System & Service Management
Data Center Technology

Data Centers
A data center (or data centre or datacentre) is a facility used to house computer systems and associated components, such as telecommunications and storage systems.
The TIA-942:Data Center Standards Overview describes the requirements for the data center infrastructure.

The four levels are defined, and copyrighted, by the Uptime Institute. The levels describe the availability of data from the hardware at a location. The higher the tier, the greater the accessibility.

Tier 1	• Single non-redundant distribution path serving the IT equipment • Non-redundant capacity components • Basic site infrastructure guaranteeing 99.671% availability
Tier 2	• Fulfils all Tier 1 requirements • Redundant site infrastructure capacity components guaranteeing 99.741% availability
Tier 3	• Fulfils all Tier 1 & Tier 2 requirements • Multiple independent distribution paths serving the IT equipment • All IT equipment must be dual-powered and fully compatible with the topology of a site's architecture • Concurrently maintainable site infrastructure guaranteeing 99.982% availability
Tier 4	• Fulfils all Tier 1, Tier 2 and Tier 3 requirements • All cooling equipment is independently dual-powered, including chillers and Heating, Ventilating and Air Conditioning (HVAC) systems • Fault tolerant site infrastructure with electrical power storage and distribution facilities guaranteeing 99.995% availability

http://en.wikipedia.org/wiki/Data_center

Hierarchical Storage Management / Tiered Storage
Hierarchical Storage Management (HSM) is a data storage technique which automatically moves data between high-cost and low-cost storage media. HSM systems exist because high-speed storage devices, such as hard disk drive arrays, are more expensive (per byte stored) than slower devices, such as optical discs and magnetic tape drives. While it would be ideal to have all data available on high-speed devices all the time, this is prohibitively expensive for many organizations. Instead, HSM systems store the bulk of the enterprise's data on slower devices, and then copy data to faster disk drives when needed. In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM system monitors the way data is used and makes best guesses as to which data can safely be moved to slower devices and which data should stay on the fast devices.
In a typical HSM scenario, data files which are frequently used are stored on disk drives, but are eventually migrated to tape if they are not used for a certain period of time, typically a few months. If a user does reuse a file which is on tape, it is automatically moved back to disk storage. The advantage is that the total amount of stored data can be much larger than the capacity of the disk storage available, but since only rarely-used files are on tape, most users will usually not notice any slowdown.
HSM is sometimes referred to as tiered storage.
http://en.wikipedia.org/wiki/Hierarchical_storage_management

Tiered Storage
Tiered storage is the assignment of different categories of data to different types of storage media in order to reduce total storage cost. Categories may be based on levels of protection needed, performance requirements, frequency of use, and other considerations. Since assigning data to particular media may be an ongoing and complex activity, some vendors provide software for automatically managing the process based on a company-defined policy.
As an example of tiered storage, tier 1 data (such as mission-critical, recently accessed, or top secret files) might be stored on expensive and high-quality media such as double-parity RAID volumes. Tier 2 data (such as financial, seldom-used, or classified files) might be stored on less expensive media in conventional storage area networks (SANs). As the tier number increased, cheaper media could be used. Thus, tier 3 in a 3-tier

system might contain event-driven, rarely used, or unclassified files on recordable compact discs (CD-Rs) or tapes.
http://searchstorage.techtarget.com/sDefinition/0,,sid5_gci1028962,00.html

Übung zum Thema „Allozierungseffizienz“
Es gelte folgendes Szenario:
- 96 Platten a 73 GB
- 6 Adapterkarten (4-Port)
- Redundante Pfade zu den Servern (max. 12)
- Anschaffungspreis: $0.1/MB
- Ähnliche Applikation pro Cluster mit 85% allocation efficiency

Möglicher Lösungsweg:
- 96 * 73GB = 7008GB → wegen Spiegelung (RAID1) fällt 50% Nutzkapazität weg → 3504GB Kapazität
- Verteilen auf 6 Cluster: 3504GB / 6 = 584GB → Effizienzverlust von 15% (wegen ungenutztem Speicher, also: Allozierungsineffizienz) → 496.4GB/Cluster

Folgender Verlust ist festzustellen:
→ Kapazitätsverlust von 87.6GB/Cluster bzw 1051GB total
→ Finanzieller Verlust von 1'051’000MB * 0.1$/MB = 105’100$ (!)
(In beiden Fällen wurde der Verlust von 50% Kapazität wegen der RAID1-Speigelung vernachlässigt).

Cost of Poor Quality
Cost of poor quality (COPQ) or poor quality costs (PQC), are defined as costs that would disappear if systems, processes, and products were perfect.

Cost element			Examples
Direct poor-quality costs	Controllable poor-quality cost	Prevention cost	• Quality planning (for test, inspection, audits, process control) • Education and training • Performing capability analyses • Conducting design reviews
		Appraisal cost	• Test and inspection • Supplier acceptance sampling • Auditing processes
	Resultant poor-quality cost	Internal error cost	• In-process scrap and rework • Troubleshooting and repairing • Design changes • Additional inventory required to support poor process yields and rejected lots • Reinspection and retest of reworked items • Downgrading
		External error cost	• Sales returns and allowances • Service level agreement penalties • Complaint handling • Field service labor and parts costs incurred due to warranty obligations
	Equipment poor-quality cost		Micrometers, voltmeters, automated test equipment (but **not** equipment used to make the product)
Indirect poor-quality costs	Customer-incurred cost		• Loss of productivity due to product or service downtime • Travel costs and time spent to return defective product • Repair costs after warranty period • Backup product or service to cover failure periods
	Customer-dissatisfaction cost		Dissatisfaction shared by word of mouth
	Loss-of-reputation cost		Customer perception of firm

http://en.wikipedia.org/wiki/Cost_of_poor_quality

Mean Time Between Failures
Mean time between failures (MTBF) is the predicted elapsed time between inherent failures of a system during operation. MTBF can be calculated as the arithmetic mean (average) time between failures of a system. The MTBF is typically part of a model that assumes the failed system is immediately repaired (zero elapsed time), as a part of a renewal process. This is in contrast to the mean time to failure (MTTF), which measures average time between failures with the modeling assumption that the failed system is not repaired.

The definition of MTBF depends on the definition of what is considered a system failure. For complex, repairable systems, failures are considered to be those out of design conditions which place the system out of service and into a state for repair. Failures which occur that can be left or maintained in an unrepaired condition, and do not place the system out of service, are not considered failures under this definition. In addition, units that are taken down for routine scheduled maintenance or inventory control, are not considered within the definition of failure.

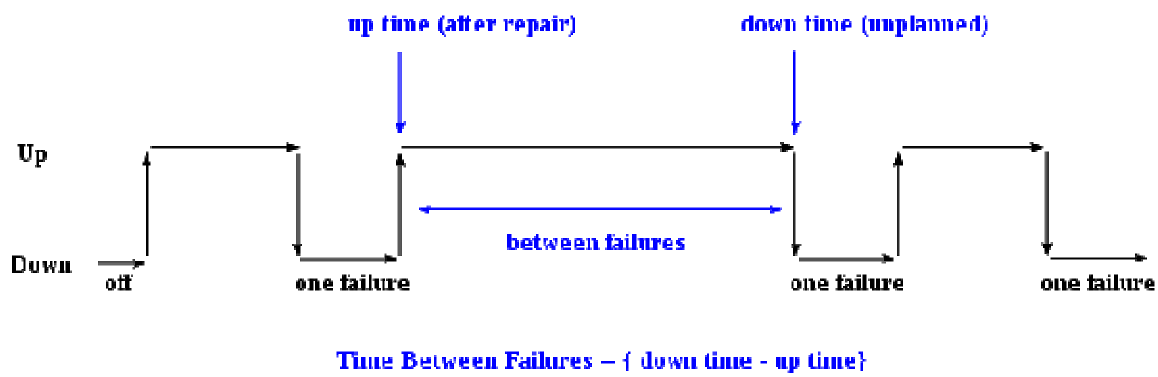

http://en.wikipedia.org/wiki/Mean_time_between_failures

Irreführende MTBF Angaben bei Harddisks
MTBF-Angaben, die bei neuesten Harddisk-Generationen mittlerweile bei über einer Million Stunden liegen, wurden laut Hitachi in der Vergangenheit überdies häufig fehlinterpretiert. So bedeuten beispielsweise eine MTBF von 1,2 Millionen Stunden nicht, dass das Laufwerk die nächsten 137 Jahre rund um die Uhr fehlerfrei funktioniert. Vielmehr versteht Hitachi unter dieser Angabe, dass bei einer installierten Basis von 1,2 Millionen Laufwerken, die unter idealen Temperatur- und Belastungsbedingungen betrieben werden, statistisch gesehen jede Stunde ein Laufwerk ausfällt. Dass ideale Betriebsbedingungen in der Praxis eher selten auftreten, bedarf keiner ausdrücklichen Erwähnung. Konkurrent Maxtor geht deshalb davon aus, dass die tatsächliche MTBF-Zeit unter realen Bedingungen gerade einmal 50 bis 60 Prozent der statistisch ermittelten MTBF-Dauer entspricht.

Doch nicht nur Umgebungseinflüsse wie Staub, Vibrationen oder eine zu hohe bzw. zu niedrige Umgebungstemperatur wirken sich negativ auf die Ausfallhäufigkeit von Festplatten aus. Auch die Belastung, dargestellt durch die Art und Dauer der Schreib-/Lese- und Suchzugriffe, kann sich insbesondere bei preiswerten ATA- bzw. SATA-Laufwerken deutlich in der Fehlerrate niederschlagen wie Festplattenhersteller Seagate in einer Studie ermittelte.

Der Analyse zugrunde lagen drei Gruppen von jeweils 100 handelsüblichen ATA-Drives, die über einen Zeitraum von 1.200 Stunden unterschiedlichen Belastungen ausgesetzt wurden. Bereits bei einem häufigen sequenziellen Zugriff verdoppelte sich die kumulierte Fehlerrate. Bei Betriebsbedingungen, die denen eines Low-End-Servers entsprechen (häufiger Random-Zugriff), vervierfachte sich die Fehlerrate der für den Desktop-Betrieb konzipierten ATA-Disks.

http://www.speicherguide.de/Infrastruktur/SATA/tabid/221/articleType/ArticleView/articleId/10508/MTBF-Zeiten-nur-Schall-und-Rauch.aspx