

Glossary

3rd-Party SCSI Copy Command The 3rd-Party SCSI Copy Command is the specification for the use of the SCSI XCOPY command in order to copy blocks from one storage device to another within a storage network. This command is, for example, the basis for server-free backup.

64b/66b-encoding Variant of 8b/10b-encoding used for 10-Gigabit networks with various cable types.

8b/10b encoding An encoding procedure that converts an 8-bit data byte sequence into a 10-bit transmission word sequence that is optimised for serial transmission. The 8b/10b encoding is used, for example, for Fibre Channel, Gigabit Ethernet and InfiniBand.

Access control The granting or refusal of a request for access to services or resources based upon the identity of the requester.

Access path Descriptively defined as the list of components that are run through by read and write operations to the storage devices and responses to them.

ACL An access control list (ACL) allows the control of access rights in a computer system. The control of the access is usually based on a user name or on groups and grants certain rights, such as permission to read or change a file.

Active An active configuration means a component that is not designed with built-in redundancy.

Active/active An active/active configuration describes a component designed with built-in redundancy, in which both subcomponents are used in normal operation. We differentiate between active/active configurations with and without load sharing.

Active/passive An active/passive configuration describes a component designed with built-in redundancy, in which the second component is not used in normal operation (stand-by).

Agent In the fields of storage networks and system management the client software of a client-server application is very often referred to as the agent. For example, we talk of the backup agent for an application-specific backup client in a network backup system or the SNMP agent for the management of storage devices.

Aggregation The combining of multiple similar and related objects or operations into a single one. Two or more disks can be aggregated into a single virtual disk or in a RAID array.

AIIM The Association for Information and Image Management (AIIM) is an international organisation that is concerned with the subject of enterprise content management (ECM). Its main goal is to merge the interests of users and ECM vendors and to establish standards.

AL_PA The Arbitrated Loop Physical Address (AL_PA) is the address of a device (host bus adapter or switch) in a Fibre Channel Arbitrated Loop.

API An application programming interface (API) is an interface that is provided by a software system so that other programs can use its functions and services. In contrast to an application binary interface (ABI), an API offers the use of the interface at the source code level.

Appliance A device for the execution of a very specific task. Appliances differ from normal computers due to the fact that their software has generally been modified for this very specific purpose.

Application server-free backup Application server-free backup refers to the backup of application data with the aid of an instant copy generated in the disk subsystem and a second server, so that the load for the backup is offloaded from the application server to the second server.

Arbitrated Loop One of the three Fibre Channel topologies. The other two are point-to-point and fabric.

Archive bit The archive bit is a bit in the metadata of a file, which can be used to accelerate the realisation of the incremental-forever strategy.

Archiving Archiving is an intelligent process where data which usually does not change any more is moved into an archiving system. Data access, search, processing and protection must be guaranteed over a long life cycle.

Asymmetric storage virtualisation Asymmetric storage virtualisation is the form of storage virtualisation within a storage network in which the data flow is separated from the control flow. The data flow runs directly between the servers and storage devices whereas the control flow, i.e. the control of the virtualisation by a configuration entity, travels outside the data path.

Autoloader Small automatic tape library with few slots and usually just one drive.

Backup The goal of backup is to generate one or more copies of data that can be used in the event of a failure to restore the original data.

Backup window Time window that is particularly favourable for the backup of the data of an application. For some applications the backup window specifies the maximum period of time that is available for the backup of data.

Bare Metal Restore Alternative expression for 'image restore'.

BC See Business continuity.

Binary encoding Physical encoding procedure for the transmission of data.

Block aggregation The bringing together of physical blocks or block vectors to form logical blocks or block vectors (block-based storage virtualisation). Two or more physical disks can thus be aggregated to form one virtual disk.

Block layer Component of the SNIA Shared Storage Model that includes block-based storage devices and block aggregation.

Block level This expression refers to the physical or virtual blocks of hard disks and tapes. For example, we talk of backup or storage virtualisation on block level.

Block level incremental backup Block level incremental backup describes the capability of a network backup system to incrementally back up only those subsections (blocks) of files or of entire files systems that have changed since the previous backup.

Block orientation Storage devices and I/O protocols that are organised in blocks are called block-oriented, for example hard disks, SCSI, iSCSI and Fibre Channel FCP. File orientation represents an alternative to this.

Bluefin See SMI-S.

Bus Physical I/O medium with several lines for parallel signal transmission.

Business continuity Business continuity (BC) describes technical and organisational strategies for ensuring continuous access to business-critical data and applications even in crisis situations.

Cache Fast storage, in which data accesses to slower storages are buffered.

Cache server Describes a component in a network that temporarily stores data for other components in order to reduce the consumption of network capacity or to provide damping for accesses to slower storage.

Cartridge Physical medium on which storage capacity is available. The storage capacity can be distributed over several sides.

CAS Content Addressable Storage is a storage system that uses a cryptographic hash or CRC that is calculated for an entity of data as an address to reference this entity of data. When data is stored in a CAS, the system calculates the hash and returns it to the application. The application uses the hash to address or reference the data for subsequent access. CAS systems enable efficient data integrity checks in which a cryptographic hash is calculated again for the data and it is compared to the content address. Based on the cryptographic hash, CAS systems also provide data deduplication functions (see Deduplication).

CDP Continuous data protection (CDP) is a backup technique that backs up data immediately at the point when the data changes, capturing every version of data that is generated.

CEE Some vendors use the term Converged Enhanced Ethernet (CEE) to refer to Data Center Bridging (DCB) and FCoE. However, some vendors include additional protocols or features to their definition of CEE than other vendors, thus customers must check carefully as to what are included in a vendor's version of CEE and what are not.

CIFS Common Internet File System (CIFS), the network file system from Microsoft for Windows operating systems.

CIM The Common Information Model (CIM) is an object-oriented description of systems, applications, networks and devices. CIM is a significant component of the Web Based Enterprise Management (WBEM), a standard developed by the Distributed Management Task Force (DMTF) for the management of storage networks, which is currently viewed as the successor to the Simple Network Management Protocol (SNMP).

Class 1, Class 2, Class 3 Different service classes for transmission in a Fibre Channel network.

CLI A command line interface (CLI) is a mechanism for monitoring and controlling hardware and software components by typing commands and receiving text-based results. See also GUI.

Cluster A compound of the same type of resources. The term 'cluster' is often used without being defined more precisely. Sometimes the term 'cluster' also denotes a single node of such a compound. Therefore when talking about clusters you should always ask precisely what is meant by the term.

CMIP The Common Management Information Protocol (CMIP) was designed at the end of the 1980s as the successor to the Simple Network Management Protocol (SNMP). In practice, however, CMIP is hardly ever used.

CN Congestion Notification (CN, IEEE 802.1Qau) is part of the DCB protocol family. It propagates congestion situations from the switches to the end devices to throttle the sending rate of end devices.

CNA An Ethernet card (Network Interface Controller, NIC) which supports conventional TCP as well as the new DCB and FCoE protocols thus being suitable for I/O consolidation is also called Converged Network Adapter (CAN).

Co-location Co-location describes the capability of a network backup system to write (co-locate) several incremental backups of one or more servers onto just a few tapes, so that the number of tape mounts is reduced if the data has to be restored.

Cold backup Cold backup describes the backup of a database that has been shut down for the duration of the backup.

Common Scratch pool Group of cartridges, the storage capacity of which is (or has once again become) completely available and to which all applications have access so that they can reserve a cartridge from it for the purpose in question.

Community name The Simple Network Management Protocol (SNMP) has no secure authentication mechanisms. Instead, so-called community names are used. Two components (for example an SNMP agent and an SNMP-capable management system) can only communicate with each other if they are configured with the same community name.

Consistency group A consistency group combines a number of instant copy or remote mirroring relationships into one unit. Consistency groups increase the consistency of the data that is distributed over several virtual hard disks and are copied using instant copy or remote mirroring.

Copy-on-demand Copy-on-demand is an implementation variant of instant copies and snapshots in which source data is not copied until it has been changed on the source.

CRC A cyclic redundancy code (CRC) represents a checksum that is calculated for an entity of data and enables the data to be checked for integrity. However, the checksum is relatively weak, which means that different data is more likely to produce the same checksum. Thus deliberate manipulation of the data can succeed in producing the same checksum again, which contradicts the data integrity. CRC is normally used to detect corruptions during data transfer. Compared to CRC, cryptographic hashes, such as Message Digest 5 (see MD5) and Secure Hash Algorithms, are more robust in regard to manipulation and are used to check the integrity of data that is stored for longer periods of time.

Credit The credit model is a procedure for the realisation of flow control. Fibre Channel differentiates between buffer-to-buffer credit for link flow control and end-to-end credit for the flow control between two end devices.

Cryptographic Hash A cryptographic hash is a unique checksum that is calculated for an entity of data using cryptographic hash functions such as Message Digest 5 (MD5), Secure Hash Algorithm (SHA) or Whirlpool. A given hash value is most likely unique

for a given set of data and can be used to prove the integrity of the data. However, a low probability exists that two distinct sets of data can result in the same hash value, which is called hash collision. Cryptographic hashes are also used in data deduplication techniques (see Deduplication).

Cut-through routing Cut-through routing is the capability of a switch, a director or a router to forward incoming data packets before they have been fully received.

CWDM Coarse Wavelength Division Multiplexing (CWDM) uses similar procedures to DWDM. The two techniques differ mainly in the division of the frequency ranges and the number of payload streams that they can transmit over a single fibre-optic cable. See also DWDM.

D2D Disk-to-disk (D2D) refers to the class of backup methods in which data is copied from one hard disk or disk subsystem to a second one.

D2D2T Disk-to-disk-to-tape (D2D2T) refers to the class of backup methods in which data is copied from one hard disk or disk subsystem to a second one and from there additionally to tape storage.

DAFS The Direct Access File System (DAFS) is a network file system that is based upon the Virtual Interface Architecture (VIA) and Remote Direct Memory Access (RDMA). DAFS aims to achieve lightweight and very fast file access within a data centre.

DAS Direct Attached Storage (DAS) is storage that is directly connected to a server without a storage network, for example over SCSI or SSA.

Data backup See Backup.

Data copying A variant of data sharing, in which common data is copied for each applications.

Data integrity Data integrity means that the stored data is always the same as its original state. Methods such as cyclic redundancy code (CRC) and cryptographic hashes are used to determine whether violations of the data integrity exist. Error-correcting methods are used to restore the data integrity if possible.

Data migration Data migration describes the transfer of data from a source location to a target location. The meaning of 'location' depends on the context and can be anything from storage media, such as magnetic tapes and hard disks, and file systems to applications such as a Document Management System (DMS). During the data migration the copy of the data at the source location is usually invalidated or deleted, whereas the copy of the data at the target is considered active.

Data scrubbing As a background process, high-end RAID controllers and disk subsystems regularly read all the blocks of a RAID array, thereby detecting defective blocks before they are actively accessed by applications. Data scrubbing significantly reduces the probability of data loss.

Data sharing The use of common data by several applications.

Data shredding Data shredding is a method of physically destroying data on a storage medium without destroying the medium itself. In the process, randomly generated data patterns are usually written over the data that is to be destroyed. This type of data shredding is only possible with rewritable storage media.

DCB The IEEE refers Data Center Bridging (DCB) to a set of protocols which enhances Ethernet to make it a suitable transmission technology for storage traffic.

DCBX Data Center Bridging Exchange (DCBX) refers to an enhancement of the Link Layer Discovery Protocol (LLDP, IEEE 802.1AB-2005) to manage components which support the DCB protocol family. In the second half of 2009 it is still open whether DCBX will become a separate standard or whether it will be integrated into Enhanced Transmission Selection (ETS).

DCE Some vendors use the term Data Center Ethernet (DCE) to refer to Data Center Bridging (DCB) and FCoE. However, some vendors include additional protocols or feature to their definition of DCE than other vendors thus customers must check carefully what is included in a vendor's version of DCE and what not.

Deduplication Data deduplication is a process in which identical data is identified and only one instance of identical data is stored. Other instances of identical data are referenced to the stored instance. This reduces the amount of stored data. The identification process might utilise cryptographic hashes such as MD5 or SHA to identify two identical sets of data. The deduplication process must maintain a list of identity characteristics (such as cryptographic hashes) for the data that has been stored. The process must also maintain a list of references for identical sets of data.

Degraded RAID array A RAID array where a physical hard disk has failed is also referred to as a degraded array. With most RAID procedures all data remains intact in spite of disk loss. However, the performance in a degraded array can sink dramatically because some of the data blocks have to be reconstructed.

DICOM DICOM stands for Digital Imaging and Communications in Medicine. DICOM is a worldwide open standard for the exchange of digital images in medicine. DICOM standardises the format for the storage of image data as well as the communications protocol for the exchange of the images. Almost all suppliers of medical image-making systems support this standard. DICOM is used as a communication protocol by Picture Archiving and Communication Systems for communication with the modalities (see PACS).

Digital signature A digital signature is a private key-encrypted cryptographic hash that is calculated for a given data object. The digital signature for a data object serves two purposes: (1) it is used to prove the integrity of the data object via the cryptographic hash; and (2) it is used to identify the originator of the data object via the private encryption key. A digital signature for a given data object is created by an originator that calculates the cryptographic hash for the object and this hash is encrypted using the originator's

private key. For validation of the integrity of the data object the encrypted hash must be decrypted using the originator's private key and the hash is then compared against a newly calculated hash for the data object. The encryption and decryption of the hash can also be based on asymmetric encryption, whereby the originator encrypts the hash using its private key and an associated public key is used to decrypt the hash. Secure electronic signatures can be generated through the use of digital signatures (also see electronic signatures).

Director A director is a switch with a higher fault-tolerance than that of a simple switch as a result of redundant components.

Disaster recovery Disaster recovery (DR) describes those measures that are needed in order to restore an IT operation in a controlled and predictable manner after an outage.

Discovery Discovery is the automatic detection of all resources (hardware, network topologies, applications) used in a storage network or more generally in a computer network.

Disk subsystem A disk subsystem is a collection of hard disks installed in a common enclosure. We differentiate between JBODs, RAID systems and intelligent disk subsystems. The storage capacity of a disk subsystem is between several Terabytes and 1 Petabyte (2009).

DMAPI The Data Management Application Programming Interface is an interface provided by some file systems to enable applications to intercept file system operations and perform additional file system independent tasks. One example for such an application is hierarchical storage management (HSM). Not all file systems provide the DMAPI interface.

DMI The Desktop Management Interface (DMI) is a protocol for the management of servers specified by the Distributed Management Task Force (DMTF). DMI is seldom used in comparison to the Simple Network Management Protocol (SNMP).

DMS A document management system (DMS) is used to manage archive data. It permits the search, linking and processing of archive data. In the context of this book the term DMS is used as a synonym for the archive management application sometimes also called Enterprise Content Management System (see ECM).

DMTF The Distributed Management Task Force (DMTF) is an association of manufacturers with the objective of driving forward the standardisation of the management of IT systems.

DR see Disaster recovery.

Dual SAN Dual SAN denotes the installation of two storage networks that are completely separate from each other. Dual SANs have the advantage that even in the event of a serious fault in a storage network (configuration error or defective switch, which floods

the storage network with corrupt frames) the connection over the other storage network is maintained.

DWDM Dense Wavelength Division Multiplexing (DMWM) increases the capacity of a fibre-optic cable by assigning several incoming optical signals (= payload streams) to certain optical frequency ranges. Metaphorically speaking, each payload stream is transmitted in a different colour. Since the signals are only optically transformed, there are no limitations with regard to data rates or data formats of the payload streams. As a result, very different payload streams such as Fibre Channel, ESCON, Gigabit Ethernet and Sonet/SDH can be transmitted simultaneously over a single fibre-optic cable.

ECC Error-Correcting Encoding (ECC) specifies coding that enables the correction of errors in a set of data. In addition to the actual data, redundant data is generated and stored together with the data. If an error occurs, the original data can be restored using the redundant data. However, this is only possible to a certain extent, which means that only a limited number of errors can be corrected.

ECM The terminology Enterprise Content Management (ECM) is not used consistently in the literature or by suppliers. Sometimes ECM is used with the same meaning as Document Management System (DMS). On the other hand, the Association for Information and Image Management (AIIM), for example, interprets ECM as technologies for the collection, management, storage, retention and provision of information and documents to support the business processes of a company. With this definition a DMS is part of the ECM solution that manages documents. In this book we concern ourselves with the archiving side of DMS and ECM so that both terms overlap and we therefore make no further distinction between the two.

Electronic signature An electronic signature is interpreted as data that enables the identification of the signer or creator of the data and can be used to check the integrity of the signed electronic data. Digital signatures in conjunction with digital certificates are often used to generate and evaluate electronic signatures. Technically, an electronic signature therefore serves the same purpose as a person's own handwriting on a document.

Electronic vaulting See Vaulting.

Element manager The element manager is a device-specific management interface that is classified as an out-band interface. It is often realised in the form of a GUI or web interface.

Emulated loop Facilitates communication between private loop devices of a Fibre Channel arbitrated loop and devices in a Fibre Channel fabric.

Enhanced shared-nothing cluster Server clusters of up to several ten servers. Enhanced shared-nothing clusters can react to load peaks with a delay.

ENode FCoE end devices (server, storage) are also called FCoE Node or just ENode.

ERP An enterprise resource planning (ERP) system supports a company by planning the maximum use of in-house resources such as capital, equipment and personnel for its business processes.

Error handler Component of a network backup system. The error handler helps to prioritise and filter error messages and to generate reports.

ESCON The Enterprise System Connection (ESCON) is a serial I/O technology for mainframes.

ETS Enhanced Transmission Selection (ETS, IEEE 802.1Qaz) is part of the DCB protocol family. It allows to share the bandwidth of an Ethernet port among priority groups whilst it assures a configurable minimum bandwidth for each priority group.

Exchange An exchange is a logical communication connection between two Fibre Channel devices.

External storage Storage (hard disks, tape drives), which is located outside the computer enclosure.

Fabric The most flexible and scalable of the three Fibre Channel topologies.

Fabric login (FLOGI) Fabric login denotes the registration of an N-Port into a fabric topology. It establishes a session between the N-Port and the corresponding F-Port of a Fibre Channel switch.

FC Abbreviation for Fibre Channel.

FCIA The Fibre Channel Industry Association (FCIA) is an association of manufacturers from the field of Fibre Channel technology.

FCIP Tunnelling protocol that transports the Fibre Channel traffic between two Fibre Channel devices via TCP/IP.

FCN A Fibre Channel Name (FCN) is a 64-bit identifier for a Fibre Channel component, which in contrast to a WWN is not unique worldwide. It has become common practice to refer to WWNs and FCNs simply as WWNs.

FCoE Fibre Channel over Ethernet (FCoE) is a protocol mapping for transferring Fibre Channel frames via Ethernet. Ethernet needs certain enhancements to be suitable for FCoE which are referred as Data Center Bridging (DCB).

FCP The Fibre Channel Protocol (FCP) is the protocol mapping that maps the SCSI protocol onto the Fibre Channel transmission technology.

Fiber Alternative name for fiber-optic cable.

Fibre Channel A technology that can realise both storage networks and data networks. Fibre Channel is currently the predominant technology for the realisation of storage networks. We differentiate between three network topologies: arbitrated loop, fabric and point-to-point.

Fibre Channel SAN A Fibre Channel network that is used as a storage network. Or the other way around: A storage network that is realised with Fibre Channel.

FICON Fibre Connection (FICON) is the mapping of the ESCON protocol on Fibre Channel.

File level The files of a file system are the object of the processing. For example, we talk of backup on file level or storage virtualisation on file level.

File orientation Storage devices and I/O protocols are called file-oriented if they are organised in files or file fragments, for example NAS servers, NFS, CIFS and HTTP. An alternative to this is block orientation.

File/record layer Component of the SNIA Shared Storage Model that maps the database records and files on the block-oriented volumes of the storage device.

FIP The FCoE Initialisation Protocol (FIP) complements FCoE for the discovery and initialisation of FCoE capable devices.

Flow control Mechanism for the regulation of the data stream between a sender and a receiver. The flow control ensures that the transmitter only sends data at a speed that the receiver can process it.

Forward recovery Forward recovery, sometimes also called ‘roll forward’, denotes the restoring of a database using a backup copy plus archive log files generated after the backup copy and the active log files that are still present.

FPMA Fabric Provided MAC Address (FPMA) is a mechanism where the FCoE switch generates a MAC address for the FCoE Line End Point (FCoE LEP) of an FCoE end device (ENode).

Frame The data packets that are transmitted in a Fibre Channel network are called frames.

GDPdU The GDPdU (Grundsätze zum Datenzugriff und der Prüfbarkeit von digitalen Unterlagen) is a set of guidelines for German auditors and tax offices that incorporates standards from the German tax and trade laws. GDPdU essentially specifies how digital data must be archived in order to be compliant with these laws. German companies and companies trading in Germany must follow these standards. Similar to the GDPdU, the U.S. Securities and Exchange Commission (SEC) has released regulations specifying legal requirements for data archiving in the U.S.

GUI A graphical user interface (GUI) is a window-based interface for monitoring and controlling hardware and software components. See also CLI.

HA see High availability.

Hard zoning In hard zoning only the end devices that lie in at least one common zone can communicate with each other. Hard zoning is often confused with port zoning.

HBA A host bus adapter (HBA) is another term for an adapter card that is fitted in a server. Examples of host bus adapters are SCSI controllers, Fibre Channel cards and iSCSI cards.

HCA Host channel adapter (HCA) denotes the connection point of a server to an Infini-Band network.

Hierarchical storage management (HSM) Hierarchical storage management (HSM) denotes the automatic movement of data that has not been used for a long time from fast storage to slower but cheaper storage, for instance from disk to tape. Thereby the movement of the data is transparent to users and applications. HSM is commonly a subfunction of network backup systems.

High availability High availability (HA) refers to the capability of a system to maintain IT operations despite the failure of individual components or subsystems.

HIS A Hospital Information System (HIS) is used to manage data, such as family history data of patients, test results, medical procedures performed and billing information, in hospitals. HIS systems are not normally used to manage picture information because this is usually handled by Picture Archiving and Communication Systems (see PACS).

HL7 Health level 7 (HL7) is an international standard for data exchange between data capturing and processing systems in the healthcare industry. The '7' in the name relates to layer 7 of the ISO/OSI reference model for communication (ISO7498-1) and expresses that the communication is specified at the applications layer. HL7 offers interoperability between hospital information systems (HIS), medical office management systems, laboratory information systems, billing systems for medical services and systems that function as electronic patient files.

Host I/O bus The host I/O bus represents the link between system bus and I/O bus. The most important representative of the host I/O bus is the PCI bus.

Hot backup Hot backup denotes the backup of a database during operation.

Hot spare disks In a RAID configuration (RAID array, intelligent disk subsystem) a spare disk is called a hot spare disk.

HSM See Hierarchical storage management.

Hub A component that is not visible to end devices, which simplifies the physical cabling of a network. In Fibre Channel networks the ring (physical) of the arbitrated loop (logical) is simplified to a star shape (physical).

I/O bus Physical communication connection between servers and storage devices, for example SCSI, Fibre Channel or iSCSI. Originally, parallel buses were used for this such as SCSI or IDE. For historical reasons, serial I/O techniques such as SSA, Fibre Channel or iSCSI are also often called I/O buses.

I/O path The path from CPU and main memory to the storage devices via system bus, host I/O bus and I/O bus.

iECM Interoperable Enterprise Content Management (iECM) is a standard developed by Association for Information and Image Management (AIIM) for the coupling and communication between multiple Document Management System (DMS) and Enterprise Content Management (ECM) systems.

IETF The Internet Engineering Task Force (IETF) is a committee that standardises the protocols for the Internet. These include TCP/IP-based protocols such as FTP, HTTP, NFS, iSCSI, FCIP, iFCP and iSNS.

iFCP Internet FCP (iFCP), a standard with the objective of replacing the network layer in a Fibre Channel SAN with a TCP/IP network.

ILM Information Life cycle Management (ILM) comprises processes, tools and methods that have the aim of ascertaining the value of information and data and adapting the cost to store this information and data to the established value. ILM accounts for the fact that the value of data changes during its lifetime.

Image restore Image restore (also known as Bare Metal Restore) denotes the restoration of a server or a hard disk partition (Windows) or a volume (Unix) from a previously generated copy of a hard disk partition or volume.

In-band management We talk of in-band management if the management of a resource takes place over the same interface over which the actual data is transmitted. Examples of this are the SCSI Enclosure Services (SES) and the corresponding services of the Fibre Channel FCP protocol.

In-band virtualisation Alternative name for ‘symmetric virtualisation’.

Incremental-forever strategy The incremental-forever strategy relates to the capability of a network backup system to calculate the last state of the file system from continuous incremental backups of a file system by means of database operations. A complete backup of the file system is only necessary the first time. After this, only incremental backups are performed. The metadata database in the backup server helps to immediately recreate the last state of the file system when restoring the file system.

InfiniBand New transmission technology that aims to replace the parallel PCI-bus with a serial network. InfiniBand may be used for interprocess communication, client-server communication and server-storage communication.

Instant copy Instant copy is the capability of a storage system to virtually copy large data sets within a few seconds.

Internal storage Storage (hard disks, tape drives) located inside the enclosure of the computer.

IPFC IP over Fibre Channel (IPFC), the protocol mapping that makes it possible to use a Fibre Channel network for IP data traffic.

IP storage General term for storage networks that use TCP/IP as a transmission technique. IP storage includes the protocols iSCSI, FCIP and iFCP.

iSCSI Internet SCSI (iSCSI) is the protocol mapping of SCSI on TCP/IP.

iSCSI SAN A storage network that is realised with iSCSI.

iSER iSCSI Extension for RDMA (iSER) is an application protocol for RDMA over TCP. iSER enables to transmit the SCSI data traffic via the quick and CPU-friendly RDMA over TCP instead of via TCP.

iSNS The Internet Storage Name Service (iSNS) defines a name service that is used by different IP storage standards such as iSCSI and iFCP.

ISL The inter-switch link (ISL) is a connection cable between two Fibre Channel switches.

Java Content Repository The Java Content Repository (JCR) is an interface specification in the programming language Java that specifies the interface and protocol between an archiving application and a Document Management System (DMS). JCR is based on the Java Specification Requests JSR-170 and JSR-283.

JBOD Just a Bunch of Disks (JBOD) is the term for a disk subsystem without a controller.

Jitter As a result of physical influences, incoming signal steps at the receiver are not the same length. This bucking within the signal sequence is called jitter.

Job scheduler Component of a network backup system. It controls which data is backed up when.

Journaling Journaling of a file system describes a method in which the file system – in a similar way to a database – first writes changes to a log file and only then enters them in the actual data area. Journaling significantly reduces the time for a file system check after a system crash.

JSR-170, JSR-283 Java Specification Request 170 (JSR-170) specifies Version 1.0 of the Java Content Repository (JCR) and JSR-283 the extended Version 2.0.

K28.5 symbol Special transmission symbol of the 8b/10b encoding, which does not represent a data byte. The K28.5 symbol includes a special bit sequence that does not occur in a bit sequence generated with 8b/10b encoding even across symbol boundaries. The K28.5 symbols scattered in a data stream allows to synchronise transmitter and receiver.

Label A label is both the sticker on the cartridge, which often has a barcode upon it, and a storage area on the tape that holds metadata.

LAN Local Area Network (LAN), a data network with low geographic extension (maximum several tens of kilometres).

LAN-free backup Backup method of a network backup system in which the backup client copies the data directly to the backup medium via the storage network bypassing the backup server and the LAN.

Latency Latency describes the time duration that passes before the input signal becomes visible in an expected output reaction.

Library partitioning Tape library partitioning statically divides a physical tape library into several logical (=virtual) tape libraries, which are perceived as independent libraries by the connected servers.

Library sharing In tape library sharing several applications dynamically share the tapes and the drives of a tape library.

Link Physical connection cable in a Fibre Channel network.

LIP The loop initialisation primitive sequence (LIP) describes the procedure for the initialisation of a Fibre Channel arbitrated loop. During the LIP procedure the data traffic on the arbitrated loop is interrupted.

Loop Abbreviation for Fibre Channel arbitrated loop.

LUN The SCSI protocol and its derivatives such as Fibre Channel FCP and iSCSI address subcomponents of a device (SCSI target) by means of the Logical Unit Number (LUN). It has become common practice to also call these subcomponents LUN. Examples of LUNs are physical or virtual hard disks exported from a disk subsystem and the tape drives and the media changer of a tape library.

LUN masking LUN masking limits the visibility of disks exported by a disk subsystem. Each computer sees only the disks that are assigned to it. LUN masking thus works as a filter between the disks exported from the disk subsystem and the accessing computers.

LUN zoning Alternative term for LUN masking. Often used in the context of more modern switches that offer zoning on the basis of LUNs and thus facilitate LUN masking in the storage network.

Magneto-optical disk (MOD) A magneto-optical disk (MOD) is a rotating storage medium that is magnetically written and optically read out. An MOD essentially consists of two layers: a reflection layer and a magnetised layer. When the disk is written, the magnetised layer is heated up through an optical laser and a magnetic field provides for the direction of the magnetisation. The direction of the magnetisation has an influence on the polarisation of the light, based on the magneto-optical Kerr effect. When the disk is read, a laser with a low temperature is used and, depending on the direction of the magnetisation, reads out light with varying degrees of polarisation that are interpreted as data bits and bytes.

MAN Metropolitan Area Network (MAN), a data network with average geographic extension (maximum several hundred kilometres).

Managed hub Fibre Channel hub with additional management functions.

Management console Central point, from which all aspects of a storage network, or all aspects of an IT system in general, can be monitored and managed.

Manchester encoding Encoding procedure that generates at least one signal change for every bit transmitted.

MD5 Message Digest Algorithm 5 (MD5) is a widely used cryptographic hash function (see Cryptographic hash) that generates a 128-bit hash value that is used as a checksum. MD5 checksums are used, for example, to test data integrity and authenticity and also for data deduplication (see Deduplication). Secure Hash Algorithm (SHA) is a modern version of cryptographic hash functions and is more robust against collisions because the hash values are longer (up to 2,048 bits in 2009).

Media changer Mechanical transport device that can transport media between slots and drives of a tape library.

Media manager The term “media manager” is used in multiple ways. In network backup systems it refers to the component which manages the hard disks and the tapes upon which the backed up objects (files, file systems, images) are stored. The IEEE 1244 standard refers the server component of a management system for removable media as media manager.

Metadata General data that contains information about actual data is referred to as metadata or meta information. Metadata can therefore be interpreted as additional data that is associated with the actual data. Metadata must also be stored either at the same storage location as the actual data or in a different location such as a database. For example, metadata associated with a file that is stored in a file system includes the name of the file, the size of the file, access rights or other file properties. Usually metadata for files is also stored in the file system. Another example is the full-text index for a file, which also represents metadata and is usually stored in a database to enable efficient searches. Metadata can also contain control information for data – for example, to interpret the data or to manage the life cycle.

Metadata controller (MDC) The metadata controller (MDC) is a management and synchronisation entity in a distributed application. For example, we talk of the metadata controller of a shared disk file system or of the metadata controller of the storage virtualisation.

Metadata database The metadata database is the brain of a network backup system. It includes approximately the following entries for every object backed up: name, computer of origin, date of last change, data of last backup, name of the backup medium, etc.

mFCP Metro FCP (mFCP) is an iFCP variant, which in contrast to iFCP is not based upon TCP but on UDP.

MIB The term management information base (MIB) stems from SNMP jargon. An MIB is a hierarchically constructed collection of variables, which describes the management options of a resource (server, storage device, network component, application).

MIB file File that contains an MIB description.

Microfilm Microfilms are dramatically reduced photographic copies of information that were originally captured on paper. Microfilms are based on a transparent medium and the information on the medium can be made visible through the use of light. The principle involved is similar to that of transparencies and slides.

Mirroring Mirroring of data on two or more hard disks (RAID 1).

Modality A modality refers to any type of device in the medical field that produces images, such as X-rays, ultrasound and computer tomography. Modalities capture the images and usually transfer them to a Picture Archiving and Communication System (see PACS) that stores and manages the data.

Monitoring Monitoring denotes the monitoring of all resources used in the storage network (hardware, network topology, applications).

MTBF Mean Time Between Failure (MTBF) indicates the average period of time between two sequential errors in a particular component or system.

MTTF Mean Time to Failure (MTTF) indicates the average period of time between the recovery of a component or of a system and the occurrence of a new failure.

MTTR Mean Time to Repair (MTTR) indicates the average period of time before a component or a system is restored after a failure.

Multipathing Multipathing is the existence of several I/O paths between server and storage system. The objectives are to increase fault-tolerance by means of redundant I/O paths, to increase the I/O throughput by means of the simultaneous use of several I/O paths, or both at the same time.

Name server In general, the term name server is used to describe an information service in distributed systems. In the case of Fibre Channel the name server (here Simple Name Server) manages information about all N-Ports connected in a fabric such as their WWPN, WWNN, Node_ID and supported service classes and application protocols.

NAS Network Attached Storage (NAS) refers to the product category of preconfigured file servers. NAS servers consist of one or more internal servers, preconfigured disk capacity and usually a stripped-down or special operating system.

NDMP The Network Data Management Protocol (NDMP) defines the interface between the client and the server of a network backup system. The objective of the NDMP is to improve and standardise the integration of NAS servers in a network backup system.

Network Management System (NMS) In SNMP jargon a Network Management System is an application that monitors and manages components by means of the SNMP protocol.

Network File System Network file systems are the natural extension of local file systems: end users and applications can access directories and files over a network file system that physically lie on a different computer – the file server. Examples of network file systems are the Common Internet File System (CIFS), the Network File System (NFS) and the Direct Access File System (DAFS).

Network backup system Network backup systems can back up heterogeneous IT environments incorporating several thousand computers largely automatically.

NFS Network File System (NFS) is the network file system originally developed by SUN Microsystems, which is currently supplied as standard with all Unix systems.

NIC Network Interface Controller (NIC), Network Interface Card (NIC); both terms for network cards.

NRO The Network Recovery Objective (NRO) is an elementary parameter for business continuity. It indicates the maximum allowed time after an outage to restart the network (LAN and WAN).

Off-site location An off-site location is a remote location at which a second copy of data that has been backed up by means of a network backup system is stored. The second copy of the data in the off-site location serves to protect against major catastrophes.

OOM Object-oriented modelling (OOM) is an object-oriented specification language, which is used for the description of the Common Information Model (CIM).

Open Systems Open Systems signifies the world of the non-mainframe server. Unix, Windows, OS/400, Novell and MacOS belong to the Open System world. Incidentally, for us 'Unix' also covers the Linux operating system, which is sometimes listed separately in such itemisations.

Ordered set 8b/10b encoded group of four transmission words that begins with the K28.5 symbol.

Out-band management Out-of-band management (out-band management for short) signifies the management of a resource by means of a second interface, which exists in addition to the data path. An example of out-band management would be the management of a Fibre Channel switch by means of an Ethernet connection and SNMP.

Out-band virtualisation Alternative term for 'asymmetric virtualisation'.

PACS A Picture Archiving and Communications System (PACS) manages and stores digital images that are captured at modalities (see Modality) of imaging medicine.

Parity Parity is a binary cross-check sum or check sum. RAID 4 and RAID 5, for example calculate and store additional parity blocks, with which the data stored upon a hard disk can be reconstructed after its failure.

Partition Part of a side, which provides storage capacity as a physical unit of the cartridge.

PCI Peripheral Component Interconnect (PCI) is currently the predominant technology for host I/O buses.

PFC Priority Flow Control (PFC, IEEE 802.1Qau) is part of the DCB protocol family. It refines the concepts of the Ethernet PAUSE frame to priority groups thus each priority group can be suspended individually without disabling the whole port.

Point-in-time restore Point-in-time restore signifies the capability of a network backup system to recreate any desired earlier state of a file system.

Point-to-point The simplest of the three Fibre Channel topologies, which solely connects two end devices (server, storage) together.

Port A port denotes the physical interface of a device (servers, storage devices, switches, hubs, etc.) to a storage network.

Port login (PLOGI) Port login denotes the establishing of a connection (session) between two Fibre Channel end devices. Port login exchanges service parameters such as service class and end-to-end credit. It is an absolute prerequisite for further data exchange.

Port zoning Zoning variant, in which the zones are defined by means of port addresses. Port zoning is often confused with hard zoning.

Prefetch hit rate The prefetch hit rate describes the success rate of a cache in shifting data from a slower storage device before a different component demands precisely this data from the cache.

Private loop A Fibre Channel arbitrated loop that is not connected to a fabric.

Private loop devices A private loop device is a device connected to a Fibre Channel arbitrated loop that does not master the fabric protocol. It is not capable of communicating with end devices in the fabric via a Fibre Channel switch connected to the loop.

Protocol converter A protocol converter connects two incompatible interfaces and translates between them.

Protocol mapping The Fibre Channel standard denotes the mapping of an application protocol such as SCSI or IP on the Fibre Channel transport layer (FC-2, FC-3) as protocol mapping.

Process login (PRLI) Process login describes the establishing of a connection (session) between two processes on the FC-4 layer of Fibre Channel.

Public loop A Fibre Channel arbitrated loop, which is connected to a fabric via a switch.

Public loop devices A public loop device denotes a device connected to a Fibre Channel arbitrated loop, which in addition to the loop protocol also masters the fabric protocol. It can communicate with end devices in the fabric via a Fibre Channel switch connected to the loop.

Quickloop Implementation variant of the emulated loop by the company Brocade.

RAID Originally RAID was the abbreviation for ‘Redundant Array of Inexpensive Disks’. Today RAID stands for ‘Redundant Array of Independent Disks’. RAID has two primary objectives: to increase the performance of hard disks by striping and to increase the fault-tolerance of hard disks by redundancy.

RDMA Remote Direct Memory Access (RDMA) enables processes to read and write the memory areas (RAM) of processes that are running on another computer. RDMA is based on VI. RDMA is aimed at the lightweight and very fast interprocess communication within a data centre.

RDMA over TCP Standardised RDMA variant that uses TCP as the transmission medium.

Real time data sharing Variant of data sharing in which several applications work on the same data set concurrently.

Regulatory compliance When archiving is subject to legal conditions and requirements, reference is made to regulatory-compliant archiving. In practice, the terminology ‘regulatory compliance’ is used with different meanings. We interpret ‘regulatory compliance’ as the obligation to maintain legal requirements governing the protection of data from deletion and manipulation.

Remote mirroring Remote mirroring signifies the capability of a block-based storage system (e.g. a disk subsystem) to copy data sets to a second storage system without the involvement of a server.

Replication Replication denotes automatic copying and synchronisation mechanisms on file level.

RIS A radiological information system (RIS) is a data processing system that manages information from the radiology area. This information is of an administrative and medical nature and encompasses such things as patients’ medical history, the scheduling of radiology equipment, test results and the data used for invoicing. In terms of its function, an RIS is closely related to hospital information systems.

RNIC RDMA enabled NIC (network interface controller), a network card that supports RDMA over TCP and, in addition to RDMA, most likely also realise the functions of a TCP/IP offload engine (TOE).

Roll forward See Forward recovery.

Rolling disaster With a rolling disaster, parts of an IT infrastructure fail gradually but not the entire data centre at the same time. This produces an inconsistency in the available data and IT services because some applications may still be running while others have already ceased operation.

RPO Recovery Point Objective (RPO) is an elementary parameter for business continuity. It indicates the maximum amount of data loss tolerated in a crisis.

RSCN The Registered State Change Notification (RSCN) is an in-band mechanism in Fibre Channel networks, by means of which registered end devices are automatically informed of status changes of network components and other end devices.

RTO Recovery Time Objective (RTO) is an elementary parameter for business continuity. It indicates the maximum length of time that is tolerated for the restart of IT operations after an outage.

SAFS SAN Attached File System (SAFS), an alternative term for shared disk file system.

SAN SAN is an abbreviation for two different terms. Firstly, SAN is the abbreviation for 'Storage Area Network'. Very often 'storage area networks' or 'SANs' are equated with Fibre Channel technology. The advantages of storage area networks can, however, also be achieved with alternative technologies such as for example iSCSI. In this book we therefore do not use the term SAN or 'Storage Area Network' alone. For general statements on storage area networks we use the term 'storage network'. Otherwise, we always state the transmission technology with which a storage area network is realised, for example Fibre Channel SAN or iSCSI SAN.

Secondly, SAN is an abbreviation for 'System Area Network'. A system area network is a network with a high bandwidth and low latency, which serves as a connection between computers in a distributed computer system. In this book we have never used the abbreviation SAN to mean this. However, it should be noted that the VIA standard uses the abbreviation SAN in this second sense.

SAN router Alternative name for a Fibre Channel-to-SCSI bridge.

SAS Serial Attached SCSI is an I/O technique that links individual hard disks and tape drives that sequentially transmit the conventional parallel SCSI protocol and, therefore, achieve higher transmission rates than SCSI.

SATA Serial ATA (SATA) is an economical I/O technology for disk attachment that transmits the conventional parallel ATA protocol serially and thus permits higher transmission rates than IDE/ATA.

Scratch pool Group of cartridges, the storage capacity of which is (or has once again become) completely available.

Scratch tape A new tape without content or a tape the content of which is no longer of interest and the whole of the storage capacity of which can be used for new purposes.

SCSI The Small Computer System Interface (SCSI) is an important technology for I/O buses. The parallel SCSI cables are increasingly being replaced by serial I/O techniques such as Fibre Channel, TCP/IP/Ethernet, SATA, SAS and InfiniBand. The SCSI protocol, however, lives on in the new serial techniques, for example as Fibre Channel FCP or as iSCSI.

SDK A software development kit (SDK) is a collection of programs and documentation for a specific software. SDKs are designed to simplify the development of applications that are run on the software concerned. Sometimes the use of an SDK is essential to the development of such applications.

SDP The Socket Direct Protocol (SDP) maps the socket API of TCP/IP on RDMA, so that protocols based upon TCP/IP such as NFS and CIFS do not need to be modified. Users of SDP benefit both from the simplicity of the protocol and also from the low latency and low CPU load obtained with RDMA.

SEC The U.S. Securities and Exchange Commission (SEC) is an independent agency of the United States government that enforces the federal securities laws and regulates the securities industry/stock market. As such, the SEC Regulations 17 CFR 240.17a-3 and 17 CFR 240.17a-4 stipulate the records retention requirements for the securities broker-dealer industry.

SES The SCSI Enclosure Services (SES) are an in-band management interface for SCSI devices.

Sequence A sequence is a large data unit in the FC-2 layer of Fibre Channel that is transmitted from transmitter to receiver in the form of one or more frames.

Server-centric IT architecture In a server-centric IT architecture, storage devices are only connected to individual servers. Storage only ever exists in relation to the servers to which it is connected. Other servers cannot directly access the data; they must always go through the server to which the storage is connected.

Server consolidation Server consolidation is the replacement of many small servers by a more powerful large server.

Server-free backup Backup method of a network backup system, in which the data is copied from the source disk to the backup medium via the storage network without a server being connected in between. Server-free backup makes use of the 3rd-Party SCSI Copy Command.

Services subsystem Component of the SNIA Shared Storage Model in which the management tasks of a shared storage environment are brought together.

SHA A Secure Hash Algorithm (SHA) refers to a group of cryptographic hash functions (see Cryptographic hash).

Shared disk file system Shared disk file systems are a further development of local file systems in which several computers can directly access the hard disks of the file system

at the same time via the storage network. Shared disk file systems must synchronise the write accesses to shared disks in addition to the functions of local file systems.

Shared-everything cluster The shared-everything cluster is the cluster configuration that permits the greatest flexibility and the best load balancing. In shared-everything clusters, several instances of an application run on different computers, with all instances providing the same services towards the outside. A corresponding load balancing software ensures that all instances are loaded to the same degree.

Shared-nothing cluster Shared-nothing clusters are a configuration of two servers in which in the event of the failure of one computer the remaining computer takes over the tasks of the failed computer in addition to its own.

Shared-null configuration The shared-null configuration is a server or an application that is not designed with built-in redundancy. If the server fails the application is no longer available.

Shared storage environment SNIA term for storage-centered IT architectures.

Side Part of a cartridge that provides storage capacity. A side contains one or more partitions. Tapes normally possess only one side. DVDs and magneto-optical media are also available in double-sided variants. Holographic storage may provide even more than two sides.

Single point of failure Single point of failure signifies a subcomponent of a system, the failure of which leads to the failure of the entire system. Fault-tolerant systems such as server clusters or high-end disk subsystems must not have any single points of failure.

Skew Skew means the divergence of signals that belong together in a parallel bus.

SLA Service Level Agreements (SLAs) describe in detail and in a measurable form what an IT service customer requires of an IT service provider. These IT service agreements are usually part of an official contract between the IT provider and the customer.

Slot Storage location for cartridges that are not being accessed.

SMI-S The Storage Management Initiative Specification (SMI-S) is a further development of WBEM and CIM by SNIA, which is specially tailored to the management of storage networks. Amongst other things, the standardised refinement of the CIM classes aims to guarantee the interoperability of management systems for storage networks.

Snapshot A snapshot means an instant copy within a file system or a volume manager.

SNIA Storage Networking Industry Association (SNIA), an association of manufacturers in the field of storage and storage networks.

SNMP The Simple Network Management Protocol (SNMP) is a standard that was originally developed for the management of IP networks. SNMP is now a widespread standard for the management of IT systems that is also used for the management of storage networks.

Soft zoning Soft zoning describes a zoning variant that restricts itself to the information of the name server. If an end device asks the name server for further end devices in the Fibre Channel network then it is only informed of the end devices with which it lies in at least one common zone. However, if an end device knows the address of a different device, with which it does not lie in a common zone, then it can nevertheless communicate with the other device. Soft zoning is often confused with WWN zoning.

SoIP Storage over IP (SoIP), the name of a product from the former corporation Nishan Technologies. According to the manufacturer these products are compatible with various IP storage standards.

SPB Shortest Path Bridging (SPB, IEEE 801.1aq) is an approach to define routing mechanism for unicast and for multicast frames which supports redundant links and parallel VLAN configurations.

SPMA Server Provided MAC Address (SPMA) is a mechanism where the management software of a server generates a MAC address for the FCoE Line End Point (FCoE LEP) of an FCoE end device (ENode).

SRM Storage Resource Management (SRM) is the category of software products that unifies storage virtualisation and storage management.

SRP The SCSI RDMA Protocol (SRP) allows the transfer of SCSI via RDMA bypassing TCP or any other intermediate protocol.

SSA Serial Storage Architecture, an alternative I/O technology to SCSI.

SSH Secure Shell (SSH) is a network protocol that, in contrast to TELNET, enables a user to register with a remote computer over an encrypted network connection and to execute programs there.

SSP A Storage Service Provider (SSP) is a business model in which a service provider (the SSP) operates a storage network, which is used by many customers. Originally it was hoped that this would result in cost benefits. In practice this business model has failed. However, it is very likely that this approach will experience a renaissance in a modified form with the increasing use of the web architecture and so-called cloud computing.

Storage-centric IT architecture In contrast to server-centric IT architecture, in storage-centric IT architecture, storage exists completely independently of any computers. A storage network installed between the servers and the storage devices allows several servers to directly access the same storage device without a different server necessarily being involved.

Storage consolidation Storage consolidation means the replacement of a large number of small storage systems by one more powerful large storage system.

Storage gateway Alternative term for a Fibre Channel-to-SCSI bridge.

Storage Hierarchie Storage Hierarchie (Tiered Storage) denotes a storage architecture which comprises multiple storage types such as disk and tape or disks with varying performance characteristics. Storage virtualisation and applications such as network backup and archiving systems can leverage a storage hierarchie and move the data between the tiers to store the data cost effective.

Storage networks The idea behind storage networks is to replace the SCSI cable between servers and storage devices by a network, which is installed alongside the existing LAN as an additional network and is primarily used for the data exchange between computers and storage devices.

Storage virtualisation Storage virtualisation (often just called virtualisation) is generally used to mean the separation of storage into the physical implementation in the form of storage devices and the logical representation of the storage for the use by operating systems, applications and users. A differentiation is made between three levels of storage virtualisation: (1) virtualisation within a storage system, for example in a RAID disk subsystem or an intelligent disk subsystem, (2) virtualisation in the form of an own virtualisation entity in the storage network and (3) virtualisation on the server by host bus adapter, volume manager, file systems and databases. A further differentiation is made with regard to the granularity of the virtualisation (virtualisation on block level and virtualisation on file level) and, for the virtualisation in the storage network, we also differentiate between symmetric and asymmetric virtualisation.

Streaming The reading or writing of large quantities of data to a tape, whereby the data is written in one go without stopping, rewinding and restarting the tape.

Striping Distribution of data over two or more hard disks (RAID 0).

Support matrix In heterogeneous storage networks, numerous components from extremely different manufacturers come together. In the support matrix, manufacturers of hardware and software components state which components from other manufacturers their components will work with.

Switch The switch is the control centre in networks such as Ethernet and the Fibre Channel fabric. It realises the routing of frames and services such as name server and zoning.

Switched hub A special kind of Managed hub, which in addition allow for the direct communication between two end devices, so that several end devices can communicate with each other in pairs within a Fibre Channel arbitrated loop at the same time.

Symmetric storage virtualisation Symmetric storage virtualisation is the form of storage virtualisation within a storage network in which the data flow between servers and storage devices plus the control flow – i.e. the control of the virtualisation by a virtualisation instance – take place in the data path.

System bus The I/O bus in a computer that connects, amongst other things, the CPUs to the main memory (RAM).

Tape library partitioning Tape library partitioning (library partitioning for short) divides a physical tape library statically into several logical (= virtual) tape libraries, which are perceived as independent libraries by the connected servers.

Tape library sharing In tape library sharing (library sharing) several applications dynamically share the tapes and drives of a tape library.

Tape mount The inserting of a tape in a tape drive.

Tape reclamation In a network backup system over time more and more data is left on a tape that is no longer required. With current technology it is difficult to write new data to these gaps on tapes that have become free. In tape reclamation the data that is still valid from several such tapes with gaps is copied onto one new tape so that these tapes can be rewritten.

Target The SCSI protocol calls the device connected to an SCSI bus a target. Examples of targets are servers, disk subsystems and tape libraries.

Target_ID Target_ID is the name for the address of a device (target), which is connected to an SCSI bus.

TCA InfiniBand calls the connection point of a server to an InfiniBand network a Target Channel Adapter (TCA). The complexity of a TCA is low in comparison to an Host Channel Adapter (HCA).

TCP/IP offload engine (TOE) A network card that realises the TCP/IP protocol stack completely in firmware on the network card. TOEs significantly reduce the CPU load for TCP/IP data traffic.

TELNET TELNET is a network protocol that, in contrast to SSH, enables a user to register with a remote computer over a non-encrypted network connection and to execute programs there. All data – including the passwords for setting up the connection – are transmitted in clear text.

Three-tier architecture Further development of the client-server architecture, in which the data, applications and the user interface are separated into different layers.

Tiered Storage see Storage hierarchy.

Translated Loop Implementation variant of the emulated loop from CNT/Inrange.

Trap A trap is a mechanism with which a resource managed by SNMP (or to be more precise its SNMP agent) informs a management system for storage networks or a general management system of state changes.

Trap recipient The trap recipient is the recipient of SNMP messages (traps). To set up a trap recipient the IP address of the computer that is to receive the trap is entered on the SNMP agent.

TRILL The Transparent Interconnection of Lots of Links (TRILL) is an approach for enabling Ethernet to operate multiple parallel links in parallel which is needed to support the multipathing concept of storage networks.

Twin-tailed SCSI cabling Cabling method in which the storage devices are connected to two servers via a SCSI bus for the benefit of fault-tolerance.

UDO Ultra Density Optical (UDO) was developed by the British supplier Plasmon and is considered the successor to MODs (see Magneto-optical disk) for professional company use. Technically, UDO is based on phase-change technology and the current UDO Version 2 is marketed as having 60 GBytes and was presented at CeBit 2007. Since September 2004 UDO has been based on the cross-vendor standard ISO/IEC 17345.

ULP Upper level protocol (ULP). Application protocol of a Fibre Channel network. Examples of ULPs are SCSI and IP.

Unmanaged hub Fibre Channel hub without management functions.

Vaulting Vaulting is the transfer of data to a secure site. With conventional vaulting, the data is first backed up on tape and the tapes are then moved to a remote location. With more recent techniques the data is copied directly over WAN connections to a different data centre (electronic vaulting).

VI The Virtual Interface (VI) denotes a communication connection in the Virtual Interface Architecture (VIA).

VIA The Virtual Interface Architecture (VIA) is a system-level I/O technology, which facilitates the lightweight and fast data exchange between two processes that run on different servers or storage devices within a data centre.

VI NIC VI-capable network card. Today VI-capable network cards exist for Fibre Channel, Ethernet and InfiniBand.

Virtualisation See Storage virtualisation.

Voice over IP (VoIP) VoIP is the transmission of telephone calls via IP data networks.

Volume A volume is a logical data container. It serves to reserve storage capacity on storage devices for applications.

Volume level Backup mode in which an entire volume (e.g. disk, partition of a disk or logical volume) is backed up as a single object.

Volume manager Virtualisation layer in the server between disk and file system that can bring together several physical hard disks to form one or more logical hard disks.

VPN A virtual private network (VPN) enables secure transmission over an insecure network. The users of a VPN can exchange data the same way as in an internal LAN. The connection over the insecure network is usually encrypted. The authentication of

VPN users can be guaranteed through the use of passwords, public keys and digital certificates.

VSAN A virtual SAN (VSAN) makes it possible to operate several virtual Fibre Channel fabrics that are logically separate from one another over one physical Fibre Channel network. In addition, separate fabric services such as name server and zoning are realised for every virtual storage network.

WAN Wide Area Network (WAN), a data network with large geographical extension (several thousand kilometres).

WBEM The Web Based Enterprise Management (WBEM) is a standard developed by the Distributed Management Task Force (DMTF) for IT infrastructure management, which is currently viewed as the successor to the Simple Network Management Protocol (SNMP). WBEM uses web techniques. A significant part of WBEM is the Common Information Model (CIM).

Web architecture Further development of the three-tier architecture to a five-tier architecture for the flexible support of Internet and e-business applications. The representation layer is broken down into the web server and the web browser and the data layer is broken down into the organisation of the data (databases, file servers) and storage capacity for data (disk subsystems and tape libraries).

WebGUI A web graphical user interface (WebGUI) is a GUI implemented using web techniques.

WORM Write once read many (WORM) means that data can only be written once and then only read. It is not possible to change or overwrite the data that complies with regulatory requirements (see Regulatory compliance). Traditionally, storage media such as CDs and DVDs – where data bits are burned – was used as WORM Media. In the newer implementations rewritable media (such as hard disk and tapes) is also being used, whereby additional control software provides WORM protection.

Write order consistency With asynchronous remote mirroring, write order consistency ensures that data that is distributed over multiple virtual hard disks is updated on the destination disks in the same sequence as on the primary disks despite the asynchronous remote mirroring.

WWN A World Wide Name (WWN) is a 64-bit identifier for a Fibre Channel component, which in contrast to FCN is unique worldwide. In practice it has become common practice to call WWNs and FCNs simply WWNs.

WWN zoning Zoning variant in which the zones are defined by WWNs. WWN zoning is often confused with soft zoning.

WWNN The World Wide Node Name (WWNN) is the WWN for a device (server, storage device, switch, director) in a Fibre Channel network.

WWPN The World Wide Port Name (WWPN) is the WWN for a connection port of a device (server, storage device, switch, director) in a Fibre Channel network.

XAM The Extensible Access Method (XAM) is an interface standard that specifies the communication between a document management system (DMS) and archive storage. XAM was developed under the umbrella of the Storage Networking and Industry Association (SNIA) with the participation of multiple storage vendors. SNIA approved the first version of the XAM architecture and specification in July 2008.

XCOPY SCSI command that realises the 3rd-Party SCSI Copy Command.

Zoning Subdivision of a network into virtual subnetworks, which can overlap.

