

Glossary of Network Storage Terms

July 2010

Version 2.0



Copyright 2010, Coraid, Inc. All Rights Reserved.

Except as specifically permitted herein, no portion of this document may be reproduced in any form or by any means without the express written consent of Coraid.

The trademarks, logos, and service marks (collectively "Trademarks") appearing on the Coraid website are the property of Coraid and other parties. ALL OF THE TRADEMARKS MENTIONED IN THIS MANUAL ARE THE PROPERTY OF THEIR RESPECTIVE OWNERS. Nothing contained in this document should be construed as granting any license or right to use any Trademark without the prior written permission of the party that owns the Trademark. Coraid and EtherDrive are registered trademarks of Coraid. Coraid trademarks include RAIDShield and VirtualStorage. Windows is a registered trademark of Microsoft Corporation in the United States and other countries. InstallShield is a registered trademark of InstallShield Corporation.

Phone

+1-650-517-9300

+1-877-548-7200

Fax

+1-650-226-3788

Coraid Incorporated
255 Shoreline Drive, Suite 650
Redwood City, California 94065
United States of America

www.coraid.com

AoE

ATA-over-Ethernet, published in 2003, an open SAN storage protocol that describes the encapsulation of block storage commands (ATA or SCSI) within an 802.3 Ethernet frame. AoE frames are registered by the IEEE as Ethernet type 0x88A2.

AoE Target (LUN)

AoE storage targets, also referred to as LUNs, are block storage devices accessed via Ethernet SAN using the AoE protocol. Each LUN is uniquely identified within the Ethernet SAN by two numbers (x=major and y=minor). Valid LUN numbers range for x = 0 to 65,535, and y = 0 to 255 (ie. x.y ranges from LUN 0.0 thru LUN 65535.255).

When using EtherDrive storage “x” is either the appliance shelf number where the “physical” LUN is located, or the “virtual” LUNs major number created by one of Coraid’s SAN virtualization appliances (ie. EtherDrive EMX or EtherDrive VSX). “y” is either the LUNs minor number assigned to a RAID group within an EtherDrive storage appliance, or the “virtual” LUN minor number created by one of Coraid’s SAN virtualization appliances (ie. EtherDrive EMX or EtherDrive VSX).

In this context, “physical” ties the LUN to a specific shelf location, and “virtual” abstracts the LUN from a shelf location.

Appliance

A ready-to-use single-purpose device. In the context of Coraid products a RAID storage appliance as opposed to a general purpose server running any user application.

ATA

Advanced Technology Attachment. ATA is a standard for connecting hard disks to a host computer. It is maintained by the X3/INCITS committee T13 (<http://www.t13.org>).

Block

A unit of data stored on a hard drive, each block size is commonly 512 bytes. Blocks are also known as sectors, which are sub-units of tracks. Block locations are uniquely addressed using a logical block address (lba) number. When a disk is access it reads or writes the entire 512 byte block of data (4096 bits).

Block Device

A storage device, such as a disk, that stores data in blocks.

Block Storage

Storage devices, such as disks, that store in blocks.

CIFS

Common Internet File System, a file-sharing protocol formerly known as SMB (Server Message Block; see also Samba). CIFS is used in Microsoft networking environments to share files over a TCP/IP network. It is the basis of Windows File Sharing.

CLN

Coraid Linux NAS, an optimized Linux file server that acts as an NAS gateway between AoE block storage and an NFS and/or CIFS file sharing network.

Cluster Filesystem

A filesystem that allows multiple computers to share files by using the same block storage device at the same time, without relying on a single server to own and manage the files. This requires the filesystem to control block address locking while data is being updated, and then communication with other cluster members to flush cache and read the new data.

Direct Attached Storage (DAS)

A storage device that is physically connected to only one computer, not to a network. Direct attached storage connects to one host using an interface such as PATA, SATA, SCSI, USB, or Firewire.

Disk Server

A computer that performs disk sharing. Contrast file server.

Disk Sharing

The use of networking to make a block storage device (normally a disk drive) directly available to one or more computers, which can format the disk and use it as if it were directly attached. Compare SAN. Contrast file sharing.

Driver

A software component that allows an operating system to access a particular piece of hardware. The AoE protocol driver bridges the operating systems block storage interface to an Ethernet network port.

EtherDrive Storage Blade

Early Coraid terminology for a network storage device consisting of a singledisk mounted to a circuit card which converts the disk interface to Ethernet. Multiple Blades slide into a rack mounted shelf, with one Blade per slot location and one Ethernet connection per Blade. Contrast lblade.

Fabric

Network infrastructure, connections, etc. See switch fabric.

FC

Fibre Channel.

Fibre Channel

A SCSI-based communication protocol developed for connecting computers to disk drives via fiber optics; versions of Fibre Channel have been adapted to Ethernet. The Fibre Channel standard is maintained by INCITS committee T11 (<http://www.t11.org>).

Filer

A file server.

File Server

A computer that performs file sharing. Contrast disk server.

File Sharing

The use of networking to make files available to computers other than the one that owns and manages them. Compare NAS. Contrast disk sharing.

File Storage

File sharing; storage of files identified by file names and managed by a server.

Filesystem

An operating system software component that manages files, mapping them to particular block storage locations.

Gateway

A device that provides protocol translation. For example, an NAS gateway translates the NFS or SMB/CIFS file sharing protocol into a disk sharing protocol such as AoE, Fibre Channel, or iSCSI; it does this by acting as a file server.

GFS

Global File System, a cluster filesystem for Linux. See cluster filesystem.

HBA

Host bus adapter; an Ethernet adapter (or other network adapter) in a computer; a NIC.

Heartbeat

A signal periodically sent out by a hardware component in order to inform another component that it is working normally. See STONITH.

Host

Any computer on a network, whether server, client PC, or virtual machine. In the context of networked storage, the host system “initiates” access to storage “targets”.

HVAC

Abbreviation for Heating, Ventilation and Cooling.

INCITS

International Committee for Information Technology Standards (<http://www.incits.org>).

Initiator

A host that requests access to storage devices; originally a SCSI-specific term but now extended to other storage protocols.

iSCSI

Internet Small Computer Systems Interface. iSCSI is a storage protocol standard that encapsulates SCSI commands into TCP/IP data packets for disk sharing over an IP network. iSCSI was originally specified in RFC 3720. Compare AoE; Fibre Channel.

JBOD

Just a Bunch of Disks. This term describes disk sharing when multiple disk drives are accessed separately, not grouped together in RAID volumes.

Journaling Filesystem

A filesystem that does not perform disk writes immediately, but instead saves them in a “journal” of work to be done, and performs the changes later, more safely and efficiently, in large blocks. The server uses the journal to present the data to the hosts as if the changes had already been committed to disk. (In accounting, a journal is a logbook of transactions that have not yet been posted to the main set of records.)

Jumbo Frame

An Ethernet data packet with more than 1500 bytes of content. Jumbo frames are supported by some Gigabit Ethernet (1000base-T) switches and network interface cards. Jumbo frames typically carry 9000 bytes of content.

Latency

The time required to store or retrieve data from a storage device (disk), measured from making the request to completing the operation.

Lblade

Lblade stand for “logical” blade. Early Coraid terminology for a RAID array constructed from “physical” blade members, now called an LUN, which may or may not correspond to a physical disk. Contrast blade.

Logical Unit Number (LUN)

The number that identifies a disk drive or other block device in a system; by extension, the disk itself as visible outside the server, whether it corresponds to a physical disk, a RAID group, or a virtual disk.

In the context of EtherDrive storage, a LUN is a grouping of contiguous numbered block addresses each address containing up to 512 bytes of stored data. LUNs can be disk drives, disk partitions, RAID arrays, or virtualized storage devices. The terms “LUN”, “hard drive”, “disk”, “volume” and “virtual LUN” can be used interchangeably.

Logical Volume Manager (LVM)

A software or hardware component that manages virtual disk partitions and presents them to the host computer as if they were physical sections of the disk. Virtual partitions can be changed without moving their actual contents. Virtual disks are LUNs created from groups of disk pooled into volume groups.

Metadata

Data describing other data. In the context of a filesystem, file metadata includes the locations of the blocks that hold the files. In a RAID group, metadata indicates the structure of the group.

MPIO, Multipath Input/Output

The use of more than one Ethernet connection (through separate Ethernet adapters) to connect a single computer or other device to a network. Using multiple paths can increase throughput and provide path redundancy.

NAS

Network Attached Storage, the sharing of files over a network by means of a file sharing protocol such as NFS or CIFS. See file sharing. Contrast SAN.

NAS Gateway

A file server that connects a file sharing network (LAN) to a separate network used to share disk drives (SAN).

Networked Storage

Any type of data storage that is accessed through a network.

NFS

Network File System, a file sharing protocol developed at Sun Microsystems in 1985.

NIC

Network interface card (Ethernet card, HBA).

Partition

A logical division of a block storage device. Typically, a newly installed disk drive is divided into one or more partitions, and then file systems are installed in the partitions. The locations of the partitions are recorded on the drive in its partition table.

PATA

Parallel ATA, the older hardware implementation of the ATA interface for connecting hard drives to computer motherboards. Contrast SATA.

Physical LUN

A fixed size LUN. LUN size is determined by the storage capacity of the disk (or disk RAID array). Contrast Virtual LUN

RAID

Redundant Array of Inexpensive Disks, a set of methods for grouping multiple disk drives for better performance or fault tolerance.

RAID Group

A collection of disk drives arranged into a RAID system. A RAID group becomes a volume once it is mounted with a filesystem.

RTO

Recovery Time Objective, a disaster recovery term indicating the maximum acceptable time to be taken to recover from a disruption.

Samba

A free, open source implementation of the CIFS (SMB) protocol that allows Linux hosts to share files and printers with a Microsoft Windows network.

SAN

Storage Area Network. Strictly speaking, SAN refers to a separate network set up for disk sharing, but the term often denotes disk sharing itself, in contrast to NAS.

SATA

Serial ATA, a newer version of the ATA interface for connecting disk drives to computer motherboards.

SCSI

Small Computer Systems Interface (pronounced “scuzzy”), a standard protocol for connecting disk drives and other devices to computers. The SCSI standard is maintained by INCITS committee T10 (<http://www.t10.org>).

Sector

A sub-unit of a track on a hard disk drive, containing a single block of data (commonly 512 bytes).

Seek-Bound, Seek-Limited

Handicapped by the time needed to move the disk heads from one location to another. This describes a disk drive whose performance is limited by inefficient storage allocation that scatters files all over the disk, or by an application that inherently needs constant access to widely scattered data.

SMB

Server Message Block, an older name for the CIFS protocol. See CIFS.

Spindle Count

The total number of physical disks (rotating spindles) in a system.

STONITH

“Shoot the other node in the head,” a somewhat comical way to describe a high-reliability system where the standby server constantly checks for the “heartbeat” (periodically repeated OK signal) of the main server and automatically powers it off and takes over if the heartbeat goes away.

Switch Fabric, Switched Fabric

The connections that link up the components of a network. In Fibre Channel, this is an elaborate mesh of switches and fiber-optic cables; in Ethernet it is a cable carrying data from any unit to any other unit.

Target

A term borrowed from SCSI that indicates the endpoint of communication, usually a disk drive or a RAID group.

Track

A circular region of a disk drive in which several sectors of data are stored. The disk rotates, bringing all those sectors under the read-write head while it stands still. To access other tracks, the head moves.

Virtual LUN

A variable size LUN. LUN size is determined by the storage logical volume manager, using extents (chunks or block addresses) allocated from a volume group (disk pool). Contrast Physical LUN

Virtualization

The replacement of any component of a computer system by something else that performs the same function but is physically not what it appears to be. For example, CPU virtualization can make one computer function as several separate computers; disk virtualization can make a group of disk drives function as one large disk drive or as a group of drives different from what is really there.

VLAN

Virtual Local-Area Network, a group of computers that communicate through Ethernet as if they were attached to a single cable, even though the actual interconnection between them is more complex.

Volume

A single block device as seen by the computer using it. It can be a disk partition, a single disk, or a group of disks assembled into a RAID volume.

XFS

A journaling filesystem created by Silicon Graphics and later ported to Linux. It is used in Coraid CLN series NAS appliances.