



TIMESCALE

Erik Nordström

hello@timescale.com

Lots of Reasons to use PostgreSQL for time-series

- Reliability
- Ease of use
- Broad ecosystem
- Flexible datatypes
- Geospatial support
- Power of JOINS



measurements		metadata	
time	temp	container_id	cargo
2017-01-01 01:03:42	56	ghi789	furniture
2017-01-01 01:03:35	64	def456	food
2017-01-01 01:03:12	72	abc123	electronics
2017-01-01 01:02:30	56	ghi789	food
2017-01-01 01:02:23	64	def456	books
2017-01-01 01:02:00	72	abc123	toys


measurements**metadata**

timestamp	temperature	cpu_1m_avg	free_mem	container_id	type	location_id
2017-01-01 01:03:42	56	100	100MB	ghi789	big	77
2017-01-01 01:03:35	64	95	350MB	def456	small	335
2017-01-01 01:03:12	72	80	500MB	abc123	medium	335
2017-01-01 01:02:30	56	120	0MB	ghi789	big	77
2017-01-01 01:02:23	64	90	400MB	def456	medium	335
2017-01-01 01:02:00	72	80	500MB	abc123	small	335

Common scenario: Measurements + Metadata

measurements

time	temp	container_id
2017-01-01 01:03:42	56	ghi789
2017-01-01 01:03:35	64	def456
2017-01-01 01:03:12	72	abc123
2017-01-01 01:02:30	56	ghi789
2017-01-01 01:02:23	64	def456
2017-01-01 01:02:00	72	abc123

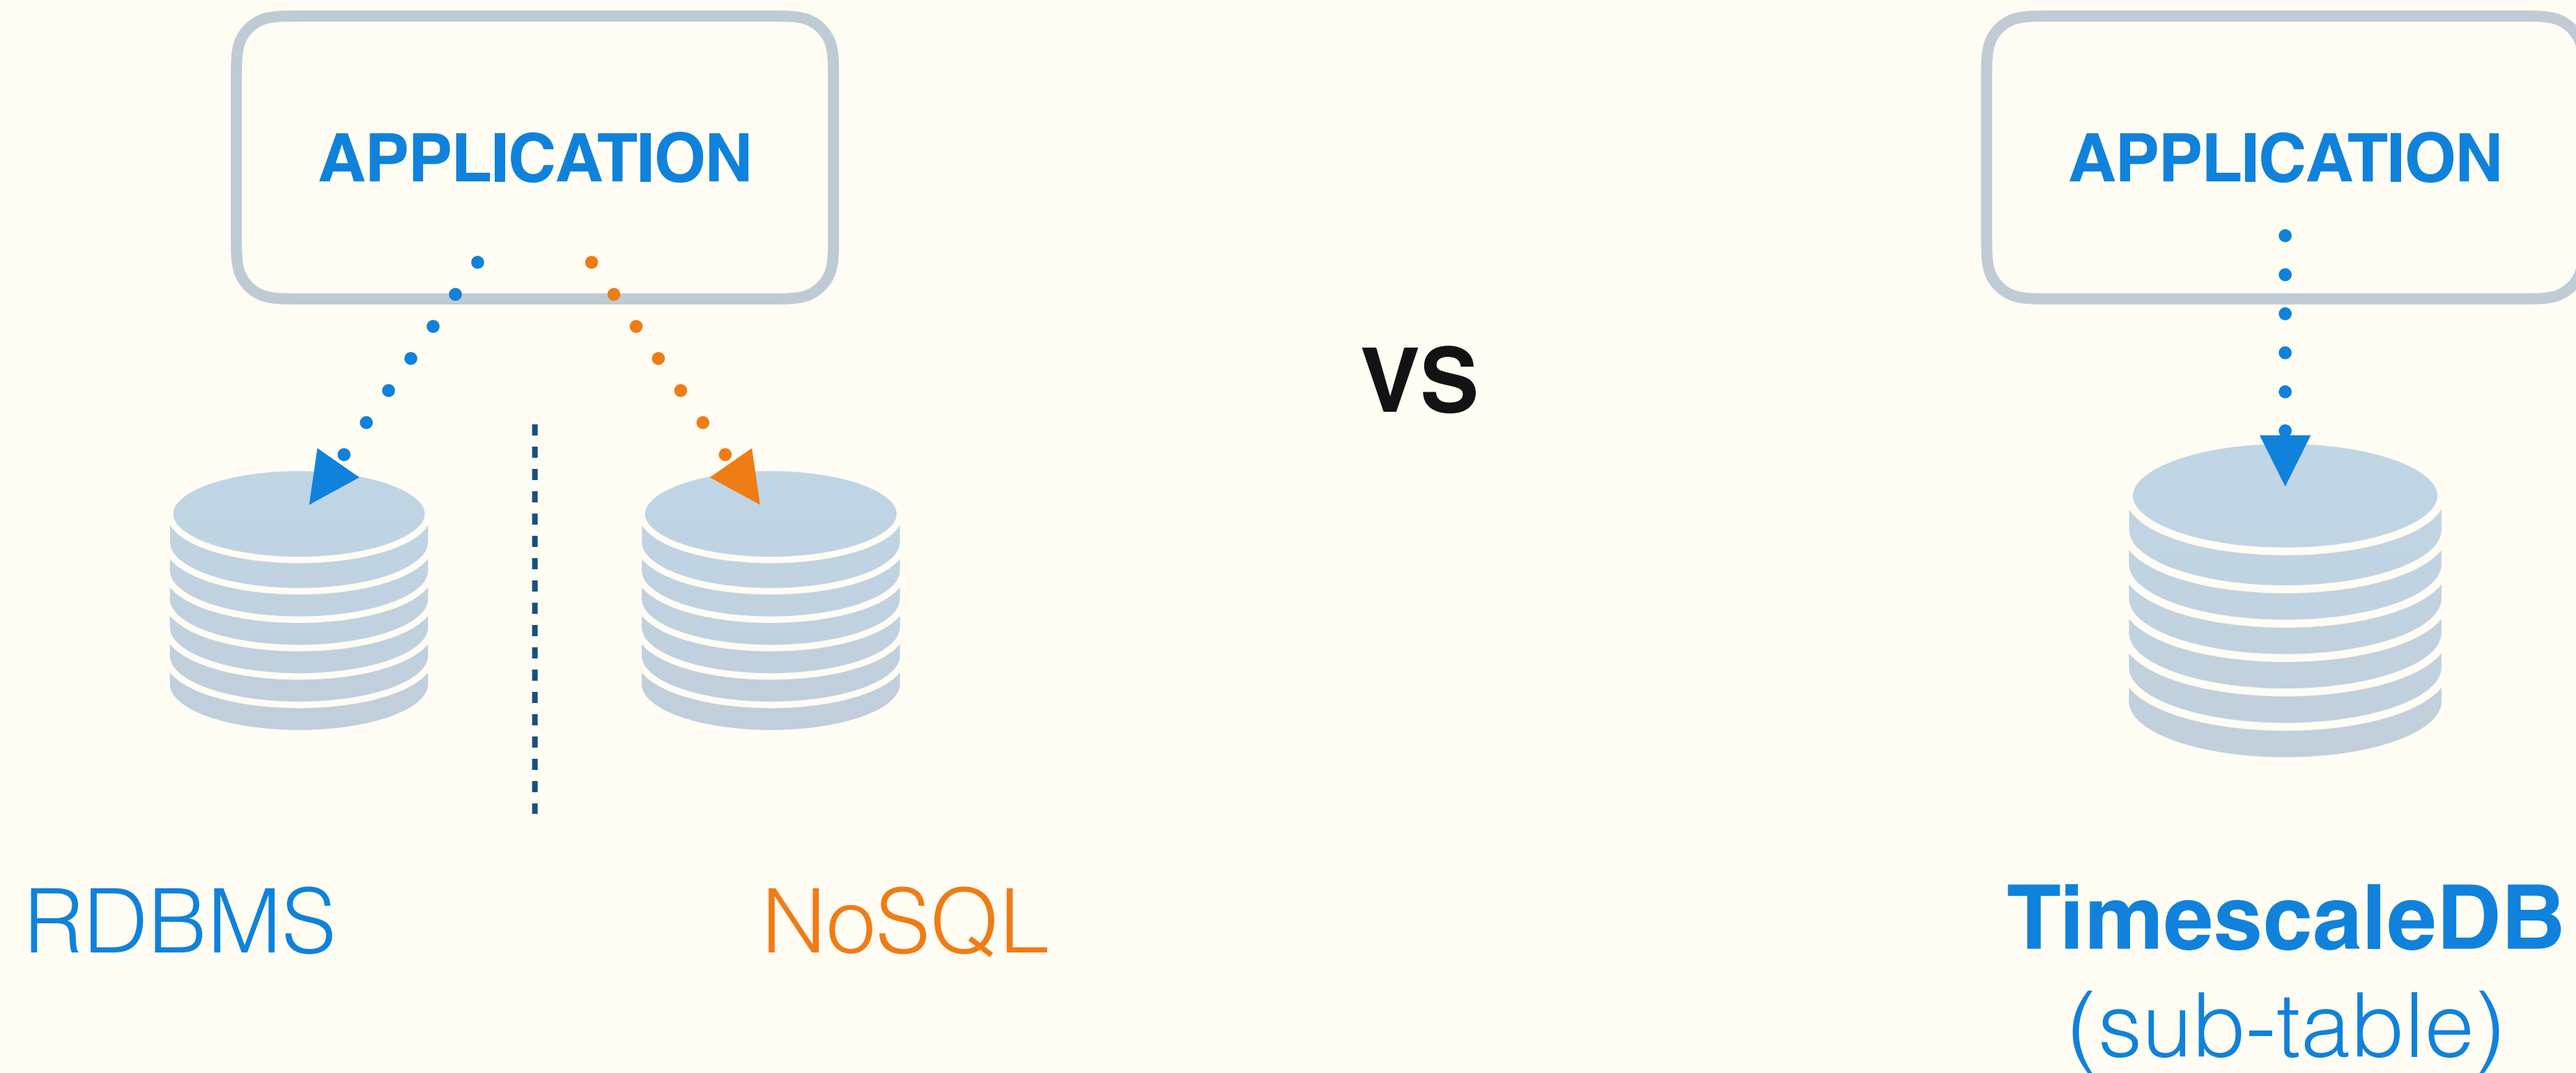


primary key metadata

container_id	type	cargo
ghi789	big	food
def456	medium	furniture
abc123	small	toys



Simplify your stack



High-Level Differences from **Plain** **Postgres**

- Insert performance
- Automatic partitioning
- Easier management
- Functions for time-based analysis
- Time-aware query optimizations
- Much faster deletes



High-Level Differences from **NoSQL**

- Secondary indexes
- Transactions
- Lower memory requirements
- No high cardinality problems
- Un-siloed data (JOINS!)
- SQL



We're **open-source!**

Apache 2.0 license



github.com/timescale/timescaledb





Using PostgreSQL, Prometheus & Grafana for Storing, Analyzing and Visualizing Metrics

Erik Nordström, PhD

Core Database Engineer

hello@timescale.com · github.com/timescale