

PRAGUE MEETUP

Real-Time Data & AI

Scale, Speed, and Insights



Michal Fizek
Development Team Leader



Exploring Aerospike Enterprise

Our Early Journey



Who we are

- RUS = *Registered User of Seznam (Registrovaný uživatel Seznamu)*.
- We maintain user profiles and sessions across Seznam services
- Small team (4 + 2 borrowed colleagues on backend)
- Our focus: stability first, then evolution



What the System Does

- Core backend for sessions & profiles of Seznam users
- Handles GDPR consent and paid products data
- Powers many services (Email, Stream, Mapy, Zboží ...)
- 500k requests per second ($\approx$ 200k UDF) – fully replicated
- Run's on Seznam's own infrastructure



Why Move to Aerospike Enterprise

- Security → encryption & compliance
- Robustness → native multi-DC replication
- Predictable performance & upgrade path
- Better hardware efficiency → fewer clusters, less virtualization noise
- Expert support & long-term stability



Early Findings and Lessons



Cluster Architecture over Datacenters

Our Current Setup

- 3 datacenters, very low latency (sub-ms)
- Data replicated to all DCs
- Goal: “*one DC down — nobody notices*”



Architectures We Considered

- XDR
- Stretched Cluster
- “Poor-man’s Stretched” (rack-awareness not available)



Architectures & Trade-offs

Architecture

Pros

Cons

XDR

Simple setup and operations •
Each DC can run
independently • Well supported

Eventual consistency across
DCs • Critical ops need extra
logic

Stretched Cluster

Strong consistency • Simplifies
app usage (no XDR logic) •
Automatic rebalancing

Sensitive to flapping links •
Needs rack-awareness • More
fragile for us now

“Poor-man’s Stretched”

Survives full DC loss • No
special license needed

Complicated placement and
monitoring • Hard to debug •
Not officially supported





Cluster Performance

Virtualized vs Bare-Metal


- Same hardware → different environments (OpenStack VM vs bare-metal)
- Simple read/write workload (no UDFs)



Benchmark Results: OpenStack



Benchmark Results: Baremetal



Results & Takeaways

Virtualized vs Bare-Metal

Environment	p99.9 write latency	Notes
<i>OpenStack (virtualized)</i>	> 200 ms (occasional spike up to 800 ms spikes)	Same HW → large overhead likely from network virtualization
<i>Bare-metal</i>	< 2 ms (2–4 ms at sometimes)	Stable latency curve until very high throughput



SEZNAM.CZ