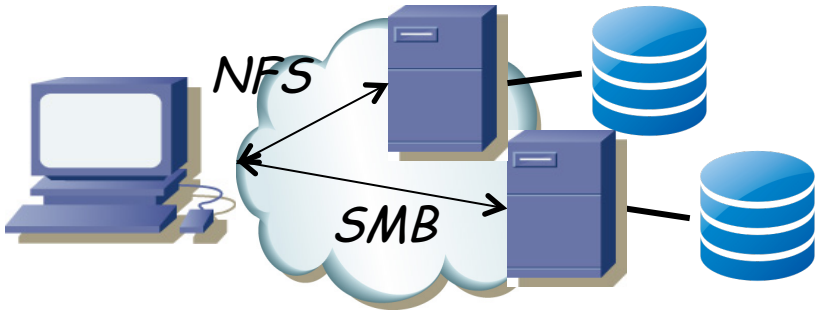

NAS

- NAS = *Network Attached Storage*
- Inicialmente el servidor es un hardware de propósito general
- Hoy en día hardware dedicado (*appliances*) para esta tarea
- Optimizado para ello
  - Simplica la gestión
  - RAIDs, discos *hot-swappable*
  - Capaces de compartir los mismos recursos mediante varios protocolos simultáneamente (SMB + NFS)
  - Mejoras de rendimiento en la implementación del soft servidor
- El disco en vez de ser local puede estar en una SAN

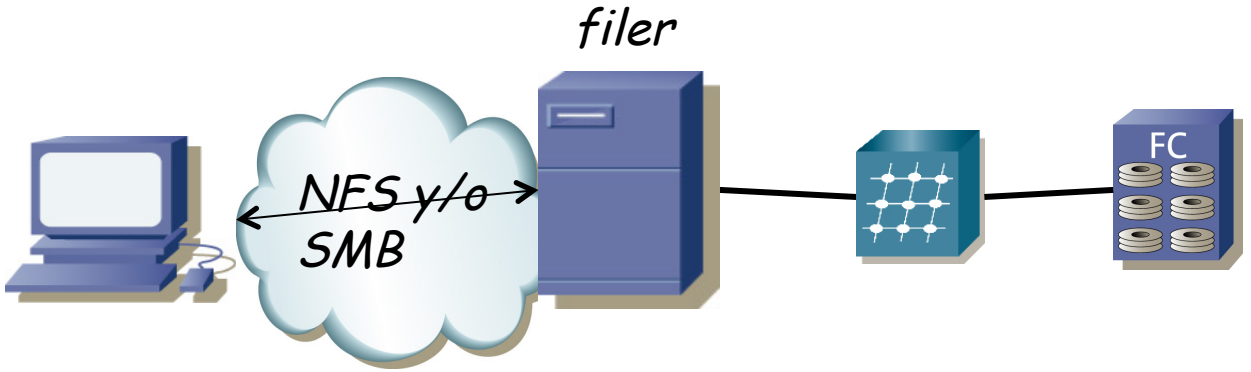

Equipos

Equipos FC

Dell SC4020

Internal Storage
24 x 2.5" drive bays

Supported expansion enclosures
Dell SC200: 12 x 3.5" drive bays
Dell SC220: 24 x 2.5" drive bays

Maximum drive count
120 (24 internal, plus 96 external)

Total storage capacity
408TB based on maximum number of drives (120) and current largest capacity supported drives

Supported drive types
HDD: 15K, 10K, 7.2K RPM; SDD: write-intensive SLC, read-intensive MLC (different drive types, transfer rates and rotational speeds can be mixed in the same system

Controllers
2 controllers per SC4020 array

Processor
Intel® Xeon® Processor E3-1265L v2, 2.50GHz, 4 cores

Memory
32GB per SC4020 array (16GB per controller)

Network/server connectivity (front-end)
8 x 8Gb FC ports per SC4020 array (4 per controller)

Internal drive connectivity (back-end)
4 x 6Gb SAS ports per SC4020 array (2 per controller)

NAS deployment connectivity (optional)
Supports file-based storage via FS8600 NAS appliance

RAID
Supports RAID 0, 5, 6, RAID 10 and RAID 10 DM (dual mirror). Any combination of RAID levels can exist on a single Storage Center. Multiple RAID levels can exist on the same storage tier within an array.

IBM DS8000

Models	DS8870 (961, 96E)
Shared SMP processor configuration	IBM POWER7+ dual 2-core, 4-core, 8-core or 16-core
Other major processors	IBM PowerPC®, ASICs
Processor memory for cache and nonvolatile storage (minimum/maximum)	16 GB/1,024 GB
Host adapter interfaces	4- and 8-port 8 Gbps Fibre Channel/IBM FICON
Host adapters (minimum/maximum)	2/16
Host ports (minimum/maximum)	8/128
Drive interface	6 Gbps point-to-point switched SAS-2 connection to an 8 Gbps Fibre Channel backbone
Device adapters	Up to 16 4-port, 8 Gbps Fibre Channel paths
Maximum physical storage capacity*	3,072 TB
Disk sizes	400 GB flash cards (High-Performance Flash Enclosure) 200 GB flash drives 800 GB flash drives 1.6 TB flash drives 146 GB (15k rpm) 300 GB (15k rpm) 600 GB (15k rpm) 600 GB (10k rpm) 1.2 TB (10k rpm) 4 TB (7.2k rpm and 3.5-inch form factor)
RAID levels	5, 6, 10
Dimensions (height × width × depth)	193.4 cm × 84.8 cm × 122.7 cm (76.14 in. × 33.39 in. × 48.31 in.) per frame, up to four frames total

Brocade 6510 Switch

Fibre Channel ports	Switch mode (default): 24-, 36-, and 48-port configurations (12-port increments through Ports on Demand [PoD] licenses); E, F, M, D, EX ports Access Gateway default port mapping: 40 F_Ports, 8 N_Ports
Scalability	Full fabric architecture with a maximum of 239 switches
Certified maximum	6000 active nodes; 56 switches, 19 hops in Brocade Fabric OS® fabrics; 31 switches; larger fabrics certified as required
Performance	Fibre Channel: 2.125 Gbps line speed, full duplex; 4.25 Gbps line speed, full duplex; 8.5 Gbps line speed, full duplex; 10.53 Gbps line speed, full duplex; 14.025 Gbps line speed, full duplex; auto-sensing of 2, 4, 8, and 16 Gbps port speeds; 10 Gbps and optionally programmable to fixed port speed
ISL trunking	Frame-based Trunking with up to eight 16 Gbps ports per ISL trunk; up to 128 Gbps per ISL trunk. Exchange-based load balancing across ISLs with DPS included in Brocade Fabric OS.
Aggregate bandwidth	768 Gbps end-to-end full duplex
Maximum fabric latency	Latency for locally switched ports is 700 ns; encryption/compression is 5.5 µsec per node; Forward Error Correction (FEC) adds 400 ns between E_Ports (enabled by default).
Maximum frame size	2112 byte payload
Frame buffers	8192 dynamically allocated
Classes of service	Class 2, Class 3, Class F (inter-switch frames)
Port types	D_Port (ClearLink Diagnostic Port), E_Port, EX_Port, F_Port, M_Port (Mirror Port); optional port type control Brocade Access Gateway mode: F_Port and NPIV-enabled N_Port
Data traffic types	Fabric switches supporting unicast
Media types	16 Gbps: Brocade 6510 requires Brocade hot-pluggable SFP+, LC connector; 16 Gbps SWL, LWL, ELWL 10 Gbps: Brocade 6510 requires Brocade hot-pluggable SFP+, LC connector; 10 Gbps SWL, LWL 8 Gbps: Brocade 6510 requires Brocade hot-pluggable SFP+, LC connector; 8 Gbps SWL, LWL, ELWL

Equipos NAS

Hitachi NAS Platform F1100 Series

Outstanding Value, Flexible and Easy to Manage

Hitachi NAS Platform (HNAS) F1100 series are entry-level members of the Hitachi NAS Platform family. They range from a self-contained 2U NAS model with 12 internal drives, to an enterprise-ready 2-node clustered NAS system with external drive arrays. They offer automatic failover, NAS virtualization, IPv6 and a 1PB file system. The series provides outstanding business value for smaller budgets. Flexibility of data protection and replication capabilities make it ideal for midsized organizations, remote or branch offices (ROBOs), or enterprises needing data protection, backup or file server consolidation.

Data Protection and File Sharing Challenges

We understand your challenges when it comes to file sharing and data protection for small enterprises, and replication from remote locations. For example, general-purpose file servers might hit their capacity limit without having a scalable upgrade path, resulting in higher maintenance and support costs. Or, many remote locations might lack a uniform method for managing data protection. For each of your challenges, Hitachi can provide a NAS solution to satisfy your data access and protection needs for Microsoft® Windows® and UNIX or Linux clients.

Business Benefits

The HNAS F1100 series provides outstanding value with a rich set of enterprise-class features at an affordable price. For example, it allows you to reclaim primary storage capacity with single instancing to remove file duplicates. And the HNAS F1140 model in particular lets you automatically place the right data on the right storage tier for optimal performance using dynamic file tiering. HNAS F1100 series also provides investment protection with easy migration and replication to higher-end HNAS 4000 models, especially for ROBOs.

Deployment Options

Within the Hitachi NAS Platform F1100 series, models are offered as preconfigured systems. HNAS F1140 nodes are clustered together to provide a symmetric active-active file server. When clustered, HNAS F1140 is configured with external storage (HUS 110), management server and Ethernet switch. The F1120 is standalone with storage and management software built-in.

Features

Hitachi NAS Platform F1100 series presents a standards-based file system interface that is feature rich, reliable and very flexible. Some of its key features include:

Scalability and Availability

- Scale up to 480TB capacity in a highly available 2-node F1140 cluster.
- Anticipate future growth with 1PB file systems and single instancing.

Advanced Data Protection Capabilities

- 992 read-only snapshots.
- Asynchronous replication between HNAS F1100 nodes.

Model	Hitachi NAS Platform F1140 Cluster	Hitachi NAS Platform F1140 Single-Node	Hitachi NAS Platform F1120 Single-Node
Configuration	2x 1U rack server + 2U Hitachi Unified Storage (HUS) 110 + 1U IP switch	1U rack server + HUS 110	2U rack server with up to 12x HDD slots
Memory	32GB	16GB	12GB
Capacity (data)	Up to 480TB	Up to 480TB	Up to 36TB
I/O Ports	8 x 1GbE and/or 4 x 10GbE	4 x 1GbE or 2 x 10GbE	4 x 1GbE or 2 x 10GbE

Hitachi NAS Platform F1100 Series

HITACHI NAS PLATFORM F1100 SERIES TECHNICAL SPECIFICATIONS

Architecture	F1140 cluster: 2x 1U with Intel E5 1.8GHz, with 4 cores per node. F1120 single: 1x 2U with Intel E5 1.9GHz with 6 cores per node. Redundant hot-pluggable fans and power supplies. Hitachi Unified Storage 110 (2U) with 12 x 3.5 in. or 24 x 2.5 in. HDD slots for F1140.Management server (1U). Ethernet switch (1U) internal management and maintenance.
Capacity (raw)	480TB NL-SAS (120 x 4TB disks in external shelf); 108TB SAS (using 900GB disks); 9.6TB SSD (using 400GB disks).
Performance	F1140: 21,000 SPECsfs_2008 for CIFS IOPS; 24,000 concurrent SMB sessions.
Thin Provisioning	Hitachi Dynamic Provisioning.
Maximum File System Size	1PB per file system, up to 256 file systems.
NAS Virtualization	Up to 4 instances.
Deduplication	File-level single instancing.
Tiered Storage	Hitachi Dynamic Tiering provides 2-level internal storage tiering.
Snapshots	992 per file system; read-only.
Data Migration	NAS migration utility (3rd-party NAS to HNAS F1140).
File Migration	Optional capability to enable connection and migration of files to Hitachi Content Platform (HCP).
Antivirus Protection	ICAP for Symantec, McAfee and Trend Micro.
Replication	HNAS F-to-HNAS F using file remote replicator software.
Management	Integrated with Hitachi Command Suite; initial setup wizard; management interface (GUI, SSL) and CLI over SSH; Simple Network Management Protocol (SNMP); Network Time Protocol (NTP); Distributed file system (DFS).
Security and Compliance	LDAP v3; POSIX ACLs, NTFS ACLs; HTTPS for communication with HCP over WAN; External Syslog server; RADIUS support.
Backup	NDMP v4; Backup and restore via HCP.
Protocols Supported	CIFS/SMBv1.0, SMBv2.0; NFSv2, v3, v4; FTP, SFTP.
Network	Network link aggregation (IEEE 802.3ad); tagged VLAN (IEEE802.1q); link alternation.

NetApp FAS8080 EX

Powered by NetApp Data ONTAP®, the FAS8080 EX is optimized at every level for enterprise applications. Forty processor cores, 256GB of high-speed DRAM, capacity for 1,440 drives, and 144TB of hybrid flash acceleration are balanced together in a high-availability (HA) pair to reliably process the massive amounts of data in a modern enterprise. With 16 onboard I/O ports as well as 24 PCIe 3.0 expansion slots, serving data to applications has never been easier. Combine this with world-class data management from Data ONTAP, and you have a system that can reduce the time it takes to complete critical operations, increase your organization's competitive advantage, and keep your enterprise applications running at top speed 24/7/365.

Performance for Extraordinary Throughput

Our leading Intel® multiprocessor architecture with its high-bandwidth DDR3 memory system maximizes throughput for business-critical workloads such as SAP®, SQL Server®, and Oracle® databases as well as computational modeling and logistics management applications. Building on a decade of multicore optimizations, Data ONTAP takes advantage of the FAS8080 EX's 40 processor cores to keep pace with growth in storage I/O demands for your most intensive SAN and NAS applications. Integrated unified target adapter (UTA2) ports support 16Gb FC and 10GbE (FCoE, iSCSI, SMB, NFS) so you're

NetApp FAS8080 EX

Capacity to Harness the Data Explosion

Storage scale-out to more than 17,000 drives and 4M IOPS of performance combined with the industry-leading management capabilities of Data ONTAP enables the FAS8080 EX to process petabytes of data for complex resource exploration, faster semiconductor chip design, or optimally managing logistics for billion-dollar-plus product operations. 103PB of data in a 24-node cluster can be managed from a single console to reduce the cost of operations and simplify warehousing years of business-critical data.

Identify Problems Before They Happen to Achieve Unparalleled Availability and Stringent SLOs

As with other business-critical arrays, the FAS8080 EX is designed to deliver 99.999% or greater availability through a comprehensive approach that combines highly reliable hardware, innovative software, and sophisticated service analytics. Advanced hardware features, including alternate control path (ACP), persistent NVRAM write logs, and an integrated service processor, provide the utmost reliability to protect your investment. All I/O devices, including embedded ports, can be independently reset, allowing the FAS8080 EX to detect, contain, and recover from faults.

NetApp FAS8080 EX

FAS8080 EX Technical Specifications

Scale-Out

	FAS8080 EX
NAS scale-out	1–24 nodes (12 HA pairs)
Maximum drives	17,280
Maximum raw capacity: hybrid[1]	103PB 1728TB flash
Maximum raw capacity: all-flash FAS	4.6PB
Maximum memory	3072GB
SAN scale-out	1–8 nodes (4 HA pairs)
Maximum drives	5,760
Maximum raw capacity: hybrid[1]	34PB 576TB flash
Maximum raw capacity: all-flash FAS	1.5TB
Maximum memory	1024GB
Cluster interconnect	2, 4, or 6 10GbE

NetApp FAS8080 EX

Per HA Pair Specifications (Active-Active Dual Controller)

	FAS8080 EX
Maximum drives	1,440
Maximum raw capacity: hybrid[1]	8640TB 144TB flash
Maximum raw capacity: all-flash FAS	384TB
Maximum Flash Cache™	24TB
Maximum Flash Pool™	144TB
Controller form factor	12U (2 enclosures)
ECC memory	256GB
NVRAM	32GB
PCIe expansion slots	24
Onboard I/O: UTA 2 (16Gb FC/FCoE/10GbE)	8
Onboard I/O: GbE	8
Onboard I/O: 10GbE	8
Onboard I/O: 6Gb SAS	8
OS version	Data ONTAP 8.2.2 and later
Shelves and media	See the Shelves and Media page[2] on *NetApp.com* for the most current information
Storage protocols supported	FC, FCoE, iSCSI, NFS, pNFS, CIFS/SMB, HTTP, FTP
Host/client operating systems supported	Windows® 2000, Windows Server® 2003, Windows Server 2008, Windows Server 2012, Windows XP, Linux®, Sun Solaris, AIX, HP-UX, Mac® OS, VMware®, ESX®

1. Hybrid configurations combine HDDs with flash using Flash Cache and/or Flash Pool. You can also create storage configurations that include separate HDD and SDD aggregates. The size limit of the SSD aggregates is the same as the all-flash maximum.
2. *netapp.com/us/products/storage-systems/disk-shelves-and-storage-media/index.aspx*

NetApp EP550 Flash Array

All-Flash FAS8000 Technical Specifications

Scale-Out

	FAS8080 EX	FAS8060	FAS8040	FAS8020
NAS scale-out		1–24 nodes (12 HA pairs)		
Maximum SSDs	2,880	2,880	2,880	2,880
Maximum raw capacity	4.6PB	4.6PB	4.6PB	4.6PB
Effective capacity[1]	16.3PB	16.3PB	16.3PB	16.3PB
Maximum memory	3072GB	1536GB	768GB	576GB
SAN scale-out		1–8 nodes (4 HA pairs)		
Maximum SSDs	960	960	960	960
Maximum raw capacity	1.5PB	1.5PB	1.5PB	1.5PB
Effective capacity[1]	5.4PB	5.4PB	5.4PB	5.4PB
Maximum memory	1024GB	512GB	256GB	192GB
Cluster interconnect	2, 4, or 6 10GbE	2 or 4 10GbE	2 or 4 10GbE	2 10GbE

Per HA Pair Specifications (Active-Active Dual Controller)

	FAS8080 EX	FAS8060	FAS8040	FAS8020
Maximum SSDs	240	240	240	240
Maximum raw capacity	384TB	384TB	384TB	384TB
Controller form factor	12U (2 enclosures)	6U	6U	3U
Effective capacity[1]	1388.5TB	1388.5TB	1388.5TB	1388.5TB
ECC memory	256GB	128GB	64GB	48GB
NVRAM	32GB	16GB	16GB	8GB
PCIe expansion slots	24	8	8	4
Onboard I/O: UTA 2 (16Gb FC/FCoE/10GbE)	8	8	8	4
Onboard I/O: GbE	8	8	8	4
Onboard I/O: 10GbE	8	8	8	4
Onboard I/O: 6Gb SAS	8	8	8	4
OS version	Data ONTAP 8.2.2 and later		Data ONTAP 8.2.1 and later	
Shelves and media	See the Shelves and Media page[2] on *NetApp.com* for the most current information			
Storage protocols supported	FC, FCoE, iSCSI, NFS, pNFS, CIFS/SMB, HTTP, FTP			
Host/client operating systems supported	Windows® 2000, Windows Server® 2003, Windows Server 2008, Windows Server 2012, Windows XP, Linux®, Oracle Solaris, AIX, HP-UX, Mac® OS®, VMware® ESX®			

1. Effective capacity is based on 5:1 storage efficiency ratios with the maximum number of SSDs installed. The actual ratio is up to 10:1 and higher, depending on workloads and use cases.
2. *netapp.com/us/products/storage-systems/disk-shelves-and-storage-media/index.aspx*

NetApp EF550 Flash Array

Attribute	NetApp EF550 Flash Array
Burst I/O rate	900,000 IOPS
Sustained I/O rate	500,000 IOPS
Sustained throughput	Up to 12GB/s
Maximum drives	120
Maximum raw capacity	192TB
Drive types supported	2.5" SSD 400GB, 800GB, 800GB (FDE), 1.6TB
Form factor	Base system: 2U//24* Expansion shelf: 2U/24*
System ECC memory	24GB
I/O interface options	(8) 16Gb FC, (8) 6Gb SAS, (8) 12Gb SAS, (8) 10Gb iSCSI, (4) 40 Gb Infiniband, or (4) 56Gb Infiniband
Operating system Management system	SANtricity OS 8.20 SANtricity Storage Manager 11.20
High-availability features	• Dual active controller with automated I/O path failover • Dynamic Disk Pools (DDP) and RAID levels 0, 1, 3, 5, 6, and 10 • Redundant, hot-swappable storage controllers, disk drives, power supplies, and fans • Automatic DDP or RAID rebuild following a drive failure • Mirrored data cache with battery backup and de-stage to flash • SANtricity proactive drive health monitoring identifies problem drives before they create issues • Greater than 99.999% availability (with appropriate configuration and service plans)
Host operating systems supported	Microsoft® Windows® Server, Red Hat Enterprise Linux®, Novell SUSE Linux Enterprise Server, Apple® Mac® OS, Oracle Solaris, HP, HP-UX, CentOS Linux, Oracle Enterprise Linux, IBM AIX, VMware® VSphere®
Included software features	SANtricity Mirroring, SANtricity Volume Copy, SANtricity Snapshot, SANtricity Thin Provisioning, Dynamic Disk Pools
Optional software features	SANtricity Drive Encryption