PingThings

The data management, analytics, and AI platform for time series measurements at scale.

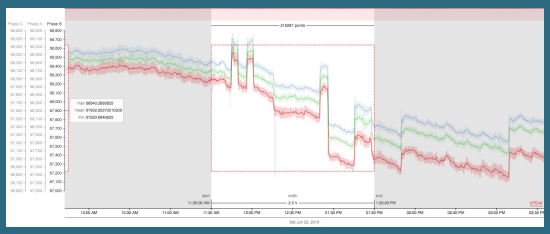

IT Features & Benefits

1 – Incredibly fast and horizontally scalable

- Horizontally scale to petabytes from millions of sensors.
- Read/write 15M/30M points per second per node.
- 10-1000x faster than the competition.

2 – Designed for real-world, messy sensor data

- Your code is versioned, why not your sensor data and analytic results? All data (and analytics) are versioned.
- Data quality is assessed on ingest with results stored for downstream consumption and reporting.
- Natively handles data dropouts and back filling data.

3 – Reliable, resilient, and cost effective

- Never lose data again; automatic data replication and self-healing for durability and high availability.
- Optional periodic backups for even more data security.
- Scale storage and compute independently for cost effective growth.
- State-of-the-art data compression.

4 – Flexible deployment options

- Fully containerized, orchestrated using Kubernetes.

amazon web services | Microsoft Azure — PUBLIC CLOUD

PRIVATE CLOUD

5 – Centralized, at the edge, or in the fog? Have your cake and eat it too.

- Fractal architecture allows you to run the platform on the edge or centralized or both.
- Run the same algorithms everywhere.
- Secure, hyper-compressed data streams make the most efficient use of potentially limited comms bandwidth.

User Features & Benefits

1 – Handles any and all time series/telemetry data

- Nanosecond time precision supports GHz sample rates.
- Ingests both real-time streaming data and massive historical archives.
- Handles multi-resolution time series natively.

2 – Interactively explore massive data sets

- Think Google Maps but for petabytes of time series.
- Zoom easily across orders of magnitude of time to examine individual events.
- Interactive visualization with <200ms queries at any scale.

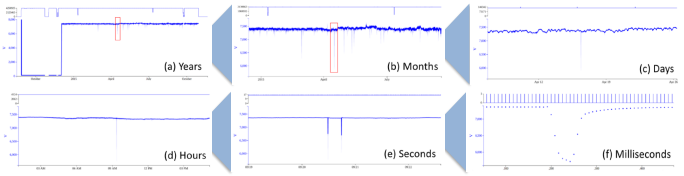

3 – Easily prototype new uses

- Robust support for Jupyter Notebooks for literate coding.
- Script and program in most languages with robust APIs.

R | python | golang | julia

4 – Traverse the data analytics pipeline quickly

- Move from prototype to production deployment with new analytics in hours not months or years.
- Purpose built for both ad-hoc and real-time analytics.
- Deploy analytics to real-time with DISTIL.
- Uses best-of-breed machine learning and deep learning to enable AI out of the box.

RAY | TensorFlow | APACHE Spark

Application Areas

Industrial IoT	Defense	Transmission Systems	Smart Cities
Industry 4.0	Sonar/Acoustics	Distribution Systems	Smart Countries
Precision Manufacturing	Cybersecurity	DER and Renewables	Smart Buildings
Semiconductors	Intelligence	Oil and Gas	Autonomous Vehicles
Telecommunications/RF	Aerospace/space	Geophysics and Seismology	Biomed/Health Sciences

(202) 991-0887 | info@pingthings.io | www.pingthings.io

The Technology Behind the Platform

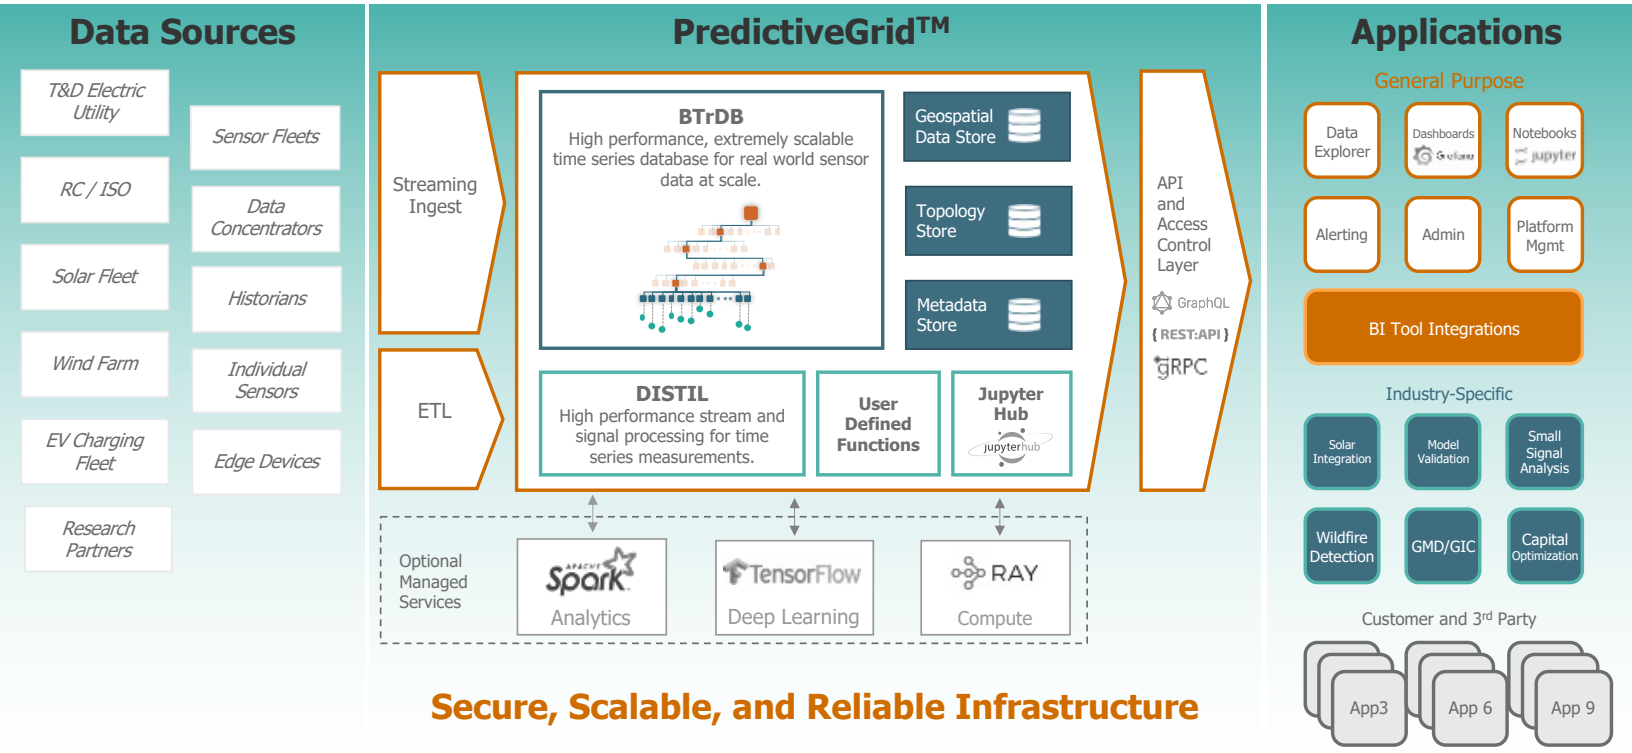

Secure, Scalable, and Reliable Infrastructure

The platform shown above can be decomposed into several functional areas. Moving from left to right in the diagram, it supports the ingestion of both streaming and historical data archives in a wide range of formats at scale via the ingest engine. Data is ingested into a database specifically designed for dense time series sensor data, the Berkeley Tree Database (BTrDB) whose development was funded by the ARPA-E Micro Synchrophasors for the Distribution System project. Further, the platform contains a distributed analytics and computational framework designed to operate across time series in parallel, executing significantly faster than real-time to handle both real-time and historical analyses and the training of machine learning and deep learning algorithms. The platform provides numerous APIs that provide not only a direct connection for web applications including a data explorer, dashboards, and Jupyter Notebooks for ad-hoc analytics but also to utility planning and operations software, allowing for the seamless integration of highly novel algorithms with the real world. The sections below will describe each component in detail.

The Berkeley Tree Data Structure

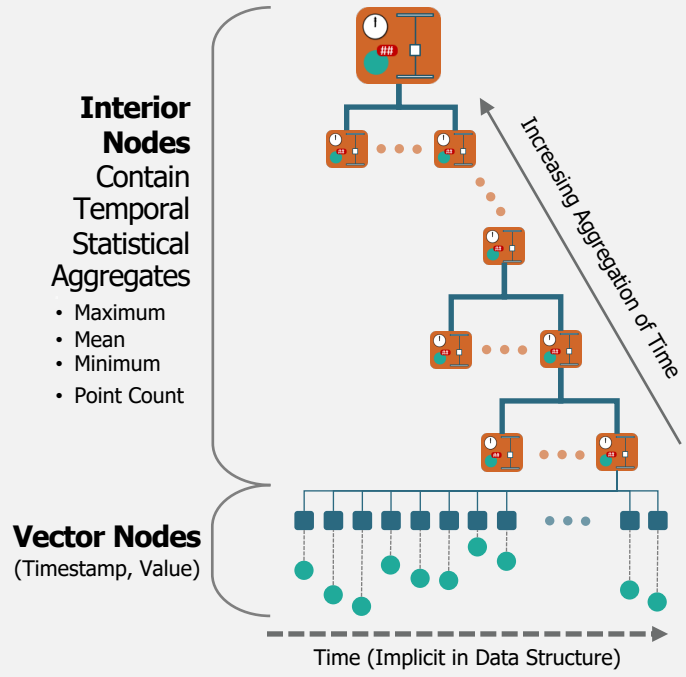

Graphic representation of the Berkeley Tree. The time-based index is implicit in the data structure itself requiring no extra storage space. The interior nodes provide near instantaneous access to time-based rollups of data. Further, the structure can easily store sensor data captured at different sampling rates or sensors that changes sampling rates.

About PingThings, Inc.

PingThings, Inc., founded in 2014, is the only startup company to ever receive seed funding from GE Ventures. PingThings' platform is an advanced sensor AI platform that is capable of ingesting, storing, accessing, visualizing, analyzing, and learning from high definition sensor data at web scale. PingThings has been awarded grants by the National Science Foundation, ARPA-E, and the Department of Energy.

1220 S Street
Suite 150
Sacramento, CA 95811

info@pingthings.io | www.pingthings.io