




quasardb

MiFID II : could a smart storage help you?

November 7th, 2016

STAC NYC

Helicopter view of “the challenges”



We made storage smarter!

We took the right amount of database technology and embedded that directly into a high performance software defined storage.

Benefits

- Improved performance thanks to direct data access
- Reduced TCO (Less technologies, less layers)
- More flexibility: storage technology evolves independently
- Reduced dependence on database technologies: basic use cases can be done directly with quasardb



Quasardb: what does “smart” storage mean?

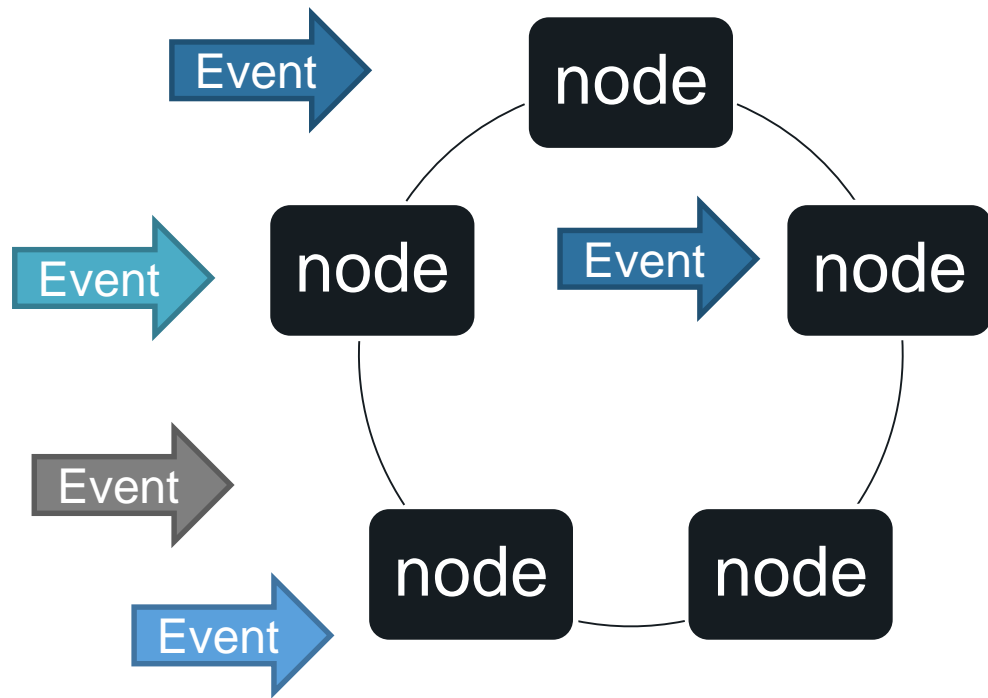
Quasardb is a **distributed, transactional key-value store** written in C++ 14.

Optimized storage for time series and other data types.

Advanced lookup based on name, tag, content...

High resolution timestamping.

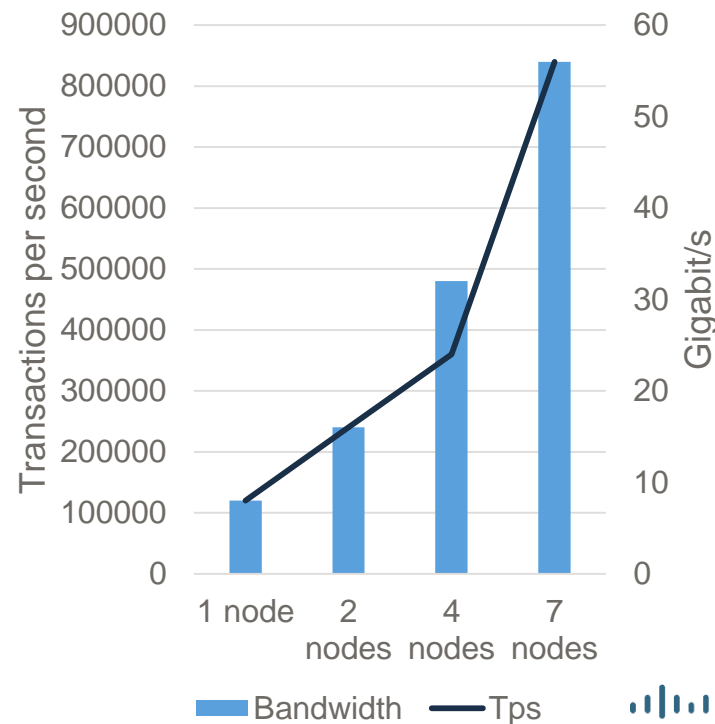
Rich API for **direct and efficient data access.**



Cisco bandwidth benchmark (not a STAC benchmark)

Loading from a quasardb cluster to an in-memory analytics server

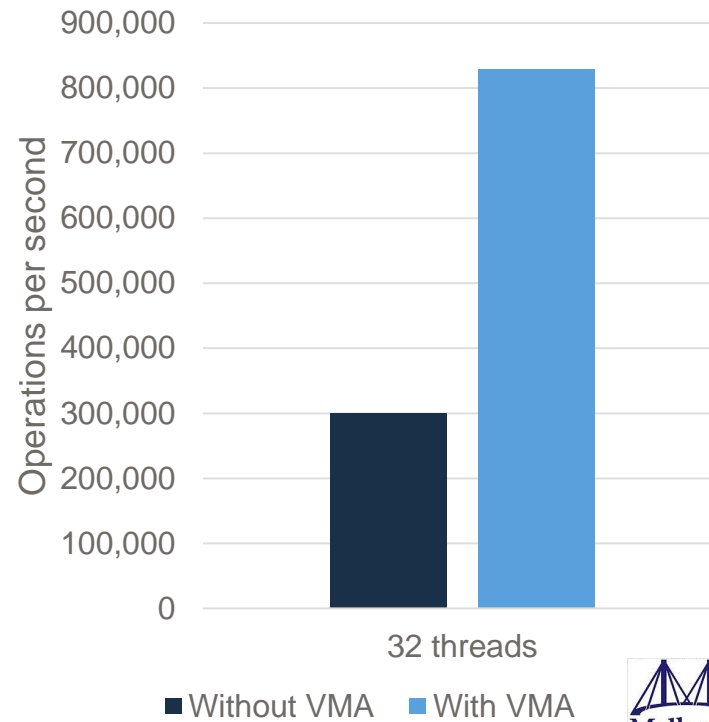
- In-memory analytics server: Cisco UCS C240M4 - 2 Intel E5-2697 v3 - 256GB – 2 x 40 GB VIC 1385
- quasardb cluster: Cisco UCS C240M4 server - 2 Intel E5-2697 v3 - 256GB - 2 x 10GB VIC - 24 1TB DD running Linux
- Switch: Cisco Nexus 9372 PQ (2 40GB connections to the simulation server – 8 10GB connections to two fabric interconnects)
- Fabric interconnect: 2 x Cisco UCS 6248 equipped with the UCS manager software



Mellanox latency benchmark (not a STAC benchmark)

Requesting 1 byte of data from one machine to the other

- Each server: 2 x Intel Xeon CPU E5-2697 - 128GB RAM
- Mellanox ConnectX-4 100 Gbit



References and partners





quasardb

WHATEVER THE AMOUNT OF DATA
WE CAN IMAGINE
THERE WILL BE MORE

WWW.QUASARDB.NET
SALES@QUASARDB.NET