



BazelCon 2026


Stop delivering your whole fleet



Sahin Yort

Sr. Engineer at Aspect

#BazelCon




“What changed?” is the wrong question

In a repo with 204 tests, measuring different kinds of changes — without execution, you can't know.

change	expected	bazel-diff	remote cache
one leaf source	1	1 ✓	1 ✓
core library	143	143 ✓	143 ✓
comment in a header	0	143 ✗	0 ✓
--copt=-DF00	0	0 ✓	0 ✓
--action_env	204	0 ✗	204 ✓


Your remote cache already knows

Ask every question, execute nothing



Build everything, then decide

Read left to right — the numbers trace the order each step runs, 1 through 9.




How do you hash a binary target?

 + hashsum aspect



added on top of the target



`GetActionResult`, never executes




`7f3e...`

one per deliverable
hit or miss




BazelCon 2026

Ask again without starting over



A hermeticity test for free



Where it bites

The cache is only as honest as your build



flags differ from CI

- ✗ every key looks changed
- ✓ probe reuses the same rc and config



no-remote / local actions

- ✗ run outside the cache
- ✓ fine if the action is deterministic



stamped outputs

- ✗ digest moves every commit
- ✓ stamp only in the release build



cache evicts entries

- ✗ looks like a miss
- ✓ rebuild, nothing ships twice



What does it save?

100 heavy tests: one runner vs. cache check + shards

the repo

-  10 services
-  100 heavy tests
-  130 CPU-min total
-  4-core CI runners

 one runner, every test

 cache check + sharded runners

no change

0 tests affected

 1 min

 1 min

one service changed

10 tests affected

 5.93 min

 7.15 min

core library changed

100 tests affected

 17.92 min

 8 min



BazelCon 2026



cold runners: every shard reruns full analysis + loading

One oracle, many verbs

Up-to-dateness is already computed — read it

