



RAISUN TECHNOLOGY SERVICES

RTS RESEARCH • WORKSHOP RECAP

From Pilots to Practice

NRC Artificial Intelligence Workshop #7 Recap

Navigating regulator transformation, industry innovation, and global collaboration

Workshop #7

NRC Data Science and AI Regulatory
Applications series

Sept 15

Hybrid public workshop, 2026

Rockville

NRC Auditorium, NRC Headquarters,
Maryland

Sam Rieck

Founder and Chief Nuclear Officer, RTS

Nuclear licensing leadership across advanced reactor developers
and the U.S. DOE. Presented at Workshop #7.

RTS Research

Workshop Recap Series

Synthesized from the official agenda and about 8 hours of notes.

AT A GLANCE

The Day in Brief

CENTRAL CONCLUSION

Nuclear AI is **rushing into the implementation phase**. The decisive questions are no longer whether AI can add value. They are where it belongs in a workflow, what evidence shows it is fit for purpose, and who keeps authority when the output matters.

Part 53

Latest nuclear AI application area

5 → 1

Months for a Duke aging-management review, as reported by Duke

10M

Relationships mapped in TerraPower's knowledge graph

KEY FINDINGS

01 Trust is a system property, not a model claim.

Speakers tied trust to curated data, explicit task boundaries, traceability, validation, monitoring, and accountable human decisions.

02 Production value is measurable.

Duke cut an aging-management review from roughly five months with ten engineers to about one month, while targeting 99% recall for critical records.

03 The industry is converging on bounded autonomy.

AI may retrieve, rank, optimize, or recommend inside a qualified envelope. Deterministic calculations, established analysis tools, or qualified professionals remain the decision anchors.

04 Data access is becoming the binding constraint.

Legacy analog assets, fragmented document systems, inconsistent metadata, and security boundaries increasingly decide what can be deployed.

05 Regulatory readiness is advancing through use.

NRC staff showed AI supporting inspection work and the AP1000 DCD review, while an OECD/NEA and CNSC tabletop tested an AI-enabled maintenance case and exposed gaps in data quality, governance, and uncertainty.

06 The next frontier is workflow design.

Strong programs start with the work, define what “good” means, build an evaluation set with subject-matter experts, and only then select or train a model.

CONTENTS

What's Inside

- 1 A Workshop About Operational Maturity
- 2 Human in Command, Bounded Autonomy, and Regulatory Evidence
- 3 The Real Infrastructure Is Information
- 4 What Deployments Are Teaching the Industry
- 5 AI Around Physics, Simulation, and Control
- 6 Part 53 as a Machine-Readable Decision Architecture
- 7 A Practical Agenda for Nuclear AI

Method and Sources

About the Author

References

HOW TO READ THIS RECAP

Sections 1 through 5 synthesize presentations and discussion from across the workshop. Section 6 summarizes the research RTS presented. Section 7 turns the day into a practical agenda. Statements are paraphrased unless marked as quotations.

1 A Workshop About Operational Maturity

The NRC’s seventh AI workshop showed a field moving past general interest and isolated demonstrations, and into real work.

Workshop #7, “Navigating the AI Landscape: Regulator Transformation, Industry Innovation, and Global Collaboration,” was a hybrid meeting held in the NRC Auditorium in Rockville. It connected three perspectives: the NRC’s internal AI maturity path, emerging regulatory-readiness work, and industry applications moving into real operations. The agenda ranged from inspection preparation and licensing information to digital overlays, cybersecurity, autonomous research-reactor control, knowledge graphs, procedure management, and agentic licensing support.

Speakers didn’t describe one universal nuclear AI stack. They presented patterns matched to different consequence levels and data environments: retrieval over controlled documents, machine learning for classification, physics-informed surrogates for core design exploration, knowledge graphs for traceability, non-invasive sensors for legacy assets, and localized compute for protected plant data.

1.1 The workshop’s practical arc

Stage	Workshop emphasis	What changed
1. Establish foundations	NRC maturity, inspection support, international collaboration	AI moved from awareness into agency and industry work processes.
2. Bound the risk	Automation bias, auditable envelopes, verification and validation	Discussion shifted from capability to evidence, controls, and consequences.
3. Deploy around constraints	Legacy data, cyber architecture, proprietary information	Architecture and data governance became part of the assurance case.
4. Prove value in work	Duke, TerraPower, Southern Nuclear, engineering tools	Teams reported production use, quantified gains, and adoption lessons.
5. Debate authority	Flash talks and open discussion	The closing question was how to delegate, bound, monitor, and recover authority.

A NOTABLE CHANGE IN TONE

Earlier nuclear AI conversations often centered on possible use cases. This workshop centered on operational lessons: what failed, what had to be measured, how users were involved, and which controls made deployment acceptable.

1.2 A risk-informed posture

The day consistently supported a graded approach. Low-consequence tasks such as discovery, summarization, and classification can tolerate different controls than safety-related calculation or

autonomous action. As consequence rises, expectations for provenance, validation, uncertainty treatment, independent confirmation, cybersecurity, and accountable approval should rise with it.

2 Human in Command, Bounded Autonomy, and Regulatory Evidence

A reviewer at the end of the process isn't the same as a human in command of it.

The workshop's most persistent debate concerned the "human in the loop." Participants agreed that placing a reviewer at the end of an AI-enabled process doesn't guarantee safety. A reviewer who sees the AI answer first can anchor on it, and a tool that is usually right teaches people that checking is rarely needed. The more useful question is where judgment and authority sit throughout the workflow.

2.1 Three control ideas converged

Control idea	Meaning in practice	Workshop evidence
Human in Command®	A qualified person or organization defines the objective, authorizes the use, owns the decision, and can stop or change the process.	RTS Part 53 remarks, Duke's lead AI scientist, Southern Nuclear's acceptance workflow.
Bounded autonomy	A model may act without step-by-step review only within a tested task envelope, with defined inputs, outputs, limits, and escalation points.	Closing discussion, CoreDesigner domain of applicability, autonomous-control testbeds.
Deterministic anchors	AI steers, retrieves, or optimizes while established calculations or one-way physical controls supply the authoritative result or boundary.	MLAPSE simulation results, physics-informed tools, data diodes for localized AI.

"The agent assembles the structure and exposes the gaps, but the human makes the safety judgment."

SAM RIECK, RTS, ON AGENTIC SUPPORT FOR PART 53

2.2 What regulators signaled

Regulators are learning through practical use rather than waiting for a complete new rulebook. NRC staff described AI assistance inside existing work: inspection preparation and documentation through the agency's internal interface called SimplifAI, and a data-centric audit approach on the AP1000 Revision 20 review that generated advanced safety-evaluation-report templates within hours, with most staff time spent on audit verification.

The international RegLab, run by the OECD/NEA and CNSC with Ennuvo, took the tabletop route. Its second exercise examined AI agents inside the INPO AP-928 work-management cycle. Its stated conclusions were that failure modes sit in the system rather than in the model alone, and that full autonomy doesn't move safety duties away from the licensee, it moves them up front into guardrails, verification and validation, and governance evidence.

The first RegLab exercise pointed the same way. It recommended good-practice guidance covering an assurance framework, explainability, uncertainty, the safety envelope, verification and validation, and graded data quality, and it warned that explainability alone isn't enough for high-consequence decisions without deterministic, physics-based models alongside the AI. Nuclearn made the parallel case for qualifying an AI agent the way a licensed operator is qualified, through training, testing, observation, qualification, requalification, and revocation.

IMPLICATION

Governance should specify more than "human review required." It should name the accountable role, the decision, the evidence the reviewer must see, the failure modes to check, the time available to intervene, and the conditions that force escalation or shutdown.

3 The Real Infrastructure Is Information

AI can't compensate for inaccessible, poorly structured, or context-free information.

Across document-centric and plant-data use cases, speakers returned to the same constraint. Nuclear organizations hold enormous stores of valuable data, but those stores are fragmented across document repositories, local folders, historical databases, drawings, workspaces, and analog instruments.

3.1 Structured retrieval at TerraPower

TerraPower described moving from semantic search over more than 100,000 project documents toward a knowledge graph. The initial search capability built ingestion pipelines and permission-aware access across data silos. The next step connected documents, requirements, components, drawings, and other entities so an agent could traverse explicit relationships rather than rely only on textual similarity. The team reported approximately 6.5 million data points and 10 million relationships.

The lesson wasn't simply scale. TerraPower emphasized source-system permissions, links back to evidence, and the need to improve graph quality by drawing relationships from authoritative requirements databases rather than relying only on computer vision and natural-language extraction.

3.2 Bringing legacy plant data online

Cypress Envirosystems framed "data starvation" as the obstacle to real-time nuclear AI. Many plants still collect pressure, temperature, level, flow, and equipment condition through analog gauges and periodic rounds. Non-invasive digital overlays can keep existing instrumentation in place while supplying higher-frequency data to historians and analytics. The company reports operating experience across 40 nuclear

plants with the data path needing no license amendment requests. In the session, the presenter described readings rising from roughly two manual observations per day to as many as 96 data points per day.

TWO KINDS OF DATA

Retrospective data supports retrieval, licensing, corrective action, and knowledge management. Real-time data supports condition monitoring, fault detection, thermal performance, and operational decision support. Mature programs need governance for both.

3.3 Security architecture is model governance

InfraShield made the case for keeping AI inside the plant boundary: data sovereignty instead of data leaving the facility, and a controlled attack surface instead of an unpredictable one, with local compute and governance layered beneath the applications people use. Its framing was that nuclear regulations govern the AI, not the other way around. In the session, the presenter went further than the slides, placing localized AI on the Level 2 plant network behind a one-way data device and citing an estimated nine-month return on investment with local hardware against continuing cloud costs.

4 What Deployments Are Teaching the Industry

The clearest lessons came from teams already running AI in production.

4.1 Duke: Define the error that matters

For subsequent license renewal, a Duke classifier AI agent helped separate aging-related condition reports from non-aging records. Because the consequential failure was missing a critical record, the team designed around recall rather than generic accuracy. In the session, Duke described targeting 99% recall and cutting an effort of roughly five months for ten engineers to about one month, while keeping engineer review and similar-document evidence in place with the lesson being: “Measure the error that matters to the work.”

Duke’s later retrieval and network-folder applications became progressively more user-led. Search alone couldn’t make forty terabytes of historical content useful. Filenames and metadata were inconsistent, context was ambiguous, and “good” depended on the task. The team’s lesson was to set golden evaluation sets and quantifiable success criteria with subject-matter experts at the start, then keep evaluating as work and data change.

4.2 Southern Nuclear: Make the output fit the work

Southern Nuclear and Everstar described AI-assisted procedure management across eight units and thousands of procedures, plus supporting documents and writer’s guides. The project began with a rubric of 111 weighted criteria, built with procedure writers and subject-matter experts and confirmed against the plant’s quality framework. AI screened procedures against that rubric and presented findings for writers to accept or reject. The Hatch phase scanned 1,334 procedures, with enhancements ranked by rubric and operational benefit, and it is the basis for taking the work to Farley and Vogtle.

The adoption lesson was practical. A technically strong result isn't enough if it arrives in the wrong form. Users wanted familiar deliverables, including Word documents with tracked changes and established review conventions. The project aimed to streamline the existing process and surface potential inconsistencies without pushing a new burden downstream.

4.3 INPO: Organize performance data for intervention

INPO described dashboards and assessment tools that organize fleet and station performance data: a Continuum Viewer with industry, functional-area, and station views, and a Performance Change Dashboard that uses AI to identify what changed between two periods. A Corrective Action Program Insights Generator was presented as coming soon. The intent isn't to label a station automatically. It's to help evaluators find patterns and focus resources.

PRODUCTION PATTERN

The strongest deployments shared five traits: a narrow workflow, a consequential error defined in advance, a subject-matter-expert evaluation set, evidence visible to the user, and a feedback loop that treats deployment as the beginning of validation, not the end.

5 AI Around Physics, Simulation, and Control

Use AI to widen the search. Use physics, qualified calculations, and controlled tests to establish the answer.

Several presentations showed how AI can speed up engineering without displacing the analysis of record. The pattern was to use AI where exploration is expensive, such as sampling, optimization, prediction, or policy testing. Teams then confirm the result through established physics, simulation, or bounded hardware tests.

5.1 MLAPSE: Simulation remains the ground truth

Numerical Advisory Solutions presented MLAPSE, a machine-learning-assisted process simulation engine. Instead of accepting a surrogate model as the final answer, MLAPSE uses machine learning to choose the next simulation cases, focusing resources where uncertainty or acceptance-criterion challenges are greatest. Confirmatory simulation output remains the analysis used for the engineering decision.

5.2 Blue Wave AI Labs: Core design in seconds

Blue Wave AI Labs described CoreDesigner.ai, trained on high-fidelity simulations run across millions of parameter combinations with licensed codes and methods. Near-instant prediction lets engineers explore loading patterns, fuel designs, control strategies, economics, and enrichments approaching 6% interactively, which the deck frames as expanding the design envelope into unknown territory. The model itself is validated against independent test data, with prediction error quantified and benchmarked against physics calculations, and promising designs are confirmed with established licensed methods.

5.3 University of Illinois: Autonomous control develops through testbeds

The University of Illinois described the CERVEROS cyber-physical testbed, which remotely controlled the PUR-1 research reactor against a setpoint derived from real AI computing load: power draw from a GPU cluster running training and inference, scaled to the reactor's rating. Control loops were evaluated at 1 Hz and 10 Hz, with end-to-end latency characterized against the reactor's response. Rather than claim near-term autonomous commercial operation, the work builds a controlled environment for studying latency, load following, and cybersecurity, with microgrid modeling and a nuclear cyber range named as next steps.

5.4 Cambrian: Optimization can reduce design and permitting burden

Other flash talks, including Cambrian Nuclear's work on site selection, applied the same principle to siting and infrastructure decisions: curate authoritative geospatial data, use AI to explore a large option space, and preserve deterministic outputs and documentation. The value proposition wasn't "AI decides the site." It was that experts can examine more alternatives, expose tradeoffs earlier, and generate repeatable evidence faster.

6 Part 53 as a Machine-Readable Decision Architecture

RTS PERSPECTIVE

Part 53 and agentic AI have reached a mutually enabling moment.

In the flash-talk session, Sam Rieck presented RTS research developed with N. Prasad Kadambi, author of NUREG/BR-0303 and co-author of the RTS policy report *Nuclear AI as the Key to Part 53*. Part 53 provides a technology-inclusive, risk-informed, performance-based structure. Agentic systems may make labor-intensive decomposition and traceability practical on a real project schedule.

6.1 The argument in four steps

Step	Part 53 and AI connection
1. Anchor	§53.230(a) sets the single primary safety function that anchors the hierarchy: limit radioactive release. The rule's supporting safety functions, including control of reactivity, heat generation, heat removal, and chemical interactions, sit beneath it.
2. Decompose	The hierarchy runs from the primary safety function to supporting functions, then to systems and components, and finally to the evidence that each one is met.
3. Assist	Agentic AI can carry out preliminary decomposition, traceability and iteration at project speed. The method has existed since at least 2002; what changed is feasibility, not theory.
4. Command	The tooling is built with a human in command, not just a human in the loop. People own the safety judgment and every generated element stays traceable.

The presentation paired prevention through an objectives hierarchy with mitigation through probabilistic risk assessment (PRA). It didn't argue that AI or an objectives hierarchy should replace PRA. It argued that both sides of defense-in-depth deserve deliberate representation, and that AI can lower the cost of maintaining the traceable structure.

The immediate obstacle is training data. RTS recommended closing that gap with curated datasets drawn from ADAMS, so an agentic workflow has authoritative material to reason over, and pairing that with a topical report on the method.


For more information, read the RTS policy brief: *Nuclear AI as the Key to Part 53*, available in September 2026.

6.2 Recommended next step

RTS recommended a topical-report pathway to define the methodology, inputs, controls, traceability expectations, and Human in Command model before a first applicant has to invent them under schedule pressure. That would give the NRC and industry a concrete artifact for evaluating performance-based licensing support tools.

“Part 53 created the regulatory opportunity, and agentic AI may now make the performance-based path practical enough to use.”

SAM RIECK, RTS



7 A Practical Agenda for Nuclear AI

The workshop didn't settle autonomy. It did produce a coherent near-term agenda.

Organizations can act now without pretending every uncertainty has been resolved. The agenda below is addressed to the three groups who can move it.

Recommendation 1

For nuclear operators and developers

- Start with one bounded workflow, and write down the consequential error before choosing a model.
- Build the evaluation set with the people who perform and approve the work. Define "good" quantitatively wherever possible.
- Preserve evidence: source links, inputs, model and prompt versions, intermediate calculations, reviewer decisions, and change history.
- Match deployment architecture to information sensitivity and plant-network level. Distinguish cloud, enterprise, plant, and safety-related environments precisely.
- Design human authority explicitly. Identify who can authorize, reject, override, pause, and recover control, and what information they need to do it.

Recommendation 2

For the NRC and standards community

- Continue tabletop reviews and publish worked examples showing acceptable data characterization, validation, uncertainty treatment, monitoring, and change control.
- Use a graded, consequence-informed assurance model that can address retrieval, decision support, engineering prediction, and autonomous action without treating them as equivalent.
- Clarify expectations for continuously changing models and for the qualified process surrounding a model, not only the model snapshot.
- Develop reference datasets, benchmark tasks, and audit artifacts that let applicants, vendors, and reviewers compare evidence consistently.

Recommendation 3

For AI vendors

- Expose the operating envelope and known failure modes in terms engineers and reviewers can test.
- Integrate into established nuclear work products rather than asking users to abandon controlled workflows.
- Treat permissions, cybersecurity, traceability, and export controls as product requirements, not deployment afterthoughts.
- Measure systemic value, including quality, uncertainty reduction, standardization, knowledge retention, and risk, alongside hours saved.

THE BOTTOM LINE

The workshop's strongest message was disciplined optimism. Nuclear organizations are showing real value today, but durable adoption will depend on evidence, architecture, and accountable command more than on model novelty.

METHOD

Method and Sources

RTS prepared this recap from four inputs: the official NRC Workshop #7 agenda and event webpage, seven session recordings totaling about 8 hours, timestamped automated transcripts of those recordings, and presenter materials supplied by Sam Rieck. Statements are paraphrased unless marked as quotations.

For this version, statements were checked against the presentations the NRC posted in ADAMS package ML26260A206. Where a figure appears in the slides, it is used as the slides state it. Figures that speakers gave

only in their remarks are identified as such, and two decks (University of Michigan and Time Horizon) are image-only and could not be checked.

DISCLOSURE

RTS presented at this workshop. Section 6 summarizes RTS's own research and is labeled as the RTS perspective. Where RTS remarks appear in other sections, they are attributed to RTS.

About the Author

Sam Rieck

Founder and Chief Nuclear Officer, Raisun Technology Services (RTS)

Sam Rieck is the Founder and Chief Nuclear Officer of Raisun Technology Services (RTS), where they lead the company's nuclear strategy and help utilities, suppliers, and advanced reactor developers implement AI and digital technologies across the nuclear sector, while holding to the industry's standards for safety, quality, and compliance.

Sam's career spans government, commercial nuclear, and advanced reactor development. Before founding RTS, they led nuclear licensing and regulatory programs for multiple advanced nuclear companies with hundreds of AI-supported engagements with the NRC and served in senior leadership roles at the U.S. Department of Energy. They speak and publish regularly on the responsible application of AI in nuclear energy, regulatory modernization, advanced reactor deployment, and spent fuel management.

Sam holds dual bachelor's degrees in Mechanical and Nuclear Engineering and in Vocal Music Performance from Kansas State University, and dual master's degrees in Nuclear Science and Engineering and in Technology and Policy from MIT.

SOURCES

References

- [1] U.S. Nuclear Regulatory Commission. “Data Science and Artificial Intelligence Regulatory Applications Workshops,” Workshop #7: Navigating the AI Landscape: Regulator Transformation, Industry Innovation, and Global Collaboration. Hybrid meeting, NRC Auditorium, Rockville, MD, September 15, 2026. Meeting notice, workshop webpage, and agenda (ADAMS ML26253A062).
- [2] NRC Workshop #7 session recordings (seven recordings, approximately 8 hours), with automated transcripts.
- [3] U.S. NRC. Workshop #7 presentation slides, ADAMS package ML26260A206 (individual decks ML26260A216 through ML26260A239), September 15, 2026.
- [4] 10 CFR Part 53, Risk-Informed, Technology-Inclusive Regulatory Framework for Commercial Nuclear Plants. Final rule, 91 Fed. Reg. 15696 (March 30, 2026); effective April 29, 2026. See particularly Subpart B.
- [5] U.S. NRC. (2002). NUREG/BR-0303, Guidance for Performance-Based Regulation. Prepared by N.P. Kadambi, Office of Nuclear Regulatory Research.
- [6] Rieck, S., and Kadambi, N.P. (2026). Nuclear AI as the Key to Part 53: Using AI Agents to Implement Risk-Informed, Performance-Based Licensing. RTS Research Policy Report, Version 1.0, September 2026.
- [7] Rieck, S. “Nuclear AI as the Key to Part 53.” Flash talk, NRC AI Workshop #7, September 15, 2026. Raisun Technology Services (RTS). ADAMS ML26260A239.

Suggested citation

Raisun Technology Services. (2026). *From Pilots to Practice: NRC Artificial Intelligence Workshop #7 Recap*. RTS Research, Version 1.1, September 2026. raisuns.com.

ABOUT RTS

RAISUN Technology Services is a nuclear AI practice built by nuclear operations professionals. RTS delivers nuclear AI advisory, research, and managed services, with the Human in Command®. Governance you can defend. Audit-ready from day one.

Schedule a roadmap session at raisuns.com. One conversation. No pitch deck.