


Reliability Hierarchies

Peter M. Chen
David E. Lowell




Computer Science and Engineering Division
Electrical Engineering and Computer Science
University of Michigan

Performance Hierarchies



Reliability Hierarchies



Write-Back Policy

When to transfer data to lower, more reliable level?

Write-through

- most reliable
- effectively eliminates upper level from reliability hierarchy

Delayed-write

- e.g. write new data to memory, then transfer to disk after 15 seconds
- trade-off between reliability and overhead (e.g. performance)

Metrics for a Reliability Hierarchy

Mean time to data loss (MTTDL)

- limited by reliability of highest (least reliable) level
- doesn't distinguish between degrees of data loss

Data loss rate

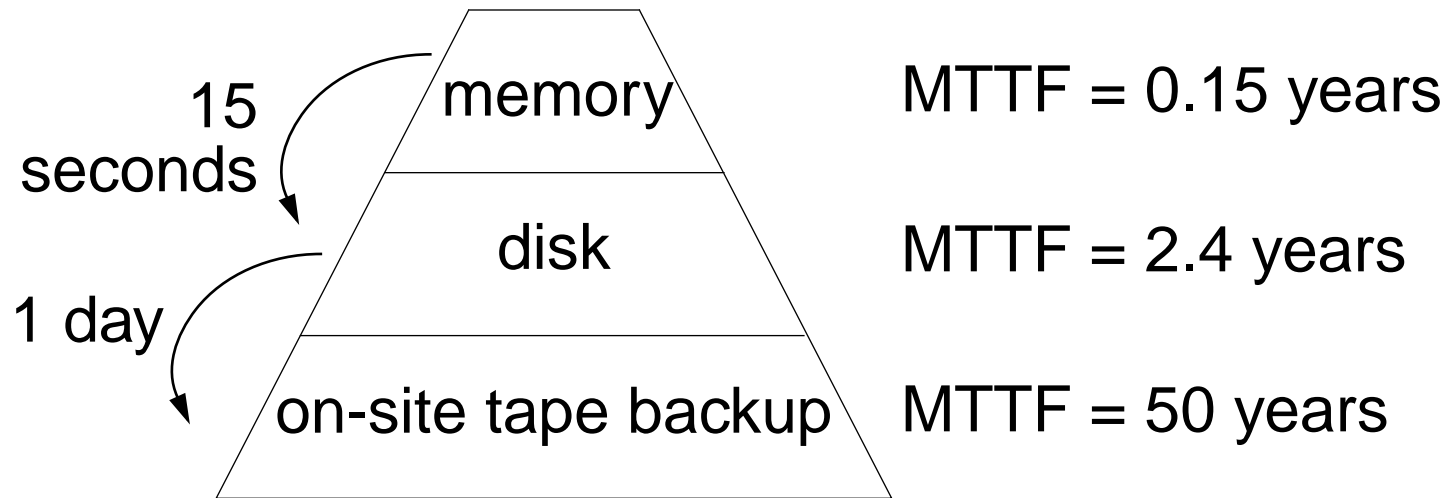
- fraction of new data lost over time

$$\sum_{\text{all levels}} \frac{\text{data loss}_L}{MTTF_L}$$

Example Faults and Storage Levels

Fault Category	Example MTTF	Storage Levels Affected by Fault				
		CPU/ memory	disk	RAID	on-site backup	remote backup
operating system	2 months	✓				
file system	5 years	✓	✓	✓		
power	10 years (UPS)	✓				
motherboard	5 years	✓				
media	5 years		✓			
catastrophe	50 years	✓	✓	✓	✓	

Analysis of Michigan Server




overall MTDDL = 0.15 years

data loss rate = 10 hours/year

Rio on PCs

New level in the storage hierarchy: **reliable main memory**


Enable memory to survive operating system crashes



Example Faults and Storage Levels

Fault Category	Example MTTF	Storage Levels Affected by Fault		
		CPU/memory	CPU/memory with Rio	disk
operating system	2 months	✓		
file system	5 years	✓	✓	✓
power	10 years (UPS)	✓	✓	
motherboard	5 years	✓	✓	
media	5 years			✓
catastrophe	50 years	✓	✓	✓

Rio's Effect on Reliability



MTTF = 0.15 years → 1.9 years

MTTF = 2.4 years

MTTF = 50 years

overall MTDDL = 0.15 years → 1.4 years

data loss rate = 10 hours/year

Conclusions

Two views of hierarchies

- trade-off between cost and performance
- trade-off between reliability and performance/
cost/power/etc.

Rio fills in the “reliability gap” between memory and disk

- hypothesis: can use Rio to store new types of data that would like higher reliability than memory but can't afford overhead of disk



<http://www.eecs.umich.edu/Rio>