

Neutron OS Data Platform

Proposed Superset Dashboard Scenarios

Date: January 14, 2026

Status: DRAFT - For Review

Reviewer: Nick Luciano

Introduction

This document outlines proposed Superset dashboard scenarios for the Neutron OS data platform. These scenarios will drive the design of our data lakehouse architecture using a test-driven approach: we define what users need to see first, then build the data pipelines to support those visualizations.

Please review each scenario and provide feedback on:

- Questions/metrics that are missing or should be prioritized
- Charts or visualizations that would be most valuable
- Filters needed for practical use
- Data sources or time ranges to consider
- Priority order for implementation

Scenario 1: Reactor Operations Dashboard

Priority: HIGH (First to implement)

User Story

As a reactor operator, facility manager, or researcher, I want to see the current and recent state of the reactor so that I can monitor operations, identify trends, and identify anomalies.

Questions This Dashboard Answers

- What is the current reactor power and how has it changed today?
- What are the current control rod positions?
- What are the fuel and pool water temperatures?
- Are there any anomalies or unexpected readings?
- What was the power history over any time span?

Proposed Charts

Chart Name	Type	Description	Priority
Power Timeline	Line chart	Real-time power (kW) over selected time range	P0
Rod Positions	Multi-line or bar	Current position of Tran, Shim1, Shim2, Reg rods	P0
Temperature Gauges	Gauge/KPI cards	Current fuel temp and water temp with thresholds	P0
Power Distribution	Histogram	Distribution of power levels over time period	P1
Detector Signals	Line chart	NM, NPP, NP signals (for advanced users)	P2

Proposed Filters

Filter	Type	Default	Notes
Date Range	Date picker	Last 24 hours	Quick presets: Today, Yesterday, Last 7 days
Power Threshold	Slider	> 0 kW	Filter out zero-power periods
Auto-refresh	Toggle	On (5 min)	For real-time monitoring

Data Sources

- serial_data/*.csv - Reactor time-series (power, temps, rod positions) or SQL database
- Refresh: Near real-time (aim for 5-minute latency)

Commented [LN1]: Yes, but currently our data is uploaded nightly to Box and LS6. So, only the recent state (yesterday and before) is currently possible. To do present state, we would need to be streaming live data from NETL to... somewhere. There is some sensitivity with this, in that the public should generally not know when the reactor is at power. Bill has said that the day after is fine. BUT, Superset requires a login, so if we are only making the data viewable through Superset with a login, we should be fine. We just currently do not have a way to stream live data. We need one.

Commented [BB2R1]: I'll build streaming data with a full access control system. Let's shoot for this.

Commented [BB3]: So, no 'react to'?

Commented [BB4R3]: @Luciano, Nick

Questions for Nick

What time resolution is needed? (per-second, per-minute, hourly?)

The raw data at its resolution should be available. Superset does automatic data decimation so that data resolution is reduced as the time scale increases.

Which metrics are most critical for at-a-glance monitoring?

Power, Control Rod positions, fuel and coolant temperatures

Should there be alert thresholds highlighted on charts?

Probably not needed

What historical range is typically needed? (hours, days, weeks?)

Depends, but probably the last few days

Scenario 2: Reactor Performance Analytics

Priority: HIGH

User Story

As a researcher or operator, I want to analyze reactor performance over time by correlating power output, xenon poisoning, fuel burnup, and control rod positions so that I can understand operational patterns, optimize startup procedures, and predict reactor behavior generally.

Questions This Dashboard Answers

- How does estimated xenon concentration correlate with power history?
- What is the current excess reactivity given fuel burnup and xenon state?
- How have control rod positions changed relative to power demand?
- Which fuel elements have the highest burnup?
- What is the typical startup time from cold critical to full power?

Proposed Charts

Chart Name	Type	Description	Priority
Power & Xenon Timeline	Dual-axis line	Power (kW) and Xe-135 concentration over time	P0
Rod Position vs Power	Scatter plot	Correlation between power and avg rod height, colored by rod	P0
Fuel Burnup Heatmap	Heatmap	Core layout showing U-235 burnup by position (hexagonal)	P1
Excess Reactivity Trend	Line chart	Calculated excess reactivity over time	P1
Temperature Correlation	Scatter plot	Fuel temp vs water temp, colored by power	P2
Startup Time Distribution	Histogram	Time to reach full power from cold critical	P2
Daily Energy Production	Bar chart	Integrated energy (MWh) by day	P1

Proposed Filters

Filter	Type	Default	Notes
Date Range	Date picker	Last 7 days	
Core Configuration	Dropdown	Latest	Select BOC/EOC configuration
Power Threshold	Slider	> 0 kW	Filter operational

Commented [LN5]: We don't (and can't) actually measure xenon concentration. We can compute a value using models, but we have no validation if those values are correct. We can (and do) measure the critical rod height in the morning. That measurement is a reflection of the xenon concentration.

Commented [BB6R5]: Added 'estimated' qualifier - or should I remove this line altogether?

Commented [LN7]: This is kinda complicated. We could estimate with models since the start of the last fuel loading.

Commented [BB8R7]: Okay - keeping as aspirational for now.

Commented [LN9]: This would be great. We can use model outputs for this.

Commented [BB10R9]: Excellent!

			periods only
Rod Selection	Multi-select	All	Filter specific rods

Data Sources

- serial_data/*.csv - Reactor time-series
- Xe_burnup_2025.csv - Xenon/Iodine dynamics (433K+ rows)
- static/core/*.csv - Core configurations with burnup
- rho_vs_T.csv - Reactivity vs temperature lookup

Questions for Nick

Is the Xenon correlation the most valuable insight here?

That depends. We can't measure xenon directly. We can measure critical rod heights which correlate with xenon. It would be valuable to have a plot like that. Not sure if it is the most valuable plot.

Should MPACT shadow predictions be overlaid on measured data?

Yes. We want plots that display both measured and modeled data together, but we always want to make sure we know what is measured and what is modeled. I am currently keeping them on different tables and they have different nomenclatures.

What burnup thresholds would trigger attention? (for heatmap colors)

Not sure. It's more about the power / burnup distribution than any given threshold.

How is "startup time" currently defined/measured?

Sam Queralt is tagging data. He has algorithms he is using to do that.

Commented [TJ11]: We can ascertain a baseline fuel "burn up" amount when starting up from a long period of Shut down, e.g. running the first day after a long s/d like the winter break. Recently we saw a difference in excess reactivity of about -49 cents (down from 612.93 cents at last core change) after the break. We attribute this to Burn Up as Xenon is negligible here. The following day, after running at 950kW the previous day, the decrease in excess reactivity (less than -49 cents) is attributed to fission products / isotopics. Similar to Nick, I am unsure how valuable this plot would be.

Scenario 3: Elog Activity Summary

Priority: MEDIUM

User Story

As a facility manager or regulatory inspector, I want to see a summary of operations log activity so that I can understand operational patterns, verify compliance, and prepare for audits. As a researcher, I want to understand how the reactor was operating and make sure key data is captured digitally rather than by handwritten logs. Also, as a researcher I want a separate log for digital twin activities that run alongside operations, but do not belong in the operations log.

Questions This Dashboard Answers

- How many log entries were created per day/week/month?
- Which operators have logged the most entries?
- What types of operations are most frequently logged?
- Are there gaps in logging that need investigation?
- What is the run history over a given period?

Proposed Charts

Chart Name	Type	Description	Priority
Entries Per Day	Bar chart	Count of log entries by day	P0
Entries by Operator	Pie/donut	Distribution of entries across operators	P1
Run Timeline	Gantt/timeline	Visual timeline of reactor runs	P1
Entry Categories	Bar chart	Breakdown by entry type (startup, shutdown, observation)	P2
Logging Gaps	Calendar heatmap	Highlight days with few/no entries	P2

Proposed Filters

Filter	Type	Default	Notes
Date Range	Date picker	Last 30 days	
Operator	Multi-select	All	
Run Number	Dropdown	All	
Keywords	Text search		Full-text search across entries

Data Sources

- Elog system (currently JSON files, migrating to immutable blockchain)
- Note: This dashboard depends on elog system development (see Elog PRD)

Questions for Nick

What categories/tags should Elog entries have?

Need to ask Jim and Rod in NETL Ops.

What constitutes a "gap" that should be flagged?

Need to ask Jim and Rod in NETL Ops.

Should this include export-to-PDF for audit preparation?

Need to ask Jim and Rod in NETL Ops.

Commented [TJ12]: I am not sure what insight could be achieved by counting the number of log entries, or having a chart of entries by operator. The console keeps track of operator hours (it even breaks rx operation down by MWh per operator), but we simply use this to track required the 4hr per quarter minimum requirement of rx operation for the operator requalification program. If the plan is to use the log to track operator log time, this could be useful.

Commented [TJ13]: Categories/Tags may have entries such as PnT sample sent/received, excess reactivity, sample reactivity, facility install / removal, sample dose rate at removal -yet this may be more useful in a sample log that is currently used at each sample insertion/removal location)

We have attempted to use the elog software, by first deciding on a method/layout. The digital Rx log should provide a clear record and provide access to the information for the DT model.

It appears most efficient if operations had the ability to manipulate this "free" software by to "play" with the fields/layout to determine the best (read: most useful) way of using it. Perhaps we could have a designer come to the NETL while operating and then design/massage the elog software into a usable form?

Commented [BB14R13]: Will look into how to provide a raw form for you to get your hands on.

Commented [TJ15]: presuming that you mean a "gap" to flag mandatory entries every :30 When operating the Rx operator notes the Rx power level and checks the status of the facility rad monitors, then logs this. A gap would mean that this :30 minute check was not performed when operating.

Commented [TJ16]: Export to PDF would work, but a simple text file for archive and future proof would also work. Once we decide on a standard format for the log (to include any changes desired) a text file could be exported for archive. This would be fine for audits and inspections if the Rx log software ever changed or was not accessible. If the normally used software was accessible, it could be used to view the logs. A periodic archive using "export to PDF" could also be set on a schedule. This item was brought up as a concern over ensuring that we had a way of archiving and maintaining the logs in the event of software failure or migration to different platform / log software.

Scenario 4: Experiment Tracking

Priority: MEDIUM

User Story

As a researcher or principal investigator, I want to track experiments from planning through completion so that I can manage research activities, correlate results with reactor conditions, and report on progress.

Questions This Dashboard Answers

- What experiments are currently planned, in progress, or completed?
- What reactor conditions were present during each experiment?
- How much beam time has each PI/project used?
- What is the backlog of experiments awaiting scheduling?
- Which experiments correlate with specific reactor runs?

Proposed Charts

Chart Name	Type	Description	Priority
Experiment Status Funnel	Funnel	Count by status: Planned → Scheduled → Running → Completed	P0
Experiment Calendar	Calendar	Scheduled experiments on calendar view	P1
Irradiation Hours by PI	Bar chart	Usage breakdown by principal investigator	P1
Experiment Timeline	Gantt	Timeline showing experiment duration and status	P2

Proposed Filters

Filter	Type	Default	Notes
Date Range	Date picker	Current quarter	
Status	Multi-select	All	Planned, Scheduled, Running, Completed, Cancelled
Principal Investigator	Dropdown	All	
Experiment Type	Multi-select	All	

Data Sources

- Experiment tracking system (to be developed)
- Correlation with reactor time-series for conditions during experiments

Commented [TJ17]: Note: we have a schedule of Authorized Experiments. These experiments are authorized by the ROC (Reactor Oversight Committee) after review through NETL Staff (Reactor Manager (RM)/HP/Director -Then ROC chair). Routine Experiments are performed as they are covered under an Authorized Experiment. Routine Experiments are approved by the RM or Senior Supervising Reactor Operator (SSRO) using a Request to Operate form completed by an experimenter / researcher. Typically a request to operate (perform a routine experiment) is permitted for the calendar year and forwarded (after annual review) to the next calendar year as needed. Each of these lists is available upon request. For planning we use a common Reactor Calendar to schedule experiments.

Reactor conditions (power and facility used / installed) are dependent on the experiment.

Beam Time for each PI (assuming you mean experiments involving use of the beam ports) is logged in the Rx Log. *note -this information may be categorized / tagged as "beam time" as you mentioned Tags / Categories above.

Your last two questions on backlog and experiment correlation may be answered with moving to a system of digital logs, digital "request to operate" and/or method of tying requested experiments with scheduled reactor runs.

Thanks for your work in thinking and moving forward with this. The experiment meta data you mention below would be very useful. We attempted this a while back using MS Access. Again, I believe it would be useful to have a modeler / designer come to the NETL and investigate and review our current system. Then begin.

Commented [BB18R17]: Glad to move the conversation forward. My intent is to gather, gather, gather and then propose a distilled set of problems/opportunities for us. If we agree on this, then we can explore the solution space in any number of ways, including having a designer propose a set of alternative concepts for us to evaluate. As we zero in on a solution concept, we'll generate some higher fidelity mocks that we can all dwell on before we build a beta for testing. This sounds like a lot but if executed well, it can go very quickly.

- Note: This depends on experiment management features (future development)

Questions for Khiloni Shah

[] How are experiments currently tracked? (spreadsheet, system, ad-hoc?)

TBD

What experiment metadata is most important?

(Nick): I would want something like:

- Sample Name (must be unique)
- Sample numeric ID assigned automatically
- Chemical Composition
- Isotopic Composition
- Density
- Mass
- Location of Irradiation (central thimble, lazy susan, etc.)
- Irradiation Facility (cadmium covered, etc)
- Datetime of insertion
- Datetime of removal
- Decay time after removal
- Count Live time
- Total counts
- Total activity
- Activity by isotope
- Measurement raw data (spectra or other)

We would want a user interface that allows some things to be prepopulated based on previous samples or based on the limited options (irradiation location, facility). We should work with NETL staff to get this right.

Should this integrate with a scheduling/calendar system?

It could. If it does, you want to make sure there was no conflicting data. The calendar is used to schedule time. It is not a reflection of what actually happened, only a reflection of what somebody intended to happen at a point in time. Knowing what actually happened is more useful.

Scenario 5: Audit Readiness

Priority: MEDIUM-HIGH

User Story

As a regulatory inspector or compliance officer, I want to verify the integrity of historical records and generate evidence packages so that I can conduct audits efficiently and with confidence.

Questions This Dashboard Answers

- Can I verify that records have not been tampered with?
- What audit events have occurred in a given period?
- What evidence packages are available for inspection?
- Are there any verification failures or anomalies?
- What is the complete audit trail for a specific record?

Proposed Charts

Chart Name	Type	Description	Priority
Audit Event Timeline	Timeline	All audit events (data changes, access, exports)	P0
Verification Status	KPI cards	Count of verified, pending, failed records	P0
Evidence Package Inventory	Table	List of available audit packages with download links	P1
Data Integrity Check	Status indicator	Visual confirmation of blockchain/Merkle verification	P1

Proposed Filters

Filter	Type	Default	Notes
Date Range	Date picker	Last 90 days	
Event Type	Multi-select	All	Create, Update, Access, Export, Verify
Record Type	Multi-select	All	Elog, Reactor Data, Experiment, etc.
Verification Status	Dropdown	All	Verified, Pending, Failed

Data Sources

- Hyperledger Fabric blockchain (audit events, Merkle proofs)
- Immudb (single-facility audit trail during development)

Commented [TJ19]: Tamper Proof of records has been mentioned, before. Can a person simply "change" a log to add or omit something that was an error. Nick seems to have this covered in that, a person may provide a supplemental comments, but the original entry would be locked as submitted.

Of course a text file (archive) could be modified easily... but a written log could simply be rewritten as well. Perhaps the elog software would provide a reasonable assurance that the log has not been tampered with (as evidenced by the fact that an entry can only be supplemented)

Audit records are maintained in the Rx control room (written / signed).

Unclear what you mean by evidence packages
Unclear on verification failures
Audits are performed iaw TS requirements (Section 6.2.4 a. through d.)

Commented [BB20R19]: This right?
<https://www.nrc.gov/docs/ML0703/ML070380197.pdf>

Commented [BB21R19]: A new capability I'm considering introducing is a blockchain-backed log. All entries would be immutable and verifiable by anyone.

- Note: This depends on blockchain infrastructure development

Questions for Jim and Rod at NETL Ops

What does a typical NRC inspection request look like?
(TBD Ask Jim and Rod in NETL Ops)

What format should evidence packages be in? (PDF, ZIP, others?)
(TBD Ask Jim and Rod in NETL Ops)

How far back do inspectors typically need to query?
(TBD Ask Jim and Rod in NETL Ops)

Commented [TJ22]: NRC inspects roughly half of the NETL required records of performance involving many aspects other than Rx ops (e.g security, HP, materials reports) every year, this places items for inspection on a two year periodicity (as to "how far back") Our inspector typically informs us via email his intent to inspect and the requested date.

Evidence Packages should include two years worth of information, but as mentioned above in another note - there is no one particular format. "evidence packages" could consist of electronic documents -this actually seems to be the preferred method currently.

Commented [BB23R22]: The ulterior motive with this is to design a 'compliance package' format and help NRC push a standard. The future system design would be such that compliance is continual and could be verified by auditors in a standard, efficient manner. For operators, maintaining compliance should be SOP. There should be minimal or zero extra effort needed to prepare for an audit.

Summary: Proposed Priority Order

Priority	Scenario	Rationale
2	Reactor Operations Dashboard	Most immediate value; validates core time-series pipeline
1	Reactor Performance Analytics	Combines multiple data sources; validates join logic
5	Audit Readiness	Critical for compliance; drives blockchain requirements
4	Elog Activity Summary	Depends on elog system development
3	Experiment Tracking	Depends on experiment management system

Next Steps

1. Nick reviews and provides feedback on this document
2. Prioritize and refine scenarios based on feedback
3. Define Gold table schemas for highest-priority scenario
4. Write dbt tests (test-driven development)
5. Build data pipeline to pass tests
6. Create Superset dashboard and iterate

Feedback Instructions

Please edit this document directly or provide comments. Key areas for feedback:

- Add/remove/modify questions each dashboard should answer
- Suggest additional charts or visualizations
- Identify missing filters or data sources
- Answer the "Questions for Nick" in each section
- Adjust priority order if needed