

LSM Compactions

Design Choice

The Design Space

Demonstration Setup

### LSMs everywhere

Fast Ingestion

+

Good Reads

+

Great Space Utilization

Periodic Compaction

### Full Compaction vs. Partial Compaction

compact all data in a level

compact only part of a level

memory buffer

level capacity

file

file to compact

files after compaction

### The Compaction Black Box

workload

LSM-tuning

performance

Which compaction strategy to use ?

### The Design Questions

How to lay out the data?

How much data to move?

Which data block to be moved?

When to re-organize layout?

### Design Primitives

Data Layout

Compaction Granularity

Data Movement Policy

Compaction Trigger

### Data Layout

leveling

tiering

hybrid

### Granularity

levels

file(s)

runs

### Data Movement Policy

Which file(s) to pick?

- Round-robin
- Least overlap
- Coldest
- Oldest
- Most tombstones

### Trigger

saturation

#runs

- Space ampl.
- Age of a file
- Read ampl.

### Primitive Ensemble

Full	L	leveling
	G	levels
	D	N/A
	T	level saturation
1-Lvl	L	hybrid leveling
	G	file
	D	least overlap w/ parent
	T	level saturation
Tier	L	tiering
	G	runs (in same level)
	D	N/A
	T	#runs / space ampl.

### Compactionary in Action

Comparative Analysis

Individual Analysis

Workload

Main Memory Parameter

Disk Parameters

Data Layout

Advanced settings

Progress

Play

Pause

Finish

Step

Vanilla-LSM

Partial Compaction

Hybrid Strategy

Build-Your-Own

Levels

#sorted runs

#compactions

4	3	4	4
2	3	4	2
97	2926	25	25

### Key Takeaways

Know your LSM-compaction

Avoid the worst designs

Adapt choices w/ workloads

Design new compactions