

Designing Storage Architectures for Digital Collections 2024

Library of Congress

Common Terms for Attendees

In conjunction with the 2024 DSA meeting at the Library of Congress, below are several common terms provided by attendees. Please note, these items are not intended as exhaustive or prescriptive; they are merely intended to provide useful context related to the meeting topics.

BaFe or Barium Ferrite: A highly magnetic material, has a high packing density, and is a metal oxide. Studies of this material date at least as far back as 1931, and it has found applications in magnetic card strips, speakers, and magnetic tapes. One area in particular it has found success in is long-term data storage; the material is magnetic, resistant to temperature change, corrosion and oxidization.

Bit Error Rate: In digital transmission, the number of bit errors is the number of received bits of a data stream over a communication channel that have been altered due to noise, interference, distortion or bit synchronization errors. The bit error rate (BER) is the number of bit errors per unit time. The bit error ratio (also BER) is the number of bit errors divided by the total number of transferred bits during a studied time interval. Bit error ratio is a unitless performance measure, often expressed as a percentage. The bit error probability is the expectation value of the bit error ratio. The bit error ratio can be considered as an approximate estimate of the bit error probability. This estimate is accurate for a long time interval and a high number of bit errors.

Block (data storage): In computing (specifically data transmission and data storage), a block, sometimes called a physical record, is a sequence of bytes or bits, usually containing some whole number of records, having a maximum length, a block size. Data thus structured are said to be blocked. The process of putting data into blocks is called blocking, while deblocking is the process of extracting data from blocks. Blocked data is normally stored in a data buffer and read or written a whole block at a time. Blocking reduces the overhead and speeds up the handling of the data-stream. For some devices, blocking reduces the amount of external storage required for the data. Blocking is almost universally employed when storing data to 9-track magnetic tape, NAND flash memory, and rotating media such as floppy disks, hard disks, and optical discs. Most file systems are based on a block device, which is a level of abstraction for the hardware responsible for storing and retrieving specified blocks of data, though the block size in file systems may be a multiple of the physical block size. This leads to space inefficiency due to internal fragmentation, since file lengths are often not integer multiples of block size, and thus the last block of a file may remain partially empty. This will create slack space. Block storage is normally abstracted by a file system or database management system (DBMS) for use by applications and end users. The physical or logical volumes accessed via block I/O may be devices internal to a server, directly attached via SCSI or Fibre Channel, or distant devices accessed via a storage area network (SAN) using a protocol such as iSCSI, or AoE. DBMSes

often use their own block I/O for improved performance and recoverability as compared to layering the DBMS on top of a file system.

Block-level Storage: A concept in cloud-hosted data persistence where cloud services emulate the behavior of a traditional block device, such as a physical hard drive. It is a form of Network Attached Storage (NAS). Storage in such is organized as blocks. This emulates the type of behavior seen in traditional disk or tape storage. Blocks are identified by an arbitrary and assigned identifier by which they may be stored and retrieved, but this has no obvious meaning in terms of files or documents. A filesystem must be applied on top of the block-level storage to map 'files' onto a sequence of blocks. Amazon EBS (Elastic Block Store) is an example of a cloud block store. Cloud block-level storage will usually offer facilities such as replication for reliability, or backup services. Block-level storage is in contrast to an object store or 'bucket store', such as Amazon S3 (Simple Storage Service) or to a database. These operate at a higher level of abstraction and are able to work with entities such as files. Instance stores are another form of cloud-hosted block-level storage. These are provided as part of an 'instance', such as an Amazon EC2 (Elastic Compute Cloud) service. As EC2 instances are primarily provided as compute resources, rather than storage resources, their storage is less robust. Their contents will be lost if the cloud instance is stopped. As these stores are part of the instance's virtual server, they offer higher performance and bandwidth to the instance. They are best used for temporary storage such as caching or temporary files, with persistent storage held on a different type of server. At one time, block-level storage was provided by Storage Area Networks (SAN) and NAS provided file-level storage. With the shift from on-premises hosting to cloud services, this distinction has shifted. Even block-storage is now seen as distinct servers (thus NAS), rather than the previous array of bare discs.

Dataverse: Can be used to mean a composable data ecosystem. May or may not be tied to products of the same name from Harvard University, Microsoft, and Ownership Labs.

Disk Array: A disk storage system that contains multiple disk drives. It includes cache memory and controllers.

Enterprise Storage: A centralized repository for business information that provides common data management, protection and data sharing functions through connections to computer systems

Error Correction Code (ECC): Used for controlling errors in data over unreliable or noisy communication channels. The central idea is the sender encodes the message with a redundancy in the form of an ECC. The redundancy allows the receiver not only to detect errors that may occur anywhere in the message, but often to correct a limited number of errors.

Failure Rate: The frequency with which an engineered system or component fails, expressed in failures per unit of time. It is often denoted by the Greek letter λ (lambda) and is highly used in reliability engineering. The failure rate of a system usually depends on time, with the rate varying over the life cycle of the system. For example, an automobile's failure rate in its fifth year of service may be many times greater than its failure rate during its first year of service. One does not expect to replace an exhaust pipe, overhaul the brakes, or have major transmission problems in a new vehicle. In practice, the mean time between failures (MTBF, $1/\lambda$) is often reported instead of the failure rate. This is valid and useful if the failure rate may be assumed constant – often used for complex units / systems, electronics – and is a general agreement in some reliability standards (Military and Aerospace). It does in this case only relate to the flat region of the bathtub curve, which is also called the "useful life period". Because of this, it is

incorrect to extrapolate MTBF to give an estimate of the service lifetime of a component, which will typically be much less than suggested by the MTBF due to the much higher failure rates in the "end-of-life wearout" part of the "bathtub curve". The reason for the preferred use for MTBF numbers is that the use of large positive numbers (such as 2000 hours) is more intuitive and easier to remember than very small numbers (such as 0.0005 per hour). The MTBF is an important system parameter in systems where failure rate needs to be managed, in particular for safety systems. The MTBF appears frequently in the engineering design requirements, and governs frequency of required system maintenance and inspections. In special processes called renewal processes, where the time to recover from failure can be neglected and the likelihood of failure remains constant with respect to time, the failure rate is simply the multiplicative inverse of the MTBF ($1/\lambda$).

Fixity Checking: The practice of algorithmically reviewing digital content to insure that it has not changed over time. See: Fixity Survey Report – An NDSA Report: <https://osf.io/2qkea/>

Flash Memory: An electronic non-volatile storage medium that can be electrically erased and reprogrammed. The two types of Flash Memory are NAND and NOR, named for logic gates. In addition to being non-volatile, flash memory offers fast read access times, although not as fast as static RAM or ROM. Its mechanical shock resistance helps explain its popularity over hard disks in portable devices, as does its high durability, ability to withstand high pressure, temperature and immersion in water, etc.

HA: High Availability.

HAMR: Heat-assisted magnetic recording is a magnetic storage technology for greatly increasing the amount of data that can be stored on a magnetic device such as a hard disk drive by temporarily heating the disk material during writing, which makes it much more receptive to magnetic effects and allows writing to much smaller regions (and much higher levels of data on a disk).

HDD: A hard disk drive (HDD), hard disk, hard drive, or fixed disk[b] is an electro-mechanical data storage device that uses magnetic storage to store and retrieve digital information using one or more rigid rapidly rotating disks (platters) coated with magnetic material. The platters are paired with magnetic heads, usually arranged on a moving actuator arm, which read and write data to the platter surfaces. Data is accessed in a random-access manner, meaning that individual blocks of data can be stored or retrieved in any order and not only sequentially. HDDs are a type of non-volatile storage, retaining stored data even when powered off.

Hyperscale: The ability of an architecture to scale appropriately as increased demand is added to the system. This typically involves the ability to seamlessly provide and add compute, memory, networking, and storage resources to a given node or set of nodes that make up a larger computing, distributed computing or grid computing environment. Hyperscale computing is necessary in order to build a robust and scalable cloud, big data, map reduce or distributed storage systems and is often associated with the infrastructure required to run large distributed sites. Companies that build data centers for hyperscale computing such as cloud providers are sometimes called "hyperscalers".

LTO: Linear Tape Open is a magnetic tape data storage technology originally developed in the late 1990s as an open standards alternative to the proprietary magnetic tape formats that were available at the time. The LTO Consortium directs development and manages licensing and certification of media and mechanism manufacturers.

MAMR: Microwave Assisted Magnetic Recording is a type of hard disk drive technology.

Meantime Between Failures (MTBF): See Failure Rate

NAND: See Flash Memory.

Network-attached Storage (NAS): File-level computer data storage server connected to a computer network providing data access to a heterogeneous group of clients. NAS is specialized for serving files either by its hardware, software, or configuration. It is often manufactured as a computer appliance – a purpose-built specialized computer. NAS systems are networked appliances which contain one or more storage drives, often arranged into logical, redundant storage containers or RAID. Network-attached storage removes the responsibility of file serving from other servers on the network. They typically provide access to files using network file sharing protocols such as NFS, SMB/CIFS, or AFP. Potential benefits of dedicated network-attached storage, compared to general-purpose servers also serving files, include faster data access, easier administration, and simple configuration.

NVMe (Non-Volatile Memory Express): A logical device interface, which has been designed to capitalize on the low latency and internal parallelism of solid-state storage devices.

Object Storage (also known as object-based storage): A computer data storage architecture that manages data as objects, as opposed to other storage architectures like file systems which manage data as a file hierarchy, and block storage which manages data as blocks within sectors and tracks. Each object typically includes the data itself, a variable amount of metadata, and a globally unique identifier. Object storage can be implemented at multiple levels, including the device level (object-storage device), the system level, and the interface level. In each case, object storage seeks to enable capabilities not addressed by other storage architectures, like interfaces that can be directly programmable by the application, a namespace that can span multiple instances of physical hardware, and data-management functions like data replication and data distribution at object-level granularity. Object-storage systems allow retention of massive amounts of unstructured data. Object storage is used for purposes such as storing photos on Facebook, songs on Spotify, or files in online collaboration services, such as Dropbox.

POSIX (Portable Operating System Interface): A family of standards specified by the IEEE Computer Society for maintaining compatibility between operating systems. POSIX defines the application programming interface (API), along with command line shells and utility interfaces, for software compatibility with variants of Unix and other operating systems.

Preservation Storage: The use of storage technology for digital preservation. In common practice, digital material has become decoupled from the underlying mechanism of its storage. With this comes consequent benefits of allowing different preservation activities to be handled independently.

RAID (Redundant Array of Independent Disks): A data storage virtualization technology that combines multiple physical disk drive components into one or more logical units for the purposes of data redundancy, performance improvement, or both. Data is distributed across the drives in one of several ways, referred to as RAID levels, depending on the required level of redundancy and performance. The different schemes, or data distribution layouts, are named by the word “RAID” followed by a number, for example RAID 0 or RAID 1. Each scheme, or RAID level, provides a different balance among the key goals: reliability, availability, performance, and

capacity. RAID levels greater than RAID 0 provide protection against unrecoverable sector read errors, as well as against failures of physical drives.

REST: A software architectural style that defines a set of constraints to be used for creating Web services. Web services that conform to the REST architectural style, called RESTful Web services (RWS), provide interoperability between computer systems on the Internet.

SMR (Shingled Magnetic Recording): A magnetic storage data recording technology used in hard drives to increase storage density and overall per-drive capacity.

Software-Defined Storage is software for policy-based provisioning and management of data storage independent of the underlying hardware. Software-defined storage typically includes a form of storage virtualization to separate the storage hardware from the software that manages it. The software enabling a software-defined storage environment may also provide policy management for features such as data duplication, replication, snapshots, and backup.

Solid State Memory: A hard disk drive technology. See Flash Memory.

SSD (Solid State Drive): A storage device using integrated circuits to store data persistently, typically using Flash memory. See Flash Memory.

Storage Area Network (SAN): A storage area network (SAN) or storage network is a computer network which provides access to consolidated, block level data storage. SANs are primarily used to enhance storage devices, such as disk arrays and tape libraries, accessible to servers so the devices appear to the operating system as locally attached devices. A SAN typically is a separate network of storage devices. A SAN does not provide file abstraction, only block level access. File systems built on top of SANs do provide file disk file systems.

Zettabyte: A billion terabytes.

Zero Trust: Zero Trust, Zero Trust Network, or Zero Trust Architecture refer to security concepts and threat model that no longer assumes that actors, systems or services operating from within the security perimeter should be automatically trusted, and instead must verify anything and everything trying to connect to its systems before granting access.