Ground Systems Architecture Workshop 2024

Evolving Telemetry Storage with Search and Time-Series

Presented by Guy Pyrzak

Other authors: Peter di Pasquale, Matt Triviski, Rob Puncel,

Jet Propulsion Laboratory
California Institute of Technology

Current Telemetry Storage Challenges… More Data!

- Deep Space Missions are lasting longer and generating terabytes of data.
- Surface Missions are getting better relays and generating more data.
- Earth Missions generate petabytes of data!
- Complex ground processing is generating 10-fold more data.
- Testbeds are generating more data.

Current Telemetry Storage Challenges…
More Commercial Off the Shelf Tools!

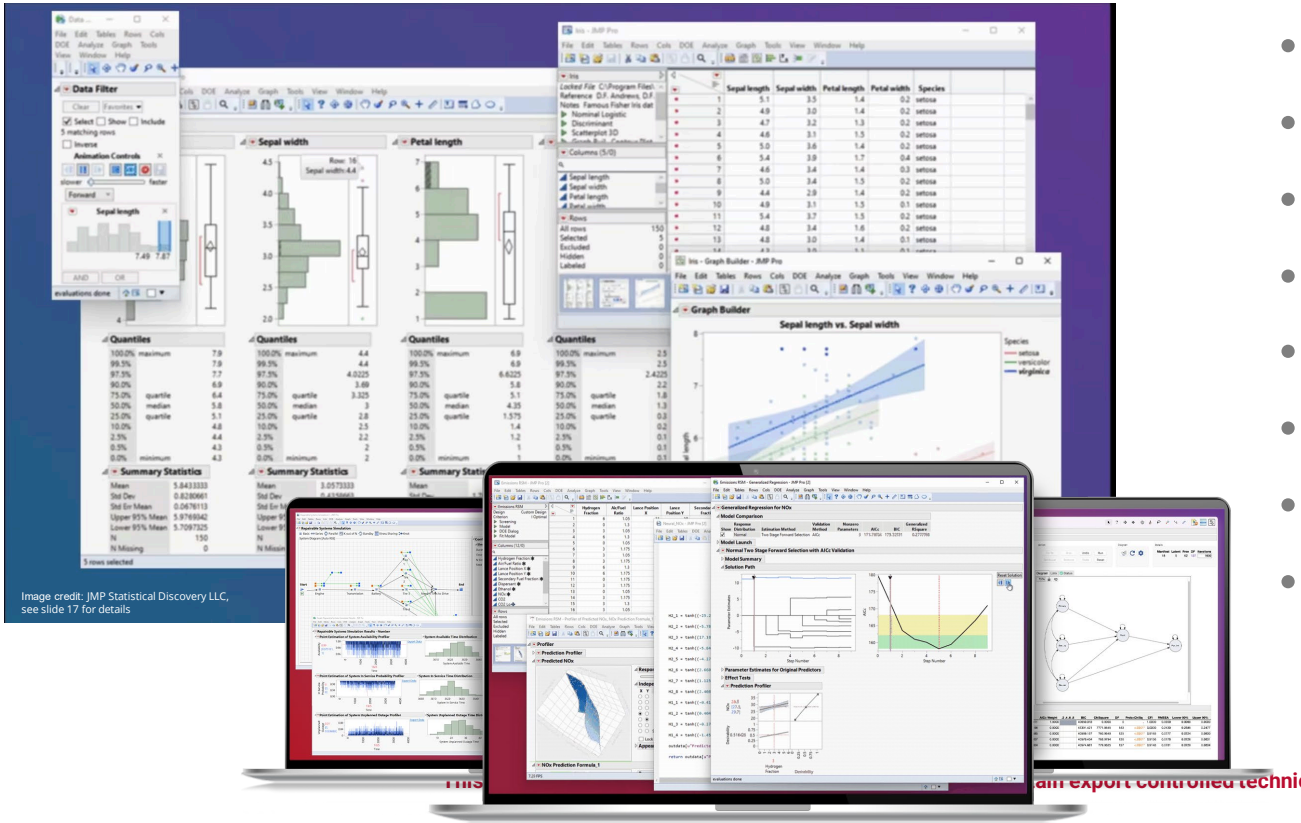

Image credit: JMP Statistical Discovery LLC, see slide 17 for details

- JMP
- Jupyter Notebooks
- Tableau
- Plotly
- PowerBI
- Grafana
- Kibana
- Straight Python

Current Telemetry Storage Challenges…
More Data Variety!

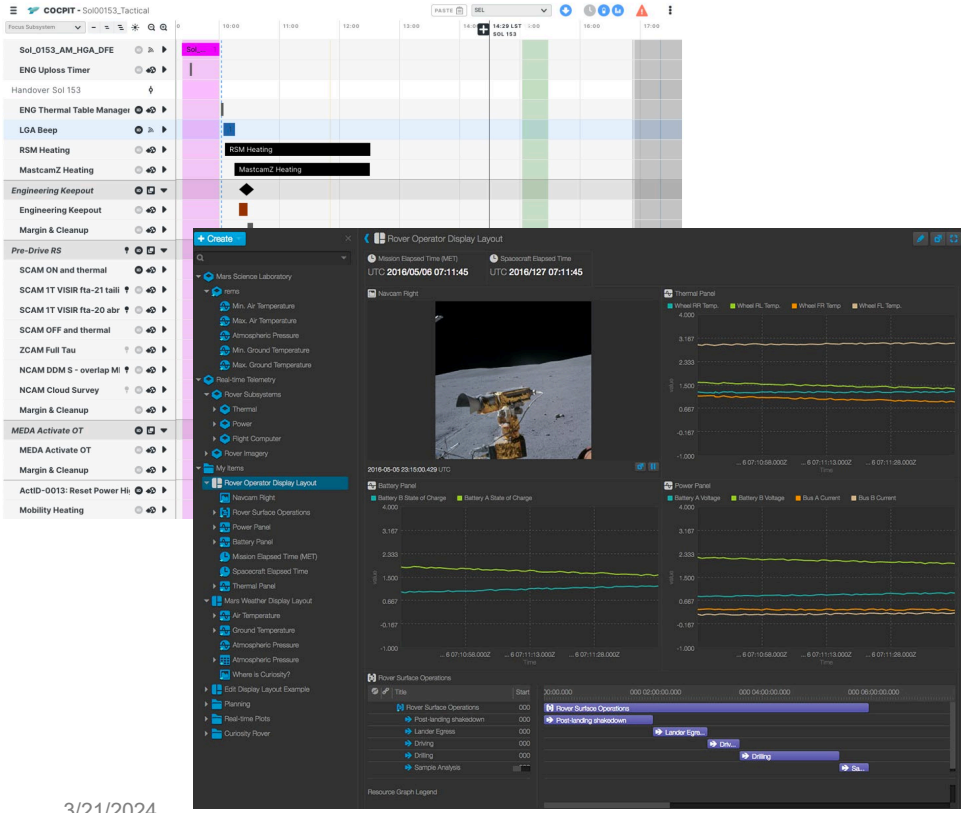

In addition to operations telemetry operators want…

- Communication/Link Data
- Operator Analysis Data
- Shift Schedule
- Uplink Data
- Testbed Data

Operators want more control and access to all the above and more!

Current Telemetry Storage Challenges…
More Security!

- Ensure all communication is encrypted
- Encryption at rest
- Zero Trust Security
- Authenticated access
- Multifactor Authentication
- Probably more!

Current Telemetry Storage Challenges… More Maintainable

Image credit: Christina Morillo, see slide 17 for details

- Missions are lasting longer
- Budgets for Phase E are getting tighter
- Need to be able to maintain performance with minimal staff.
- Ensure storage stays security compliant!
- Ensure scaling and archiving are cost effective

Where we started

- Many individual databases using MariaDB, Mongo, Redis, Postgres, or network file systems

Pros!

- Developers have ultimate freedom.
- Easy for developers to use on their own.
- No Co-ordination needed.

Challenges

- Hard to scale for volume & performance
- Lacks common security model.
- Operators can’t query directly OR need to learn LOTS of querying methods.
- No integrated archive strategy
- Self Hosted
- Complex schemas to maintain
- Security approach unclear

Elastic Search

Distributed Search/Analytics Engine with full text search and no schema.

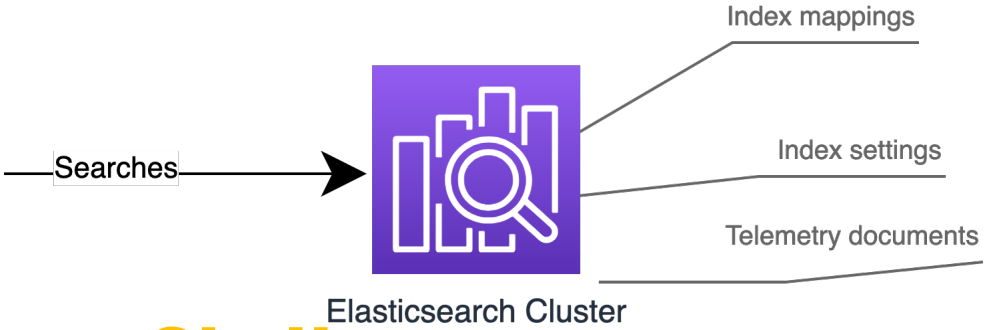

Pros!

- Very Performant & Scalable
- Used widely in industry
- Fed ramp approved hosted options
- Securable Web API
- Great Documentation & Examples
- Integrates with many COTS solutions.
- Supports MANY data

Challenges

- **Expensive to host lots of small datapoints.**
- Operators have direct access to database.
- Not all COTs analytics tools support interface natively.
- Not familiar to JPL developers.
- Profit Driven commercial ownership

Time Series DataBase

JPL In-House database created to address the deficiencies of Elastic Search.

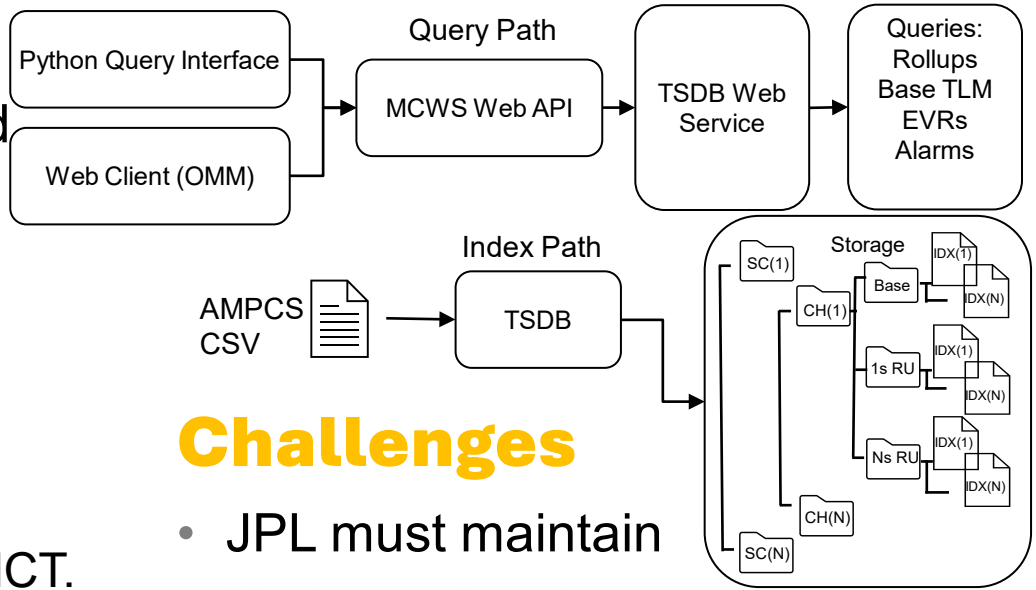

Pros!

- **Extremely cost effective &**
- **Highly Performant**
- Works well with AMMOS OpenMCT.
  - Supports rollups for any time ranges, as well as LAD, Alarm, and EVR queries.
- Securable Web API
- Design makes it easy to archive!

Challenges

- JPL must maintain
- No hosted options
- Tied to JPL's AMMOS Mission Control ecosystem.
  - COTs analytics tools don't support JPL's web API.

Influx DB

An emerging time-series database choice for small satellites and rovers that's prized for its simplicity

Pros!

- **Simple to set up and deploy**
- Has been used with OpenMCT, Grafana, and other tools.
- Is the default for some AMMOS open-source tools (AMMOS Instrument Toolkit)

Challenges

- Typical deployments are custom for each mission, so maturity of the deployments and capabilities vary dramatically.

Timescale Database

Open Source time series database built on Postgres

Pros!

- Super performant
- Much more cost effective (we think)
- Highly scalable
- Hosted Postgress options
- Postgres is familiar to developers
- Works with most COTS tools because it's SQL under the covers
- Open Source!
- Full text search indexing

Challenges

- No web API
- TBD security model
- Operations team are not familiar with SQL queries.

Chillax

In-House add on to Time Series DB to plug up any

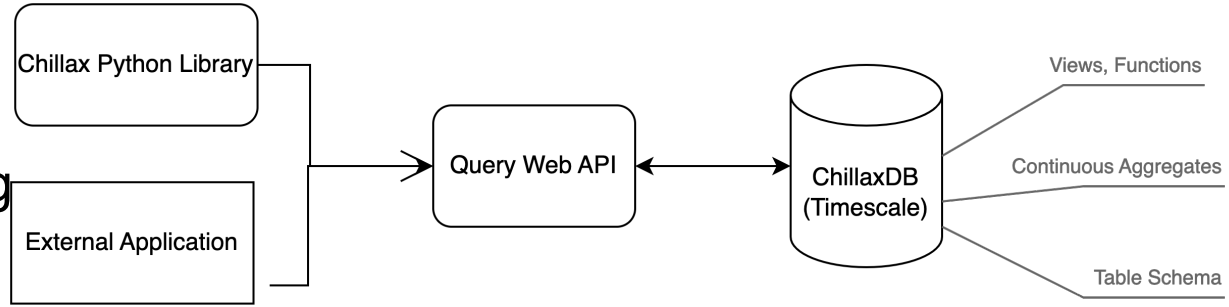

Pros!

- All the Pros of Time Series DB plus…
- Securable Web API leveraging GraphQL
- Python Library design for Operations team
- Adds-ons allow for even more customization.
- Full text search options being explored

Challenges

- In house web API and software maintained by JPL.
- Tuning and understanding full text search capabilities
- Queries not by time partition

In Operations…

Elastic Search

- MSL
- SMAP
- Mars 2020
- Ingenuity

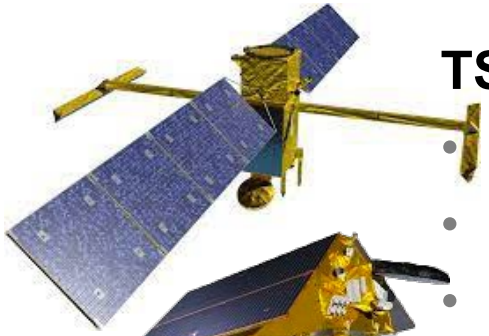

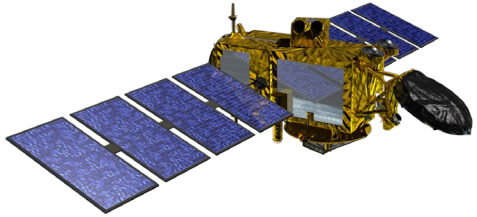

TSDB

- SMAP
- SWOT
- Sentinel6A
- CAL
- NISAR
- Jason3

Chillax

- Psyche
- Europa Clipper (soon)

Influx

- SunRISE
- CADRE
- DSOC

Chillax Current Status

- Development ongoing for features supporting external integrations, simulated and testbed data, custom extensions
- In use for Psyche operations by telemetry viewing applications
- Initial deployments on Europa Clipper

Chillax Work ahead

Potential Impact

- 100s of users daily
- Future missions
- Consolidate existing with existing TSDB and AMPCS software

Roadmap

- Multiple time systems with a single datastore
- Federated query across mulitple venues / environments
- Additional data types

Maturity

- Managing data for balancing cost, performance, and completeness especially over long time spans

Conclusions

5 very different approaches to providing centralized telemetry searching solution to missions, each with a variety of trade offs.

Clearly this need will increase as missions gather more data and operator's data analytics literacy increases

For more information on TSDB contact Peter di Pasquale:

peter.di.pasquale@jpl.nasa.gov

For more information on Chillax contact Rob Puncel:

robert.puncel@jpl.nasa.gov

Image Credits

- Slide 3: JMP Screenshots
  - *Shown with permission from JMP Statistical Discovery LLC.*
- Slide 6: **Christina Morillo**
  - https://www.pexels.com/photo/software-engineer-standing-beside-server-racks-1181354/
- Slide 5: TheDigitalArtist
  - https://pixabay.com/photos/cyber-security-internet-network-4610993/

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov