Grounded Reasoning and Artificial Intelligence for National Security (GRAINS) Workshop

September 15-17, 2026

OBJECTIVE: The GRAINS Workshop brings together experts to advance the use of formal methods in rigorously ensuring the security of AI-enabled national security systems and critical infrastructure. The goal is to foster collaboration, identify challenges and success stories, and develop actionable plans with clear ownership to drive innovation, inform policy, and expand programmatic opportunities.

VISIT INFORMATION

ATTIRE

Business Casual.
Long pants, closed-toe and closed-heel shoes are required.

CONTACT

Heather McVey
Protocol Conference Coordinator
(208) 351-5889

Host: Gregory Shannon
Protocol Officer: Heather McVey
Final Agenda

DAY ONE: TUESDAY, SEPTEMBER 15, 2026

ATTIRE FOR TODAY: Business casual. Long pants and closed-toe, closed-heel shoes required.

Engineering Research Office Building | 2525 Fremont Ave., Idaho Falls, ID 83402
Conference Room 159ABC

TIME	EVENT/TOPIC	SPEAKER
7:45 – 8:30 a.m.	**Event Badging & Networking** Heather McVey, *Protocol Conference Coordinator, INL*	
8:30 – 8:40 a.m.	**Welcome and Introductions** Gregory Shannon, *INL Laboratory Fellow, Chief Cybersecurity Scientist, National & Homeland Security; Chief Science Officer at CyManII*	Gregory Shannon
8:40 – 8:50 a.m.	**Welcome to INL** Wayne Austad, *Chief Technology Officer, National & Homeland Security; Chief Research and Development Officer, CyManII*	Wayne Austad
8:50 – 9:00 a.m.	**Workshop Goals, Process, Facilitation** Darcie Martinson, *Facilitator* Jodi Grgich, *Facilitator*	Darcie Martinson Jodi Grgich
9:00 – 10:00 a.m.	**Keynote 1: Who Do You Trust?** Boyd Multerer, *CEO Kry10*	Boyd Multerer
10:00 – 10:15 a.m.	**Short Talk #1: Accountable Partial Understanding for Agentic AI Security** Josiah Dykstra, *RTX/BBN*	Josiah Dykstra
10:15 – 10:30 a.m.	**Short Talk #2: Controller Learning with Automated Requirements, Invariant Checking and Traceability (CLARITY)** Max Taylor, *Assistant Professor of Computer Science, Boise State University*	Max Taylor
10:30 – 10:45 a.m.	**Comfort Break, Move to Breakout Rooms**	
10:45 – 11:30 a.m.	**Breakout #1: Who/What Do You Trust?** Greg Shannon Doug Ghormley, *Senior Scientist, Sandia National Laboratory* Max Taylor	Greg Shannon Doug Ghormley Max Taylor
11:30 – 11:45 a.m.	**Comfort Break, Finish Report-Out Slide**	
11:45 a.m. – 12:30 p.m.	**Report Outs, Discussions, Results** Darcie Martinson Jodi Grgich	Darcie Martinson Jodi Grgich

12:30 – 12:45 p.m.	**Lunch Served**	
12:45 – 1:30 p.m.	**Lunch Presentation and Discussion: INL & N&HS Overview** Wayne Austad	Wayne Austad
1:30 – 2:30 p.m.	**Keynote 2: Verifying, Understanding, and Implementing Critical Systems with Inherent Uncertainty** Steven Holtzen, *Assistant Professor, Northeastern*	Steven Holtzen
2:30 – 2:45 p.m.	**Short Talk #3: State Machines for Embedded Systems** Andrew Johnson, *PhD Candidate Princeton University*	Andrew Johnson
2:45 – 3:00 p.m.	**Short Talk #4: Software Understanding for National Security (SUNS)** Doug Ghormley	Doug Ghormley
3:00 – 3:15 p.m.	**Comfort Break, Move to Breakout Sessions**	
3:15 – 4:00 p.m.	Breakout #2: **Who? What Is Your Adversary?** Wayne Austad Rob Armstrong Christopher Harrison	Wayne Austad Rob Armstrong Christopher Harrison
4:00 – 4:15 p.m.	**Comfort Break, Finish Report-Out Slide**	
4:15 – 5:00 p.m.	**Report Outs, Discussions, Results** Darcie Martinson Jodi Grgich	Darcie Martinson Jodi Grgich
5:00 p.m.	**Travel to Collaborative Computing Center (C3) for Reception**	
Collaborative Computing Center | 955 MK Simpson Blvd., Idaho Falls, ID 83402 Atrium		
5:15 – 7:00 p.m.	**Welcome Reception** *Photography Opportunity in C3 Atrium	

DAY TWO: WEDNESDAY, SEPTEMBER 16, 2026

ATTIRE FOR TODAY: Business casual. Long pants and closed-toe, closed-heel shoes required.

Engineering Research Office Building | 2525 Fremont Ave., Idaho Falls, ID 83402
Conference Room 159ABC

TIME	EVENT/TOPIC	SPEAKER
8:00 – 8:30 a.m.	**Guest Arrival & Badge Check** Heather McVey, *Protocol Conference Coordinator, INL*	
8:30 – 8:50 a.m.	**Welcome Co-Organizer Reflections**	Co-Organizers
8:50 – 9:00 a.m.	**Facilitators' Guidance for the Day** Darcie Martinson Jodi Grgich	Darcie Martinson Jodi Grgich
9:00 – 10:00 a.m.	**Keynote 3: Automated Reasoning at Cloud Scale** Robert Jones, *Sr. Principal Applied Scientist, AWS*	Robert Jones
10:00 – 10:15 a.m.	**Comfort Break**	
10:15 – 11:00 a.m.	**Presentation and Discussion: DOE Genesis Mission & INL Prometheus Project** Logan Browning, *Visualization Engineer, Nuclear Science & Technology, INL* Daniel Schwen, *Distinguished Computational Scientist, Nuclear Science & Technology, INL*	Logan Browning Daniel Schwen
11:00 – 11:15 a.m.	**Short Talk #5: Autonomous Reactor Control** Ryan Stewart, *Data Scientist INL*	Ryan Stewart
11:15 – 11:30 a.m.	**Short Talk #6: Argo Demonstration** Rebecca Andrews, *Forward Deployed Engineer, Atalanta*	Rebecca Andrews
11:30 a.m. – 12:15 p.m.	**Breakout #3: What are Viable AI-OT Frameworks?** Lance Joneckis, *Senior Cybersecurity Researcher*, *INL Laboratory Fellow, National & Homeland Security* Rob Armstrong Christopher Harrison	Lance Joneckis Rob Armstrong Christopher Harrison

12:15 – 12:30 p.m.	**Lunch Served, Finish Report-Out Slides**	
12:30 – 12:15 p.m.	**Working Lunch: Report Outs, Discussions, Results** Darcie Martinson Jodi Grgich	Darcie Martinson Jodi Grgich
1:30 – 2:30 p.m.	**Keynote 4: Neural Network Verification for Critical Infrastructure: From Robustness to Full-Stack Assurance** Anne Tumlin, *PhD Candidate Vanderbilt University*	Anne Tumlin
2:30 – 2:45 p.m.	**Short Talk #7: Gradual Capability-Typed Agents: Securing AI Intent Through Capabilities of Thought** Jon Aytac, *Senior Staff R&D S&E Computer Science, Sandia National Laboratory*	Jon Aytac
2:45 – 3:00 p.m.	**Short Talk #8: Inference Efficiency** Christopher Harrison, *Senior Scientist, Sandia National Laboratory*	Christopher Harrison
3:00 – 3:15 p.m.	**Comfort Break, Move to Breakouts**	
3:15 – 4:00 p.m.	**Breakout #4:** **How do Architectures Handle AI Failures?** Greg Shannon Doug Ghormley Max Taylor	Greg Shannon Doug Ghormley Max Taylor
4:00 – 4:15 p.m.	**Break, Finish Report-Out Slide**	
4:15 – 5:00 p.m.	**Report Outs, Discussion, Results** Darcie Martinson Jodi Grgich	Darcie Martinson Jodi Grgich
5:00 p.m.	**Adjourn**	

DAY THREE: THURSDAY, SEPTEMBER 17, 2026

ATTIRE FOR TODAY: Business casual. Long pants and closed-toe, closed-heel shoes required.

Engineering Research Office Building | 2525 Fremont Ave., Idaho Falls, ID 83402

TIME	EVENT/TOPIC	SPEAKER
7:00 – 7:15 a.m.	**Guest Arrival & Visitor Badge Check** Heather McVey, *Protocol Conference Coordinator*	
7:15 – 8:15 a.m.	**Travel to EBR-I** INL Transportation	
Experimental Breeder Reactor–I Museum (EBR-I)		
8:15 – 9:30 a.m.	**Ingress & Tour** Break into 2-3 groups	
9:30 – 10:00 a.m.	**Egress & Travel to MFC**	
Materials and Fuels Complex		
10:00 – 10:10 a.m.	**Ingress MFC & Split into Two Groups**	
Group 1		
10:10 – 10:20 a.m.	**Walk to Hot Fuel Examination Facility (HFEF)**	
10:20 – 10:50 a.m.	**Tour HFEF**	
10:50 – 11:30 a.m.	**Tour Neutron RADiography Reactor (NRAD)**	
11:30 – 11:40 a.m.	**Walk to Manipulator Repair Group**	
11:40 a.m. – 12:10 p.m.	**Tour to MRG**	
Group 2		
10:10 – 10:20 a.m.	**Walk to Manipulator Repair Group**	
10:20 – 10:50 a.m.	**Tour MRG**	

10:50 – 11:00 a.m.	Walk to Hot Fuel Examination Facility (HFEF)
11:00 – 11:30 a.m.	Tour HFEF
11:30 a.m. – 12:10 p.m.	Tour Neutron RADiography Reactor (NRAD)
12:10 – 12:20 p.m.	Egress MFC, Load Bus
12:20 – 1:00 p.m.	Travel to Idaho Falls
1:00 p.m.	Conclude

Keynotes

Keynote #1, 9:00 a.m., Tuesday September 15th

Boyd Multerer, *CEO Kry10*

Title: Who do you trust? Formal Verification for OT systems defending against modern cyber threats and deployed AI workloads.

Abstract: For operational technology (OT) operators, AI raises the question of trust on two fronts simultaneously. Firstly, AI tools have materially expanded the adversary's reach. Vulnerability discovery, exploit generation, and targeted attack capability are now accessible to any motivated individual. Edge devices and OT systems that were previously difficult targets face a broader and faster-moving threat environment. Secondly, AI as a deployed workload introduces a different trust question: the AI running on your systems, and increasingly, the AI running your systems. Can they be trusted to keep operations running?

In this session, we outline how formal verification can address both fronts and deliver what OT operators actually need: greater efficiency and agility that comes from having confidence in their systems. For OT operators, this confidence ensures faster installation and update, and faster reaction to threats. The proven foundation provided by formal verification enables operators to demonstrate what has and has not changed — compressing approval cycles without reducing assurance and creating the conditions for workloads whose behaviour can be contained and, when necessary, safely removed.

Applying FM isn't magic and it's not impossible. The keynote will look at frameworks for approaching this problem space and challenge participants to consider what a solution might look like in their context.

Bio: Boyd Multerer brings a practitioner's perspective to GRAINs. As one of the architects of the Xbox at Microsoft, Boyd built systems relied on by hundreds of millions of users, where the cost of getting trust wrong was immediate, and public. Now as CEO of Kry10, having spent seven years building formally verified operating infrastructure for OT and critical devices, he has worked from first principles on what it takes for software to be genuinely trustworthy in high-stakes environments.

Keynote #2, 1:30 p.m., Tuesday September 15th

Steven Holtzen, Assistant Professor, Northeastern University

Title: Verifying, Understanding, and Implementing Critical Systems with Inherent Uncertainty

Abstract: How do we develop and deploy trustworthy software systems with uncertain behavior? Uncertainty is omnipresent in today's software systems: it comes from interactions with the physical world, randomized algorithms like in cryptography, and increasingly from AI components like large language models (LLMs). Uncertainty poses significant technical challenges for ensuring trustworthiness: it is difficult to test software with uncertainty, since the system by design behaves differently on each invocation. Moreover, AI can fail: it is typically the goal to ensure that this failure happens with sufficiently low probability. As uncertainty increasingly proliferates through our software, we need to ask: what should our specifications be for software with uncertainty, and how do we ensure that our software meets those specifications?

We will outline emerging approaches to formal reasoning for probabilistic systems. Our aim is to:

1. Provide a common programming foundation for software with uncertainty based on probabilistic programming languages (PPLs).
2. Give a formal language for specifications about the behavior of probabilistic programs based on probabilistic separation logic.
3. Give tools for verifying and implementing probabilistic programs that meet those specifications.

Ultimately, these three components come together to enable a new workflow for implementing and verifying high-consequence systems with uncertainty.

Bio: Steven Holtzen is an assistant professor at Northeastern University. His research focuses on designing tools and foundations for reasoning about programs with probabilistic behavior. He received his PhD. in computer science from the University of California, Los Angeles in 2021. His work has been recognized by two ACM SIGPLAN distinguished paper awards and an NSF CAREER award. He was a Member of Technical Staff at Sandia National Laboratories, as a member of the cyber data and analytics team from 2015 — 2021.

Keynote #3, 9:00 a.m., Wednesday September 16th

Robert Jones, Sr. Principal Applied Scientist, AWS

Title: Automated Reasoning at Cloud Scale

Abstract: Amazon has what is likely the largest team of automated reasoning experts ever assembled. Why? Amazon works backwards from customers, and those customers want the assurance that only automated reasoning can provide. This is especially critical as AI agents write production code — and become production code themselves. I'll survey why our customers care about automated reasoning, with selected examples from three areas: foundational capabilities, solution building blocks, and futures. We'll take a deeper dive into a couple of those examples based on audience interest.

Bio: Robert Jones is a Senior Principal Applied Scientist at AWS, where he leads the solver/Lean science team in the Automated Reasoning Group. He works on neurosymbolic AI and automated reasoning in the cloud, for the cloud. Prior to joining AWS, Robert worked in hardware verification, modeling, and security. He holds a PhD in Electrical Engineering from Stanford University. He enjoys family, reading, photography, playing the piano, and getting outside in the Pacific Northwest.

Keynote #4, 1:30 p.m., Wednesday September 16th

Anne Tumlin, PhD Candidate Vanderbilt University

Title: Neural Network Verification for Critical Infrastructure: From Robustness to Full-Stack Assurance

Abstract: Neural network verification has progressed from robustness proofs for small classifiers toward richer architectures, broader properties, and increasingly complex cyber-physical applications. This talk traces that trajectory and considers what verification must achieve to provide meaningful assurance for critical infrastructure. Using the International Verification of Neural Networks Competition (VNN-COMP) as a measure of the field's progress, we review the foundations of contemporary verification methods and persistent challenges in scalability and expressive specifications. The discussion then turns to emerging applications relevant to the security and resilience of U.S. critical infrastructure and national security, including graph neural networks (GNNs) for energy delivery and grid operations, neural ODE and PDE models for physical processes, and closed-loop autonomy in operational technology environments. Looking forward, we argue that integrating language models and agentic AI into operational workflows will demand meaningful guardrails, and that developing verification approaches capable of certifying them is a grand challenge for the field. The talk concludes with a vision for full-stack assurance of AI-enabled critical infrastructure through compositional guarantees spanning the computational and physical stack and supported by independently checkable evidence.

Title: Anne Tumlin is a fourth-year PhD student at Vanderbilt University whose research develops formal verification methods for learning-enabled systems, with a focus on certifying the safety, robustness, and reliability of neural networks and their applications to critical infrastructure. Her work explores verification across a range of learning architectures and uncertainty models, with particular interest in applications to the electric grid and grid security. Anne is a Department of Energy Computational Science Graduate Fellow, a Russell G. Hamilton Scholar at Vanderbilt University, and a year-round intern at Sandia National Laboratories. She completed her CSGF practicum at Sandia in summer 2025 and now works on the verification of learned components within grid-security architectures. Anne is advised by Taylor T. Johnson in the VeriVITAL Lab and by Tyler Derr in the Network and Data Science Lab at Vanderbilt University. Her work at Sandia is conducted under the mentorship of Georgios Fragkos, Shamina Hossain-McKenzie, and Logan Blakely.