

The Trident Historizer is a time-series data storage and aggregation module designed for high-performance and scalability. It is capable of reading and writing millions of data points per second and provides a built-in health dashboard, direct query tool, REST API, and live tag values for easy monitoring of the system. The system is composed of the following components:

1. **Database** The Trident Historizer uses a NOSQL database cluster for storing time-series data. The database cluster is designed to be masterless, highly available, and scalable. It provides replication and redundancy for data durability and fault tolerance.
2. **Direct Query Tool** The direct query tool provides a simple interface for querying time-series data. It is designed to be fast and efficient, with support for powerful data aggregations and scaling. The tool can be used for ad-hoc trending and data export.
3. **Health Dashboard** The health dashboard provides a detailed overview of the entire cluster, including the Triton Edge Messengers and their performance metrics. The dashboard provides information on the connection state, store and forward buffer size, number of tags being polled per second, number of messages being sent per second, CPU utilization, and up-time. The dashboard also provides consolidated performance data for the Triton Edge Messengers tag subscription rate and the health of the brokers and clustered storage nodes.
4. **Duplication** The duplication component allows the system to connect to other historian databases to pull data into the Neptune Cluster system. This improves the speed of accessing data and offloads burden from legacy systems. The duplication component is designed to be highly available and scalable.
5. **REST API** The REST API provides a web-based interface for connecting Neptune Cluster to third-party applications. The API supports querying time-series data, ad-hoc trending, and data export. The API is designed to be scalable and secure, with support for authentication and authorization.
6. **Tag Values** The tag values component provides a dashboard showing live tag values for validation and troubleshooting. The component is designed to be scalable and efficient, with support for filtering and sorting.

Component	Feature	Description
Database	Storage	NOSQL database cluster designed for scalability and fault tolerance
Direct Query Tool	Querying	Fast and efficient querying interface with support for powerful data aggregations and scaling
Health Dashboard	Monitoring	Detailed overview of the entire cluster, including Triton Edge Messengers and performance metrics
Duplication	Integration	Connect to other historian databases to pull data into Neptune Cluster system
REST API	Integration	Web-based interface for connecting Neptune Cluster to third-party applications
Tag Values	Monitoring	Dashboard showing live tag values with support for filtering and sorting