

Fundamentals Of Storage

Disk Hardware, Connectors, OS Storage Subsystems, File Systems, RAID, Tx, Caching, Hashing, B-Trees

This course provides a comprehensive tour of the important storage fundamentals that all technologists need to understand in order to build, provision and operate modern IT solutions that involve storage (which of course means all such solutions). There's a lot more to storage than the hard disk on a PC or simple APIs such as fopen() / fwrite() that developers might use. In this course we explore storage end-to-end and see how it interacts with remote disks, hypervisors, the cloud and various in-memory representations. We see how the distinction between storage, database and memory is being blurred. We see how system-wide reliability and

performance demands impart on storage. We see how modern OSes and their powerful storage stacks allow rich journaling file systems and databases to be built that deliver a wide variety of enhanced storage features. We also explore specific technical approaches to transacting, caching, hashing and B-Trees.

A good understanding of storage fundamentals along with clear knowledge of the storage-related technical options available to choose from helps all involved in engineering to make the optimum design decisions.

Contents of One-Day Training Course	
<p>Target Audience Developers, devops, IT professionals, engineering managers – all of whom need an understanding of the core building blocks of storage platforms.</p> <p>Prerequisites Good all-round knowledge of modern computing infrastructure at a technical level.</p>	<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> <p style="text-align: center;">Storage Concepts</p> <p>Latency Throughput Reliability “bit-rot” and other issues</p> <p style="text-align: center;">Modern Storage for ...</p> <p>File Object XML SQL Graph</p> <p style="text-align: center;">SCSI – Talking To Storage</p> <p>Small Computer System Interface (SCSI) is a comprehensive protocol / command / interface between controller and disk</p> <p style="text-align: center;">Disk hardware and connectors</p> <p>Traditional disk SATA / SAS SSD and Flash NVMe and PCI Express RAM disk</p> <p style="text-align: center;">Tape, Optical</p> <p>Managing larger amounts of data Tape and optical formats used for backup, transporting and long-term archiving</p> <p style="text-align: center;">DAS, NAS, SAN</p> <p>Directly Attached Storage Network Attached Storage Storage Area Network</p> <p style="text-align: center;">Storage Subsystem In An OS</p> <p>Storage and networking are the largest subsystems inside any OS Exploring an OS storage architecture</p> </div> <div style="width: 48%;"> <p style="text-align: center;">Modern File Systems</p> <p>Journaling 128-bit sizes Virtual File Systems And Client File APIs Detailed look at internals of a modern file system – ZFS</p> <p style="text-align: center;">RAID</p> <p>Costs and benefits of various RAID types Stripping When errors occur in hardware and in software</p> <p style="text-align: center;">Transactioning</p> <p>Transaction principles Two phase commit Hierarchical transaction groups Savepoints Resource manager Enlistment</p> <p style="text-align: center;">Caching</p> <p>Putting data close to where it is used Cache lifetimes Keeping data in cache fresh Caching architecture</p> <p style="text-align: center;">Identity And Hashing</p> <p>How to identify storage objects Role of hashing algorithms Importance of hashing for de-duplication and other aspects of storage</p> <p style="text-align: center;">B-Tree</p> <p>B-Tree is the most important data structure for storage architectures Performance characteristics of B-Trees Variations of B-Tree layout</p> </div> </div>