Exceptional service in the national interest

Digital Assurance for High Consequence Systems (DAHCS) Mission Campaign FY26 Call for LDRD Proposals

Sandia National Laboratories
Digital Assurance for High Consequence Systems Mission Campaign Team

December 5, 2024

Table of Contents

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2024-16572O

Investment Area Background

The ***Digital Assurance***[1] for ***High Consequence Systems***[2] (DAHCS, pronounced "Dax") Mission Campaign (MC) invites research ideas that help to ensure that the use of digital technologies does not weaken our nation's high consequence systems[3]. Digital technologies offer significant benefits in speed, cost, and flexibility, and we seek to reap those benefits without introducing new system failures. Today, we lack the means to efficiently evaluate digital technologies with rigor and confidence to assure they operate as intended.

To address this gap, we seek to fund research discovering metrics and *principles* (abstractions, approaches, and assumptions) that are unique to assuring that embedded cyber-physical controllers[4] do not *fail to function* (i.e., their availability and reliability is assured, including in the presence of up to one human or digital insider threat). We call for proposals ranging from foundational research to mission-focused engineering that seek to address this gap by creating tools, techniques, and methods to confidently characterize, evaluate, and manage digital risks to our high consequence systems:

- developing the scientific foundation needed to create rigorous, rapid, cost-effective, generalizable digital assurance across many types of high consequence systems and their lifecycles (including design, qualification, and sustainment),
- creating an ecosystem of capabilities that gives us confidence in the digital assurance of our high consequence systems,
- and proving that we can:
    - characterize the digital technologies within our systems at any point in their lifecycles,
    - assess the risks to our systems from digital technologies, moving well beyond vulnerability-focused security,
    - and select among design and implementation options that appropriately manage and accept digital risks while balancing against other systems-level trade-offs (for example, resilience, reliability, safety, size, weight, power, cost).

Ultimately, DAHCS must produce solutions that can be adopted by system organizations and that enable decision makers to make confident, evidence-based system trade-offs that consider mission risk from current and future digital threats.

Note: vulnerability detection, information technology (IT) systems, systems of systems, and scaling existing algorithms are out of scope *unless unique to HCS and their embedded controllers*.

---

[1] For our purposes, ***digital assurance*** includes processes, measures, and/or controls applied to digital technologies to ensure that a given system fulfills its intended purpose, even given current and future digital threats [NNSA SD 452.4-1 Nuclear Enterprise Assurance (NEA), 1/27/2022]. We include digital technologies both *within* and *influencing* HCS, and we include threats such as active adversaries, cyber-attacks, supply-chain issues for components and tools, an insider, natural environmental hazards (both digital and physical), and both unintended behaviors (e.g., from errors) and emergent behaviors.

[2] For some systems, the cost of failure is catastrophic, e.g., death or existential threat to a nation. These ***high consequence systems*** (HCS) are created to serve very specific missions. In the DAHCS MC, we focus on four HCS types: nuclear deterrence, hypersonics, satellite, and individual critical infrastructure (e.g., nuclear power generator) systems.

[3] See the DAHCS whitepaper (2024) for more information (using slightly older terminology); our soon-to-be-released definitions page will provide more context. We hope to have this available in January.

[4] Including hardware, software, firmware, integration, and the ecosystem from requirements to retirement.

DAHCS MC Roadmap:

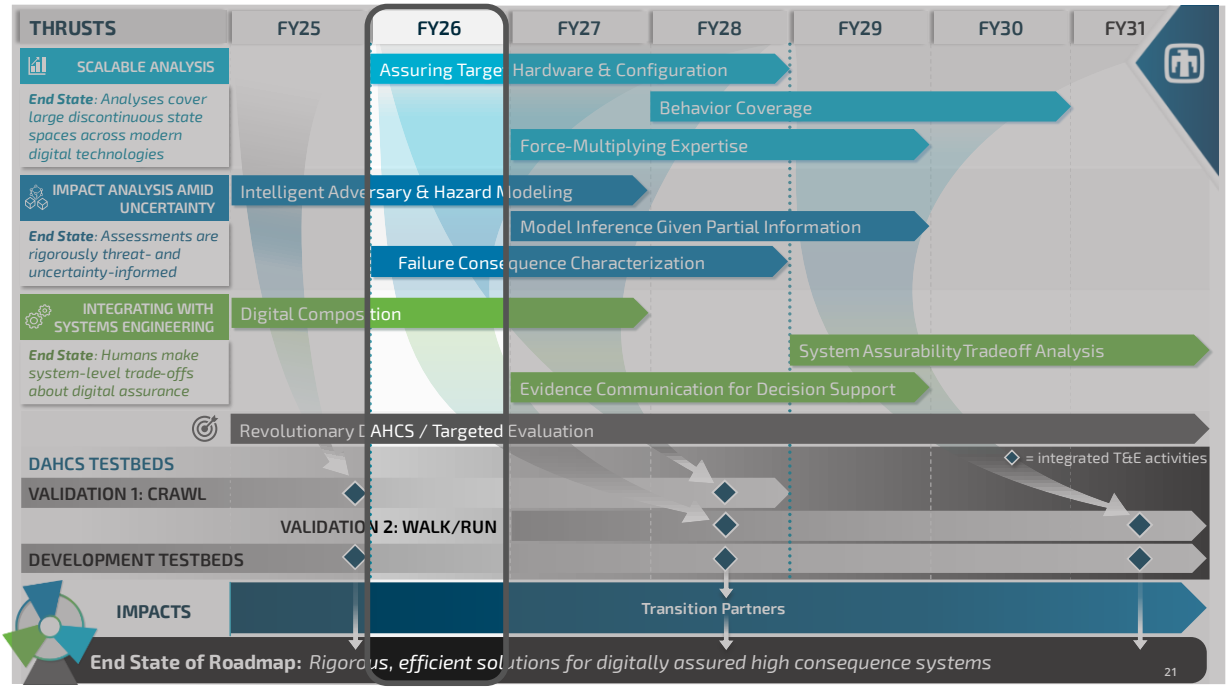

The DAHCS MC inspirational roadmap contains three broad technical thrusts, which are further subdivided into research challenges. **We strongly encourage all revolutionary ideas ("Revolutionary DAHCS")** that enable end-to-end digital assurance of HCS or create, evaluate, and use evidence in *assurance cases*[5] (which make *claims* about digital technologies both within and directly influencing an HCS). The DAHCS technical thrusts are:

1. ***Scalable Analysis:*** Enable at least two orders of magnitude improvement (speed, cost, scale, etc.) in creating and evaluating DAHCS arguments, focusing on assuring the digital behaviors of cyber-physical embedded systems.
2. ***Impact Analysis Amid Uncertainty:*** Measure and increase confidence in an assurance case and its evidence despite incomplete information (e.g., by identifying what additional information is needed to increase confidence by how much).
3. ***Integrating with Systems Engineering:*** Support systems-level decisions about digital assurance and residual risks, including making trade-offs among digital technologies and digital design options.

[5] D. J. Rinehart, J. C. Knight and J. Rowanhill, "Current Practices in Constructing and Evaluating Assurance Cases With Applications to Aviation," NASA, 2015.

For the FY26 cycle, we place extra emphasis on four research challenges[6]:

1. ***Intelligent Adversary and Hazard Modeling:*** How do we explicitly account for adversary goals, choices, and capabilities, particularly to support threat-informed prioritization and dynamic reassessment? Solutions might include methodologies and metrics that incorporate information about adversaries to enable well-characterized, repeatable, rapid, full-stack digital assurance.
2. ***Failure Consequence Characterization:*** How do we enable *end-to-end* reasoning about consequences of failures, including understanding *direct impacts* such as the impact of an incorrect input or a single timing delay, understanding *aggregate failures* like radiation-induced bit flips plus timing delay, and understanding *indirect impacts* such as follow-on failures caused by an upstream failure? Which changes might influence a particular behavior relevant to the system and use case, and which are irrelevant? Which consequences should be prioritized? Solutions might include new methods and metrics that rapidly detect and characterize the impacts of such changes or that extend our ability to characterize consequences.
3. ***Assuring Target Hardware and Configuration:*** How do we confidently assert a device and its configuration (e.g., firmware) do not contain unexpected characteristics or exhibit unexpected behaviors due to an adversary's influence? The problem's scope encompasses a wide range of hardware (custom integrated circuits to commercial-off-the-shelf microelectronics to finished circuit boards) and all phases of the lifecycle. Our goal in this thrust is to ensure integrity and authenticity of target hardware and binary data needed to operate.
4. ***Digital Composition:*** How do we rigorously combine information from disparate techniques to support system-specific, risk-informed decisions? To reason about system-level assurance, we seek new methods and metrics that rapidly stitch, fix, and compare assurance evidence. We seek full-stack solutions that characterize and limit digital risk from emergent behaviors and enable multi-fidelity reasoning across vast, interconnected hardware and software behavior spaces.

DAHCS will execute a highly collaborative portfolio of LDRD projects that enable transformative DAHCS solutions, characterize fundamental bounds on what we can confidently build and maintain, and reduce barriers (technical, behavioral, etc.) for technology adoption.

We recommend that proposers explicitly address alignment to our Scenario-Based Test & Evaluation (T&E) activities and clearly connect their research outputs to a claim or question about the component/testbed/system. The DAHCS T&E team will test tools, techniques, and methods developed under MC funding on testbed controllers (UUR testbed: the commercial satellite application Satellite Identification and Location, SIDLOC, [website]; potentially custom ASICs in the future) to create assurance cases for three scenarios:

1. ***Rapid Reassessment:*** Provide, within two weeks, an updated assurance determination and proposed actions given a technical surprise (e.g., a new threat, a failed test)
2. ***Rapid Build:*** Build, within six months, a new controller with requirements altered from a prior design but with as much digital assurance as possible within the timeframe
3. ***100% Solution:*** Aim to build, at whatever cost, an entirely cyber-secure, digitally assured controller (we assume this is impossible, but we aim for it)

Participation in T&E activities is critical to evaluating the generalizability, interoperability, scalability, and/or rigor of DAHCS research.

[6] Upcoming web releases will contain a variety of resources, including one-page descriptions of the research challenges. We hope to have this available in January.

General Guidance

A successful DAHCS MC proposal should consider and/or include the following:

- Demonstrate relevance to Sandia and DAHCS MC strategic objectives, research challenges, and scenarios.
- Highlight how results could be integrated into an assurance case, Digital Thread, and/or other digital systems engineering design and integration tools and approaches.
- Highlight plans to characterize why and under what assumptions/conditions an approach or method "succeeds" or "fails".
- Include steps to address anticipated downstream barriers to integration and adoption.
- (Encouraged) Include external collaborations to broaden the expert community addressing DAHCS MC objectives. Potential non-Sandia partners may reach out to the DAHCS MC team or the Sandia University Partners Network (SUPN) for help finding potential Sandia PI partners.
- Consider creative project plans to enable higher risk research (e.g., official fast-fail go/no-go criteria) and/or to increase student engagement (e.g., options that enable students to direct part of the research). We want to hear your ideas.
- Budget for extensive inter-LDRD collaboration.
- Consider explicitly relating to current and completed projects.
- Plan to deliver artifacts supporting research reproducibility and legacy.
- Describe a complete, self-contained problem statement and research plan.
- Consider using our testbed controllers, but, if needed, at least demonstrating utility with clearly defined testbeds, metrics, and methods.

A successful DAHCS MC LDRD project will accomplish at least one of the following:

- Discover, via case study, whether a given approach is warranted for a DAHCS problem.
- Create and/or characterize a novel and rigorous approach to a well-scoped DAHCS problem.
- Demonstrate/prove that a capability can be generalized and/or scaled repeatably.
- Pioneer a revolutionary solution that may not be explicitly in the roadmap.

Project Selection Criteria

- **Programmatic**
  - Alignment: Addresses why the proposal requires LDRD funding from this investment area, aligning with multiple items explicitly in the call, the spirit of the call, and lab priorities
  - Impact: Describes impact to investment area, Sandia missions, and the nation, including potential outcomes and deliverables
- **Technical Science/Innovation**
  - Merit: Describes truly novel, leading-edge research and/or development
    - States a clear research question and interesting hypothesis, including relevant metrics
    - Describes anticipated knowledge to be gained and/or new DAHCS capabilities
    - Improves technology readiness level (TRL) and/or ***human readiness level (HRL)*** ***[SAND2020-5713C]***
    - Shows strong possibility of publication, patents, etc., and extends state-of-the-art approaches
  - Feasibility: Describes a sound but aggressive project plan
    - Shows strong balance of ambition (high technical risk) and practicality (thoughtful risk mitigation)
    - Describes clear go/no-go, **fail-fast** decision points
  - Qualifications & Budget:
    - Requests appropriate resources
    - Describes a diverse, multi-disciplinary team with a compelling mix of staff (including experts, new staff, and consultants).
- **Mission Campaign Alignment**
  - HCS Differentiation: Describes how the research is specific to digital assurance for HCS, i.e., makes or uses assumptions about characteristics of HCS and/or their lifecycles that are unique to HCS
  - Test & Evaluation (T&E):
    - Describes a strong test and evaluation (T&E) plan, including integration with our T&E team to support creation of a DAHCS ecosystem
    - Clearly demonstrates impact to HCS missions and/or testbeds
    - Initiates steps to move beyond an individual research project towards integration with other DAHCS investments and adoption by system organizations
  - MC Advances: Demonstrates at least one (if not all) of the following advances:
    - ***Generalization:*** Easily generalizes to HCS beyond those demonstrated in T&E plan
    - ***Interoperability:*** Creates strong ties between currently incompatible methods and highlights plan or vision for integration with other DAHCS projects
    - ***Scalability:*** Ensures rapid determination of digital assurance
    - ***Rigor:*** Significantly advances credibility of digital assurance arguments

Contact Information

DAHCSMC-help@sandia.gov

Developing the scientific foundation needed to create **rigorous**, **rapid**, **cost-effective**, **generalizable** digital assurance **across high consequence systems' lifecycles**