CS2640 Midterm (Practice) — Solution Key

Modern (Computer) Storage Systems

Spring 2026

Part I — Multiple Choice Answer Key

1	B	**2**	C
3	B	**4**	A
5	B	**6**	A
7	C	**8**	B
9	B	**10**	B

11. HDD vs SSD

- HDDs suffer mechanical seek and rotational latency; SSDs provide much lower access latency.
- Both can achieve high bandwidth for sequential reads/writes, but SSDs typically achieve a 10-100x higher bandwidth.
- Random I/O degrades sharply on HDDs; SSDs handle random reads much better, though random writes trigger garbage collection and write amplification.
- HDDs usually win on cost per TB; SSDs win on latency and IOPS.

12. SSD internal behavior

- Flash (NAND) cannot overwrite data in place, so updates are written out of place to fresh pages and the FTL remaps logical addresses to new physical pages.
- Invalid pages accumulate, so garbage collection must copy still-valid pages out of blocks before erasing them.
- That extra movement increases write amplification.
- As free space shrinks, garbage collection becomes more urgent and more expensive, causing unstable latency and throughput. Overprovisioning allows flash to avoid / reduce urgent garbage collection.

13. Filesystem design tradeoff

- Update-in-place filesystems preserve current-location layout and can be simpler to reason about, but require careful crash-consistency handling and can fragment small updates.
- Log-structured filesystems batch writes into large sequential appends, improving write efficiency, but they pay cleaning/segment-reclamation costs and may age poorly under some workloads.

14. Distributed storage semantics

- Consistency across replicas: systems must decide when an update is visible and how to reconcile failures or partitions.
- Failure detection and recovery over a network: nodes may be slow, partitioned, or fail independently, making coordination, re-replication, and metadata placement substantially harder than in a local filesystem.

15. Reliability metrics

- Redundancy alone is not enough because failures are often correlated across disks, racks, firmware versions, or environmental conditions.
- During rebuild, the system is in a degraded state for a long window, so a second failure is more dangerous.
- Latent sector errors or unreadable blocks may appear exactly when data is needed for rebuild.
- Human error, software bugs, and controller or firmware bugs can defeat nominal redundancy.

16. RAID reasoning

(a) RAID-0 offers full capacity and high performance but no redundancy; any disk failure loses the array. RAID-1 with four mirrored pairs cuts usable capacity in half, offers strong read performance, and can survive one failure per mirror pair. RAID-5 provides the capacity of seven disks, tolerates one disk failure, and pays parity-update overhead on writes.

(b) RAID-0 is riskiest because any single disk failure causes data loss.

(c) Risk increases during rebuild because the array is degraded, rebuilds add heavy load, and another failure or latent read error can cause permanent loss.