

VLDB Access Metrics

Andrew Deason

June 2022

OpenAFS Workshop 2022

Background

- VLDB: **V**olume **L**ocation **D**atabase
- Served by the `vlserver` process on database servers
- `/afs/cell/vol.foo` → needs volume `vol.foo`
- Lookup `vol.foo` in the VLDB to know what file servers to contact

Background

- In 2019, Mark Vitale presented “VLserver memory cache”¹
 - Included tools to look at cache
 - Cache contents → popular volumes (names)
- In 2021, I talked about ubik-kv²
 - No more cache
 - No more data!
- Shouldn't need to rely on storage layer for stats
 - DTrace/systemtap/etc³
 - wireshark
 - audit logs

¹<https://workshop.openafs.org/afsbpw19/2019/schedule/vlserver-memory-cache>

²<https://workshop.openafs.org/afsbpw21/talks/next-gen-ubik-and-the-vldb/>

³<https://workshop.openafs.org/afsbpw19/2019/schedule/dtrace-systemtap-probes-in-openafs/>

What stats?

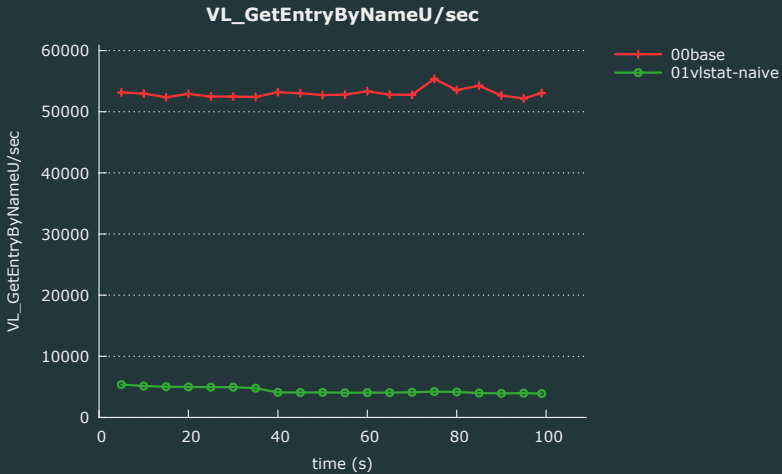
- New subsystem: `v1stat`
- Popular volumes (hits)
 - Alternative to looking at fileserver stats
 - Possible client errors / weird behavior
- Popular bad volumes (negative hits, misses)
 - Misspelled/deleted/etc mountpoints
 - or rare db errors
- Popular clients
 - Traversing the tree
 - Badly behaving clients

Storage

- Not fixed-size stats, sorted
- In-memory → limits, dropping data
- On-disk → higher limits, persistent
- Don't use ubik
 - RO sites
 - Write-heavy, don't need consistency
 - Reuse kv layer

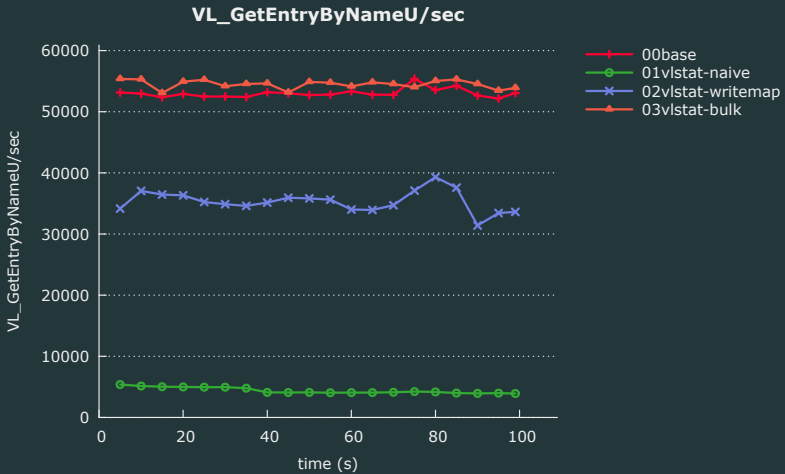
Storage

Write-heavy performance, not great:



Storage

But it gets better:



Limits

- Volume hits, bounded
 - $O(VLDB)$
- Volume misses, unbounded
 - Need some kind of limit
 - LRU
 - Should be uncommon
- Clients, semi-bounded?

Example

Looking at the top N entries:

```
$ openafs-ctl vldb-stats-get
total getentry volume hit records: 50001
top volumes:
- 538731264 (vol.vwRfStcXJ), 44 hits
  first request: host 127.0.0.1, time 1652895592.151 (Wed May 18 13:39:52 2022)
  latest request: host 127.0.0.1, time 1652993063.317 (Thu May 19 16:44:23 2022)
- 542141952 (vol.0tJ0FrdHD), 44 hits
  first request: host 127.0.0.1, time 1652895592.341 (Wed May 18 13:39:52 2022)
  latest request: host 127.0.0.1, time 1652993063.335 (Thu May 19 16:44:23 2022)
[...]
total getentry volume miss records: 1024 (dropped 66428)
top volume misses:
- vol.lVBH_zzz, 1 hits
  first request: host 127.0.0.1, code 363524, time 1652899551.681 (Wed May 18 14:45:51 2022)
  latest request: host 127.0.0.1, code 363524, time 1652899551.681 (Wed May 18 14:45:51 2022)
- vol.pXSfEzzz, 1 hits
  first request: host 127.0.0.1, code 363524, time 1652899551.681 (Wed May 18 14:45:51 2022)
  latest request: host 127.0.0.1, code 363524, time 1652899551.681 (Wed May 18 14:45:51 2022)
[...]
total getentry host records: 2
top hosts:
- 127.0.0.1, 1719334 hits
  first request: vol.foo, code 363524, time 1652815660.132 (Tue May 17 15:27:40 2022)
  latest request: vol.MUtqnS0_0, code 0, time 1652993069.378 (Thu May 19 16:44:29 2022)
- 192.0.2.5, 18 hits
  first request: foo3, code 363524, time 1649446338.496 (Fri Apr 8 15:32:18 2022)
  latest request: foo3, code 363524, time 1651785934.618 (Thu May 5 17:25:34 2022)
```

Rotation

- Need some kind of reset/rotation
 - Volume has 5MM hits, but from when?
 - Automatic daily (?) reset
 - Similar to volume `dayUse/weekUse` fields
- For long term trends, export to some other system

```
$ openafs-ctl vldb-stats-get --day 2
```

```
$ openafs-ctl vldb-stats-export --format json
```

Postamble

Contact

adeason@dson.org

adeason@sinenomine.net

Slides

<http://dson.org/talks>

Questions?

The background of the slide features a repeating pattern of question marks in a light gray color, set against a dark gray background. The question marks are arranged in a grid-like fashion, slightly offset from each other, creating a subtle texture.