



STS Database

Seismic Time Series Database

New Generation Time Series Database for Seismic Data

KEY FEATURES

- Clustered
- Linearly scalable
- Fault tolerant
- Embedded real time data backups
- Earthworm & Seiscomp compatible

KEY BENEFITS

- All your data in a single database
- No data sharding
- High availability
- No single point of failure
- High storage density - miniseed data
- Virtually no maintenance needed

STS Database is specifically designed to store very large seismic data sets in a single, clustered and auto-replicated NoSQL database.

STS represents a major leap forward from traditional relational databases such as Winston or file based storage mechanisms, by embracing new technologies and a modern design.

Its clustered nature along with a no single point of failure architecture, maximizes availability while at the same time improves performance linearly with each STS server added to the cluster.

Embedded real time data backups means all data is automatically replicated upon arrival in at least two different servers.

An STS cluster can tolerate server failures with no service downtime as long as failures doesn't occur in servers containing all copies of the same piece of data. In a cluster of N servers, the maximum tolerable limit for server failures is $N / 2$.

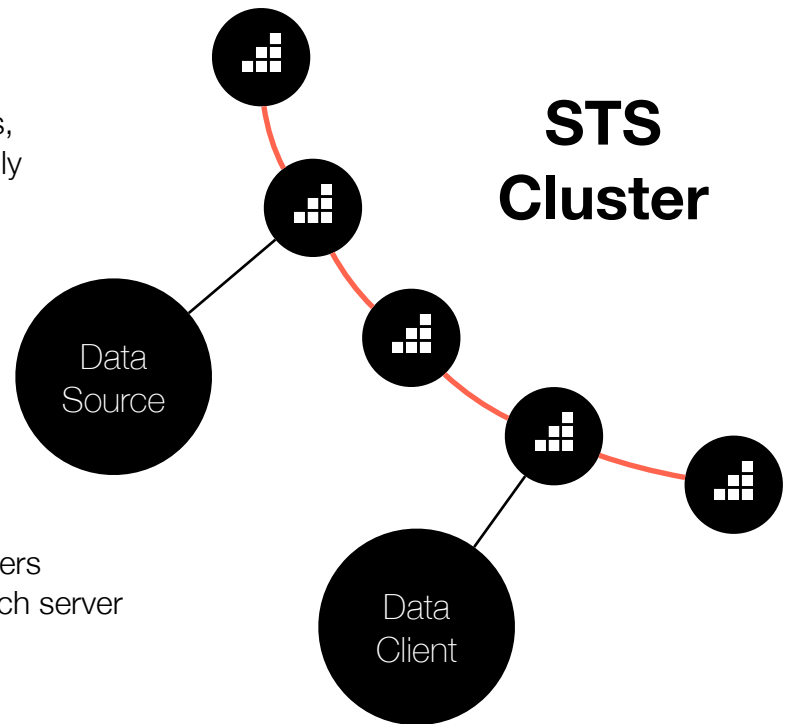


Cluster Advantage

STS servers coordinate data ingestion. Only one server per cluster connects to each individual data source. If an ingestion server fails, another server within the cluster will automatically take its place and connects to the orphaned data source.


STS servers in a cluster use a topology that resembles a ring. This ring is auto-healing, which means that in case of server failure, neighboring servers will automatically close the gap with no overall service downtime.

Data clients can connect to any of the STS servers in the cluster, all data is available no matter which server they connect to.



Data Storage Strategy

Data is distributed and replicated among all servers within the cluster at insertion time. When a new STS server is added, it will gracefully join the cluster and start storing its share of the data. This process is called storage auto-rebalance and is automatic.



Writing and reading from an STS cluster is extremely fast. This is achieved by writing data in disk sequentially, this way when writing or reading, the amount of disk seeks needed are kept to a minimum.



SPECIFICATIONS

Clustering

Topology	Ring (auto-healing)
Elasticity	Hot server insertion and removal

Management & Configuration

Management	Web dashboard
Configuration files	None
Console access	None

Data Ingestion

Real Time	EW, SeedLink and BRP clients
Bulk	From miniseed files

Data Egress

Protocol	Winston compatible
Output format	Winston compatible
Note	Other protocols and formats available upon request

System Requirements

Processor	4 cores @ 2GHz minimum
RAM	16 GB minimum
Disks	2 disks 7200 RPM Disk 1 at least 32 GB Disk 2 > 1 TB recommended
OS	Debian 10 GNU / Linux

Licensing

Container	1 license per container
Server leasing	1 license per physical server

Support & Documentation

Support	Via email Premium options available
Documentation	Available online