U.S. DEPARTMENT of ENERGY

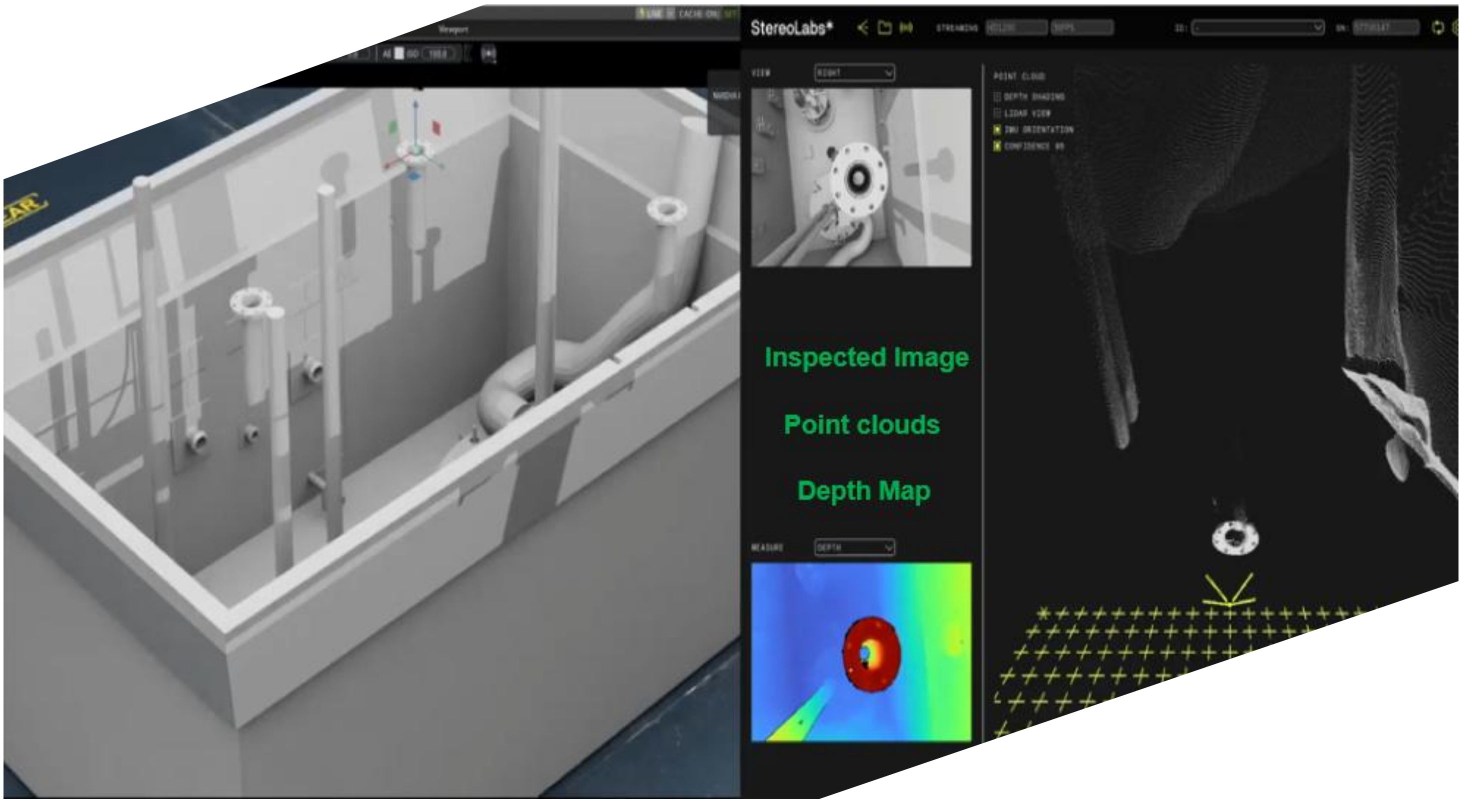

Digitally Optimized Autonomous Robotic Systems for Hanford Waste Tank Handling, Year 2 Summary

Automated Pit Exploration System (APES) Year 2 Report
Hanford Tank Waste R&D Award 278709

June 2026

Disclaimer

Authors

The authors of this report are:

Anthony D'Andrea, Idaho National Laboratory

Kaleb Houck, Idaho National Laboratory

Pengcheng Cao, Idaho National Laboratory

Celing Li, Idaho National Laboratory

Leonel Lagos, Florida International University

Anthony Abrahao, Florida International University

Munim Rayhan, Florida International University

Derek Hansen, Brigham Young University

Elise Bangerter, Brigham Young University

Acknowledgments

This report was prepared by the Idaho National Laboratory, Florida International University, and Brigham Young University for the U.S. Department of Energy, Office of Environmental Management.

This document was prepared with the assistance of an AI-based drafting tool.

List of Acronyms

AI	artificial intelligence
APES	Automated Pit Exploration System
BYU	Brigham Young University
CAD	Computer-aided design
DOE	U.S. Department of Energy
DT	Digital twin
FIU	Florida International University
FK	Forward kinematics
GPU	Graphics processing unit
H2C	Hanford Tank Waste Operations and Closure
IMU	Inertial measurement unit
INL	Idaho National Laboratory
LiDAR	Light Detection and Ranging
ML	machine learning
MR	mixed reality
PDU	Process Development Unit
PTZ	Pan-tilt-zoom |
R&D	research and development
RGB	Red, green, blue
ROS 2	Robot Operating System 2
SLAM	Simultaneous Localization and Mapping
URDF	Universal robot description files
XR	Extended Reality

Executive Summary

The Automated Pit Exploration System (APES) is an autonomous, multi-robot inspection system developed by the Idaho National Laboratory (INL) in collaboration with Florida International University (FIU) and Brigham Young University (BYU) to support pit-inspection operations at the Hanford waste-tank farms. Rather than relying on a single complex robotic platform, APES coordinates multiple task-specific robots: a custom cable-driven manipulator arm, a commercial mini crawler, a pan-tilt-zoom camera, and a Nexxis Argus light detection and ranging (LiDAR) mapping system, each deployed through a 4-in. view port and powered from a Ford F-150 Lightning electric pickup truck.

Year 2 efforts advanced all major subsystems. The FIU cable-driven robotic arm progressed from conceptual design to a refined prototype with improved rigidity, payload capacity, and additional axes of rotation. The Nexxis robotic crawler was procured and delivered, replacing the previously planned university partnership to develop a snake-robot. The Nexxis Argus LiDAR mapping system was acquired and integrated for pit-scale three-dimensional mapping. Software development advanced significantly, with a distributed ROS2 architecture validated through real-world testing, a digital-twin framework built on DeepLynx and Apache Airflow, and an operator interface prototype developed collaboratively with BYU in Unity. A multi-robot LiDAR-scanning case study was conducted using a Boston Dynamics Spot quadruped and a Yahboom Rosmaster R2, demonstrating successful point cloud alignment and merging a key proof of concept for the pit-inspection mission. INL also established a high-fidelity Isaac simulation environment with realistic lighting for perception and autonomy development while FIU synchronized the manipulator model with MoveIt 2 for planning and inspection-coverage analysis. While half of Year-2 funding was held up due to the current political climate, project leadership reviewed tasks to shift focus and budgets to continue the most-impactful research.

Table of Contents

List of Figures

List of Tables

1 System and Operational Requirements

Work conducted within the Hanford waste-tank farms is governed by rigorous safety and operational constraints that must be carefully considered throughout system development. Although the Automated Pit Exploration System (APES) produced under this effort is intended as an advanced prototype platform, designing in alignment with these requirements is critical to support future deployment and acceptance at the Hanford Site. These requirements are established through collaboration between the Hanford Site and the Idaho National Laboratory (INL) engineering and project-management teams.

To ensure consistency across project phases, the functional-requirements table originally developed in Year 1 is included again in this Year 2 report as a reference point. This table will be revisited and revised as needed to capture any updates to system requirements, design changes, or refined compliance criteria identified during ongoing work. Retaining and updating this information provides traceability while ensuring the design remains aligned with current expectations.

Table 1. Requirements Compliance Table

ID	Requirement	Compliance Level (0–10)	Notes
1	The system will operate within environmental conditions described in Section 1.1.	10	All conditions can be met.
2	The system will adhere to design requirements within Section 1.2.	9	The project scope is to build and test a prototype system. Durability and drop tests will be limited as necessary to complete a demonstration system.
3	The system will adhere to visual inspection requirements within Section 1.3.	9	Video streaming capabilities are limited by cameras that can fit down the viewport. Frame rate and resolution are trade-offs that need to be made. Greater resolution is not always preferred over reasonable frame rates. Multiple camera systems will be used including RGB-D, Monocular, PTZ.

ID	Requirement	Compliance Level (0–10)	Notes
4	The system will adhere to data collection requirements within Section 1.4.	8	Research into radiation sensors is ongoing. Sensors identified to fit within the system may fall short of detecting up to 1,000.
5	The system will adhere to software and controls system requirements within Section 1.5.	10	All conditions can be met.
6	Section 1.6 contains optional system functionality (outside the original scope of the proposal) that the Hanford Site has expressed interest in a pit inspection system doing.	Optional	INL will investigate and implement additional functionality as time and budget allow. Other material considerations for a final build would be evaluated and refined at a later date for a full deployment system. In Year 2 design review Hanford representatives have clarified that 100% visual inspection is the main focus of the technology, features should not be added that risk lowering the quality of the primary goal.

1.1 Environmental Conditions

The waste-tank farm pits at the Hanford Site present a highly hazardous operating environment that significantly limits direct human interaction. These pits are underground radiological spaces characterized by beta-radiation levels ranging from approximately 5 to 1,000 mrem/hour (primarily from strontium) and gamma-radiation levels ranging from 1 to 60 rem/hour (primarily from cesium), posing serious health risks from even short-term exposure. In addition to radiological hazards, the environment includes extreme and variable temperature conditions, with operating ranges of 14–104°F and storage conditions of 0–120°F. The pits are completely devoid of natural or installed lighting, requiring all inspection and operational equipment to supply their own illumination. High humidity levels, reaching up to 90% non-condensing, combined with the presence of liquids, further complicate operations and necessitate robust, sealed equipment with ingress protection ratings of IP65 or greater for manipulators and IP67 or greater for sensitive instrumentation. Communication is also challenged by the approximately 2-foot-thick concrete covers, which significantly attenuate wireless signals, requiring hardwired tethering for reliable data transmission.

Due to this combination of radiological exposure, confined access, environmental extremes, and limited visibility, minimizing human entry into these pits is critical for worker safety and operational efficiency. Robotic systems provide a practical and necessary solution by enabling remote inspection and manipulation within these hazardous conditions. By removing personnel from direct exposure while maintaining operational capability, robotic platforms not only reduce health risks but also allow for more consistent, repeatable, and extended-duration work in environments that would otherwise be unsafe or impractical for human workers.

1.2 Standard Design Requirements

The design of the system is driven primarily by the requirement to achieve complete, high-quality visual inspection of the tank-farm pits, ensuring that 100% of the pit surfaces can be viewed and inspected. All other data-collection capabilities and object interaction are considered secondary to this core objective. To support this, the system is designed to be deployed through a nominal 4-in. viewport; no component shall exceed a 3.8-in. diameter to accommodate variations in field conditions. Once deployed, the system must be capable of thoroughly scanning a representative pit area, with performance validated through testing in a full-scale mock-up environment. Inspection operations are expected to be efficient, with a full pit scan completed within one hour and requiring minimal entries into via build in view ports—ideally no more than three—to reduce operational complexity and exposure risk.

The system integrates multiple sensor payloads to collect and map pit data, including real-time visual imaging, radiation detection, LiDAR mapping, and environmental monitoring such as temperature, pressure, and humidity. However, these sensing capabilities are intended to supplement and enhance the primary visual inspection mission. To ensure visibility in the unlit pit environment, the system shall provide its own lighting. Additionally, an overhead pan-tilt-zoom (PTZ) camera offers a third-person perspective to support operators during deployment and operation.

From a mechanical and deployment standpoint, the system is designed to operate within strict spatial and structural constraints while remaining simple to deploy. Supporting structures are limited both below and above the cover block, and the entire system shall be designed to be assembled and operated by two personnel without the need for heavy equipment such as cranes. Deployment and retrieval shall be automated through the above-ground structure, which shall be lightweight, modular, and compatible with a single ground vehicle platform for transport and power. The system shall incorporate both a robotic arm for overhead inspections and a crawler configuration for accessing the bottom of the pit, ensuring full coverage of the inspection area. While the robotic arm should be capable of limited manipulation tasks, such as collecting small contamination samples and handling payloads less than 10 lb, these functions are secondary to the primary inspection mission. Durability is also a key consideration, with the system designed to withstand handling impacts. Nevertheless, it will not be drop tested. The system should improve reliability and maintainability by using commercial-off-the-shelf components wherever practical.

1.3 Visual Inspection Requirements

The visualization strategy for the APES system is centered on capturing clear, high-resolution imagery and video to enable complete visual inspection of 100% of the pit. Achieving full coverage with reliable, high-quality visual data is the primary objective, and the system is designed to ensure that all areas of interest can be observed in detail to identify damage or wearing of the enamel pit lining. This visual inspection can also be done to monitor equipment inside of the pit to detect any abnormalities. To support this, the platform incorporates high-resolution cameras capable of producing up to 9-megapixel still images and 4K video output, with the ability to both livestream and record footage for later analysis. Adequate onboard lighting is integrated to ensure image clarity in the completely dark pit environment, and camera systems include zoom and focus capabilities to allow operators to inspect specific features more closely. Additional focus has been placed on the lighting, with the coordinated efforts of all of the robots and the PTZ camera to more fully light the area and reduce shadows and glares.

In addition to standard imaging, the system includes depth sensing and LiDAR technologies to generate spatial awareness and point-cloud data, which can assist operators in understanding geometry and navigation within the pit. Depth sensing is achieved using a ZED X mini RGB-D sensor camera, which combines conventional color imaging with distance measurements to capture the three-dimensional (3D) structure of the environment. However, these capabilities are intended to complement, not replace, the primary visual-inspection mission. Earlier design considerations included the integration of an infrared (IR) camera to identify subsurface defects such as cracks or coating degradation. While this approach offers potential benefits, it is no longer a primary focus moving forward because project resources are prioritized toward improving the quality, coverage, and reliability of visible-spectrum inspection methods.

The physical configuration of the system supports comprehensive visual access throughout the pit. The robotic arm is designed to carry the RGB-D camera, monocular camera, and LiDAR all at once for overhead and targeted inspection while the crawler platform provides close-up views of lower regions and the pit floor with a standard monocular camera. To further enhance coverage and reduce blind spots, the team is investigating the addition of a pan-and-tilt mechanism to the crawler-mounted camera, allowing for greater flexibility in viewing angles without requiring repositioning of the robot. Together, these design choices ensure that the system can deliver thorough high-quality visual inspection across the entire inspection area.

1.4 Data-Collection Requirements

The APES platform is designed to support both manual operation and autonomous inspection, with a strong emphasis on automated data collection to improve efficiency and reduce operator burden. The system integrates a suite of environmental and radiological sensors, including radiation, temperature, vapor, dew point and humidity, and pressure, all of which are deployed on the robotic arm to enable comprehensive data acquisition inside of the pits. During operation, sensor data is continuously collected and made available to operators in real time while also recorded for post-mission analysis. These data are spatially correlated with generated maps of the pit and presented within a digital-twin environment, allowing for intuitive visualization, historical tracking, and enhanced situational awareness. Although data collection will be primarily performed inside of the pit, data collection can continue outside of the pit environment as well, but to a lesser extent.

To ensure meaningful radiological characterization, the radiation-sensing system is designed with an angular resolution of 5 degrees or more and will be capable of detecting a wide range of dose rates, from approximately 5 to 1,000 rem per hour. By combining autonomous deployment, synchronized data visualization, and integrated mapping, the system enables more-consistent, repeatable, and comprehensive inspections while minimizing the need for direct human involvement in hazardous conditions.

1.5 Software/Control Systems Requirements

The APES platform is a highly integrated system composed of multiple hardware and software subsystems working together to achieve autonomous inspection of hazardous pit environments. Successful operation requires seamless communication between components responsible for robot control, data processing, simulation, visualization, and operator interaction. To support this, the system leverages a suite of specialized software tools unified through a cohesive architecture, with the goal of presenting all critical functionality within a single, intuitive operator interface. Ease of use is the priority of the design, as experience has shown that overly complex systems are unlikely to be adopted. If operators do not like using systems, they often simply do not use them, opting for more familiar manual operation methods. Developing a streamlined and accessible interface is therefore essential to enabling real-world deployment and overcoming barriers to technology transition.

At the core of the system architecture is the Robot Operating System (ROS 2), which provides the primary communication framework for coordinating robotic subsystems and sensor integration. Visualization and operator feedback are supported through the Unity engine, while Deep Lynx serves as the backbone for digital twin (DT) development and metadata management. Advanced simulation and machine learning capabilities are enabled through the NVIDIA Isaac platform, allowing for simulated testing, training, and performance optimization in virtual environments that can then be transitioned to physical components. A comprehensive digital twin of both the system and the pit environment is maintained to support real-time monitoring as well as predictive modeling, giving operators insight into system performance over time, potential wear points, and maintenance needs.

Operational awareness is further enhanced through the use of mixed reality (MR) headsets, which provide immersive, real-time visualization of system data and environment mapping. Control inputs shall utilize standard devices such as a mouse and keyboard and game controller to provide an intuitive method of interacting with the software. The system is designed to maintain accurate positioning within the confined pit space, with built-in path planning and collision avoidance capabilities that incorporate operator feedback. All communications and control signals are transmitted through tethered connections to ensure reliability within the signal-restricted environment, eliminating dependence on wireless systems.

1.6 Additional Interest Areas

Engagement with operators and review of Hanford Site documentation have identified several additional features and design considerations that are desirable for a fully deployable system. These features were not part of the original project scope; therefore, they are not fully addressed within the current demonstration effort. However, they provide important guidance for future development and may be incorporated as the system matures, should time and budget allow. The APES platform has been designed with an open end-of-arm tooling interface to support future expansion, enabling additional tools and capabilities to be integrated without requiring significant redesign of the core system.

Several functional enhancements have been identified as valuable for long-term deployment, particularly those that would allow interaction with in-pit infrastructure and debris. These include the ability to interface with and manipulate components such as valves, hoses, tubing, and miscellaneous obstructions, as well as the potential to install devices, such as test plugs (e.g., 2- or 3-in. Griptight test plugs), or to deploy cleaning systems, such as spray-based or laser ablation end effectors. While some of these concepts may be partially explored in a prototype environment, full implementation is expected to occur in later design iterations of a site-ready fully deployable system.

In addition to functional considerations, there are material and durability requirements that would apply to a deployable system at Hanford. These include adherence to site-specific standards such as TFC-ENG-STD-34, "Standard for the Selection of Non-Metallic Materials in Contact With Tank Waste," as well as a preference for stainless-steel construction where practical. Operational-durability expectations, such as the ability to withstand a 3-ft drop, have also been identified. Material specifications are being reviewed to meet these criteria, but a drop test of the specialty prototype equipment cannot be supported under this project. Similarly, the use of site-approved fluids, such as Quintolubric 888-46 for fluid-controlled components, represents a requirement for deployment that will need to be incorporated into later-stage designs if later tasks require fluid-controlled components. At this time no fluid-controlled components are on the robots.

Collectively, these considerations define a pathway for transitioning the system from a demonstration platform to a fully deployable solution that meets Hanford Site operational standards and expectations. This progression also establishes a foundation for future commercialization, where partnerships with industry can support scaling and deployment across the Hanford waste-tank farms and potentially other DOE environmental-management sites.

2 Design Concepting and Research

2.1 System Support Structure/Platform and Housing

During Year 1, the design philosophy centered on developing a large, integrated support structure capable of spanning the full width of the tank pits. This approach sought to provide a single platform to organize, house, and deploy multiple robotic subsystems. Concepts such as adjustable steel-frame structures and vehicle-based deployment systems were explored to balance configurability, system integration, and operator usability. While these approaches offered advantages in terms of robustness and centralized support, they introduced significant challenges related to deployment complexity, cost, and site constraints.

As the design matured and further engagement with Hanford Site stakeholders occurred, the limitations of large-scale support systems became more apparent. In particular, systems that spanned the full pit width would require heavy equipment, such as cranes, for installation, positioning, and removal. Although the use of such equipment is feasible within the tank farms, it introduces substantial overhead in the form of work planning, safety approvals, and operational controls. These requirements add both time and cost to deployment and were identified by Hanford personnel as undesirable.

In response, the Year 2 design philosophy shifted toward a significantly smaller, more modular approach that eliminates the need for heavy equipment. Rather than spanning the entire pit, the updated system interfaces directly with the pit from above via the existing viewport. This change reduces system footprint and simplifies deployment logistics.

The resulting concept utilizes a dual-stage track system that can be assembled in place using minimal tools by a small team of two to three operators. The design prioritizes ease of use, rapid deployment, and minimal site disruption. The system is being developed with the goal that two operators can fully assemble and position it in under 2 hours without the need for cranes or other heavy machinery. This evolution reflects a broader shift from maximizing structural integration toward emphasizing deployability, operational efficiency, and alignment with site constraints.

The design for the entrance and extraction system centers around a main mounting place that interfaces with the view port flanges. The plate is meant to mount onto the flange at any angle in a 360-degree rotation, this will allow the robot to be positioned around other obstacles that exist on the top of the tanks. From the flange, truss stabilizing arms and a vertical rack and rail bed are bolted to the mounting plate. The trusses use a quick-locking design similar to that used in stage lighting setups. The mounting flange along with truss supports allow the load of the system to be distributed across the top of the pit block. Additional riggers can be tied into this structure for extra distribution and support. With this design philosophy, the main structure of the assembly can be built in place in under 15 minutes without the need for any special equipment or lifting vehicles.

A cart is then installed on the vertical rack and rail bed that carries a motor for movement of the secondary z mast stage (Figure 1). This stage consists of a square tube with an additional rack and rail system mounted onto it. The primary stage can move up and down independent of the secondary stage movement. The secondary stage is designed for a 2.5-ft length to introduce the robot below the full concrete cover block. The arm will mount to the end of the secondary-stage tube, and the cables that control the arm will be passed through the tube and up to the control actuator system that will remain above the pit during operation (Figure 2). The dual-stage system allows operators to position the assembly at a reasonable working height while maintaining a height to allow the arm to be introduced fully into the pit.

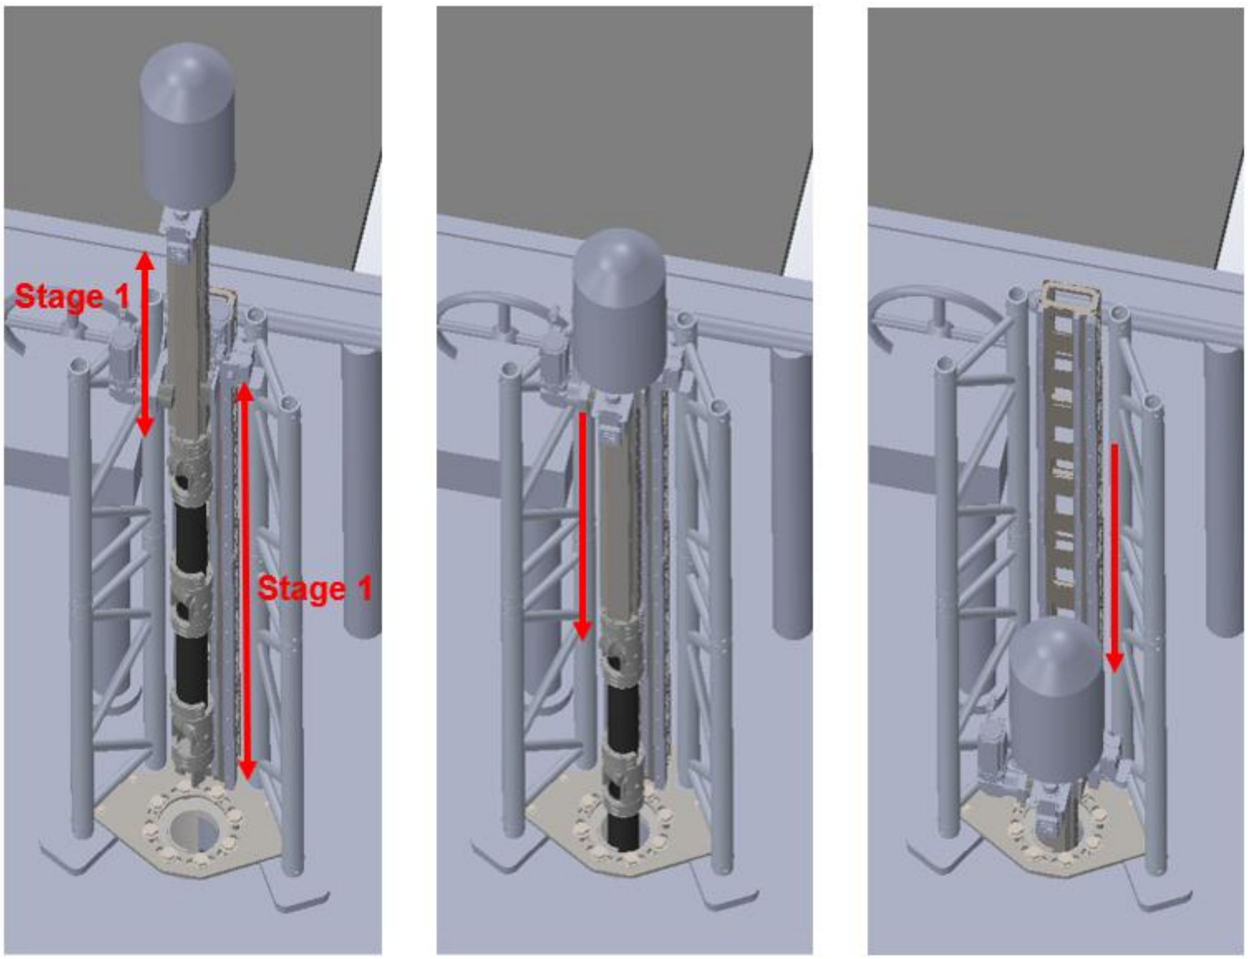

Figure 1. Build in place dual stage z mast design.

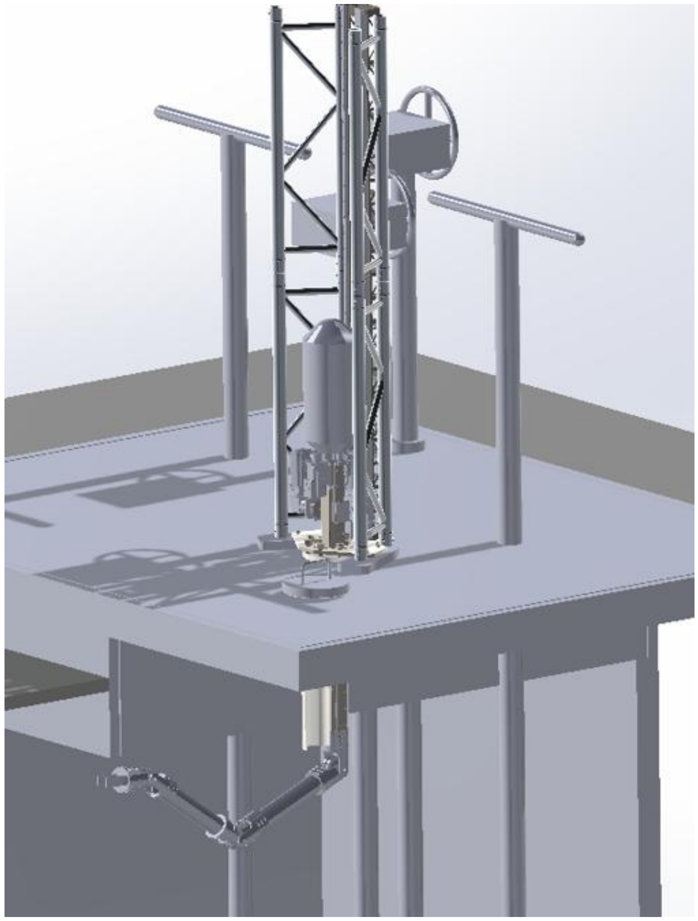

Figure 2. Support structure with robot.

2.2 Electric Truck as Transportation and Power Delivery for Robots

Among the components of the multi-robot inspection framework, the Ford F-150 Lightning electric truck plays a dual role: supplying electrical power and serving as an onboard computing resource. Bench testing confirms that the vehicle's output remained consistent and well-regulated, sustaining operations up to a threshold of 7.4 kW without instability. When the entire system was run under simulated field conditions, the energy drawn from the truck's battery was surprisingly low, pointing to a generous operational window and adequate reserves for driving between sites. Taken together, these findings suggest that modern electric vehicles are well-positioned to replace conventional generators in off-grid and logistically challenging deployments, offering a cleaner for alternative short term remote robotic deployments. Figure 3 presents a conceptual rendering of the truck's configuration carrying the system payload. Further technical details benchmarking power-output quality and capacity can be found in Ref. [1].

Figure 3. Conceptual design of the electric truck-powered system.

2.3 Custom Robotic Arm

A joint effort between INL and Florida International University (FIU) is underway to engineer a purpose-built robotic arm tailored for confined-space inspection tasks. The device is designed with several articulated joints and must be compact enough to pass through a 4-in. access port, after which it navigates around internal obstructions to carry out a range of assessment objectives. Its sensor suite includes both a depth 3D imaging system and a standard monocular camera to support visual surveys and spatial mapping. Beyond vision, the platform integrates LiDAR for environmental awareness, alongside instruments measuring radiation levels, temperature, pressure, and humidity. FIU's development process yielded several prototype configurations, two of which are highlighted below.

The first concept, internally called the scissor-lift design (Figure 4), achieves movement through a combination of vertical descent, lateral translation, and axial rotation once deployed inside the confined space. A vertically oriented linear actuator, aligned with the system's central support column, drives the scissor mechanism through its extension and retraction cycle. A sensor-carrying carriage is mounted to this assembly, moving in tandem as the mechanism operates.

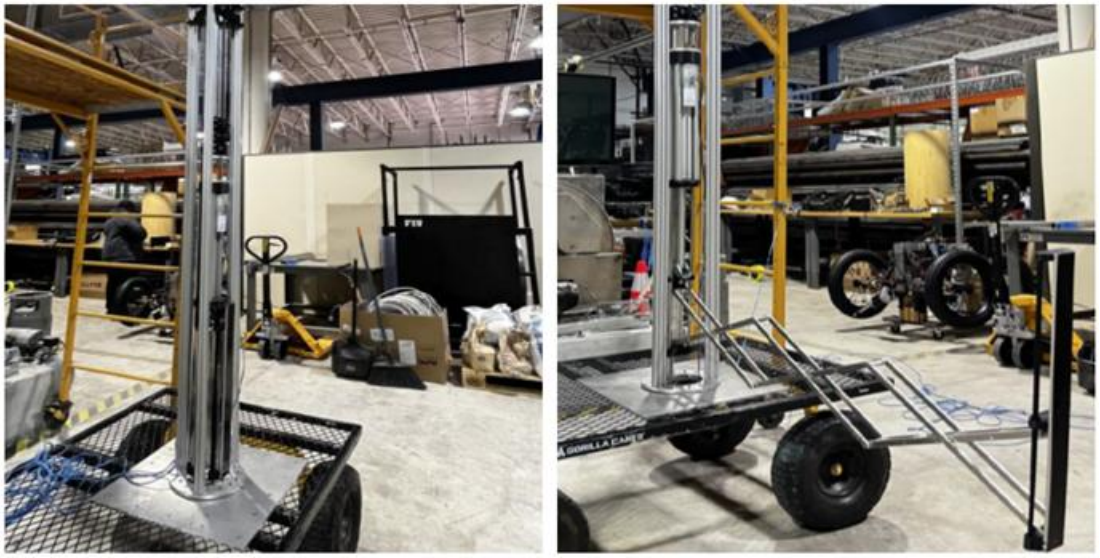

Figure 4. Scissor-lift arm design, retracted on right, extended on left.

The second prototype takes a more-conventional approach to robotic-arm architecture and is known as the cable-driven arm (Figure 5). It features a primary rotational axis complemented by a series of joints that can be added or removed to adjust the arm's reach and range of motion as needed. Rather than embedding motors at each joint, actuation is achieved remotely through a network of cables and pulleys routed back to the external support structure, where linear motors apply the necessary force. This arrangement enables considerable maneuverability without sacrificing the structural capacity needed to support heavy end-effector payloads.

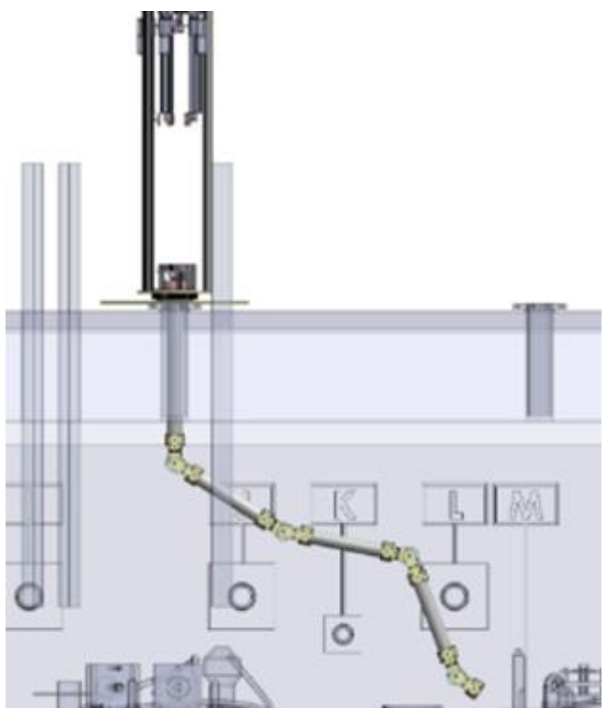

Figure 5. Cable-driven arm design concept by the FIU team.

After evaluating both functional prototypes, the cable-driven configuration was selected for continued development toward the final system. Its superior adaptability in navigating tight or irregular spaces, combined with its ability to handle substantial loads, made it the more-capable and practical choice for real-world deployment.

2.4 Additional FIU Mini-Crawler

Once the decision was made to pursue a mini-crawler platform for the project, FIU identified an existing design, a small crawler system that could be adapted for deployment within the tank pits. The platform is designed with modularity in mind, allowing for straightforward integration of various sensing payloads, including cameras, LiDAR, and other compact environmental sensors. It has dual tracks on each side, with independent controls that enable tight turning and enhanced maneuverability in confined spaces.

The crawler is primarily constructed from 3D-printed components, making it both cost-effective and easy to manufacture and replace (Figure 6). Power and communications are supplied through a tethered connection, which also provides a reliable means of retrieval by physically extracting the robot if necessary. The crawler's design is particularly advantageous in contaminated environments, where the system's low cost allows it to be treated as a disposable asset if required, with rapid reproduction of replacement units. Additionally, the control architecture is built on robotic operating system 2 (ROS 2), ensuring compatibility and seamless integration with the broader robotic system and enabling coordinated operation with the other robots.

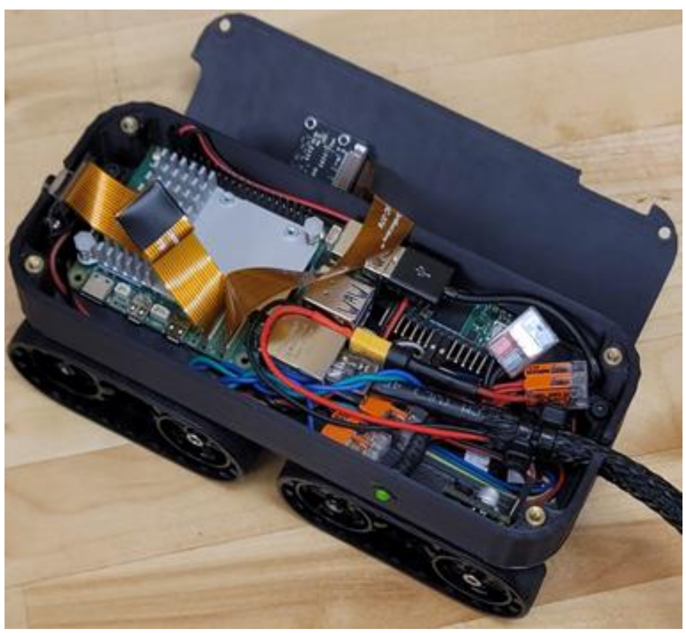

Figure 6. Mini-crawler developed by FIU.

2.5 Nexxis System Manual Controls

The Nexxis equipment comes as either a fully ROS 2-compatible or a standalone system. A hardshell case with an enclosed computer, keyboard, and screen makes up what is called the Juicebox controller. This provides a rugged field deployable operator interface and control system for all of the Nexxis equipment (Figure 7). Out of the box there are configurable options for deploying the crawler, Argus module, and PTZ camera, either together or individually. The Juicebox offers connectivity through ROS 2 to the rest of the control and visualization systems, meaning ROS packages can be created and sent to other systems. The following sections outline the various components in the Nexxis system.

Figure 7. Nexis system components and Juicebox controller.

2.6 Mini Robotic Crawler

Developed for deployment in confined and difficult-to-access locations, such as underground vaults, the compact robotic crawler is a capable and adaptable inspection platform. Locomotion is handled by a twin-track drivetrain, which provides reliable traction and steering across irregular terrain (Figure 8). Onboard lighting and imaging sensors work together to maintain visual clarity and ensure consistent data acquisition, even in poorly lit, narrow spaces. The system supports two distinct physical arrangements: a side-by-side track layout optimized for agile maneuvering and an in-line track configuration that allows the unit to pass through access openings as narrow as 4 in. Rather than being a custom-built solution, the crawler is a commercially available unit sourced from Nexxis, which simplifies component sourcing, reduces downtime, and keeps routine maintenance manageable. The platform is also engineered for parallel operation alongside a separate robotic arm entering through an alternate access point although it is equally capable of carrying out assessments on its own, giving operators the flexibility to scale the inspection approach based on site conditions and task requirements.

Figure 8. Nexxis robotic crawler.

2.7 Pan-Tilt-Zoom Camera

Among the supporting components of the inspection system is a commercially sourced PTZ camera, a relatively straightforward piece of equipment that nonetheless plays a critical role in overall situational awareness (Figure 9). Its primary contribution is providing an external, overhead perspective of the robotic assets operating within the pit, a vantage point that proves far more informative than relying solely on robot-mounted imaging. The unit also incorporates adjustable supplemental lighting, giving operators the ability to fine-tune brightness levels and minimize the visual interference caused by glare or shadows in the enclosed space. Operationally, the PTZ camera is flexible enough to be used during preliminary site assessments before any robots are deployed or positioned through a dedicated access port to run in parallel with active robotic operations.

The value of this external perspective became apparent through experience with prior robotic deployments, where first-person cameras fixed directly to a robot consistently fell short in conveying spatial relationships and depth to the person at the controls. A broader, third-person field of view addresses this limitation directly, giving operators a comprehensive understanding of the pit environment, clearer awareness of where each robot is positioned, and the context needed to intervene with manual control whenever the situation demands it.

Figure 9. Nexxis PTZ scope camera.

2.8 Nexxis Argus LiDAR Mapping System

Nexxis's Argus 3D SLAM is an industrial-grade platform built around simultaneous localization and mapping (SLAM) principles, purpose-designed for environments where autonomous inspection and data gathering are required. At its core, the system leverages LiDAR sensing to continuously generate high-fidelity 3D point clouds, allowing it to orient itself and operate reliably in settings where GPS signals are unavailable or where hazardous conditions make human entry impractical. One of its key contributions to robotic systems is the enhancement of machine vision capabilities that equip robots with the spatial awareness needed for real-time positioning and self-directed navigation without external guidance.

Beyond navigation, the platform supports precise positional referencing for robotic arms and attached tooling, which translate directly into more-consistent inspection outcomes and the ability to geotag collected data for future repeatability. An integrated path-planning module further extends its utility by autonomously guiding robotic assets through structurally complex or obstructed spaces, ultimately raising the bar for safety, operational efficiency, and the integrity of acquired data within asset management programs [2].

For this project, INL commissioned a tailored variant of the Argus platform, scaled to fit within the pit's access viewport (Figure 10). This custom configuration is being evaluated in both a standalone capacity and as a sensor module mounted directly on the robotic arm, with the goal of advancing the system's overall mapping fidelity and degree of autonomous functionality.

Figure 10. Nexxis Argus, custom-designed for s 4-in. view port.

3 Software Architecture and Design.

3.1 Software Architecture, Data Collection and Processing

The system's digital-twin framework is distributed across two computing platforms: a standard PC and an Nvidia Orin Jetson. The platforms work in concert to handle the demands of real-time robotic inspection. INL's own open-source data-management ecosystem, DeepLynx, forms the backbone of this architecture. Workflow orchestration is handled through Apache Airflow, packaged within a component called DeepLynx Run, while artificial-intelligence (AI) model development for the robotic platforms takes place within Nvidia Isaac Lab [3], with model versioning and tracking managed via the open-source tool MLFlow. The Nvidia Jetson serves as the primary interface between the digital twin and the physical hardware, maintaining direct communication with all deployed robots and sensors through a ROS 2 data bridge (Figure 11).

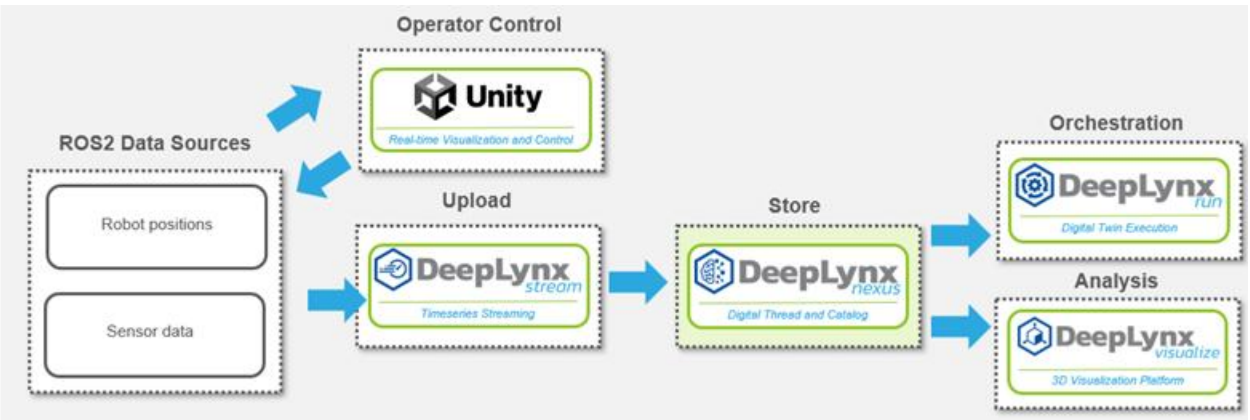

Figure 11. Digital-twin architecture.

To support simulation-based development and future autonomy training, INL constructed a detailed Isaac Sim/Isaac Lab environment that goes well beyond basic kinematic modeling. The virtual space incorporates accurate representative pit geometry, sensor positioning, and lighting conditions representative of the dark, cluttered operational setting anticipated at Hanford. Given that high-quality visual inspection is the system's foremost objective, realistic illumination was treated as a functional simulation requirement rather than an aesthetic consideration. Notably, the interplay between the robotic arm's end-effector lighting and the PTZ camera's illumination is actively modeled to assess whether the combined output produces inspection-grade imagery and to generate training data suitable for eventual sim-to-real transfer.

The first step in Nvidia Isaac Lab simulation efforts was the establishment of a functional simulation environment within the program, including importing models and configuring test scenarios. A computer-assisted design (CAD) model of a laboratory space was developed to enable a simulated scan of the room, providing a controlled digital environment for initial testing (Figure12). The Boston Dynamics Spot robot was selected as the robotic test platform due to out-of-the-box support within Isaac Sim and the ability to modify it with a backpack that housed the needed sensors. Within Isaac Sim, a simulated Livox Mid-360 LiDAR sensor was integrated to replicate the real-world sensing configuration. Using Isaac ROS bridges and simulation tools, a scanning routine was executed in the virtual environment to generate a representative point cloud of the lab space. This simulated workflow was then transferred to the physical Spot robot, which performed an equivalent LiDAR scan of the real environment. While the robot control implementation remained relatively simplified and rudimentary at this stage, the experiment successfully demonstrated a key capability of the project: interoperability between physical robots and Isaac ROS, and the feasibility of simulation to real-world (sim-to-real) transfer.

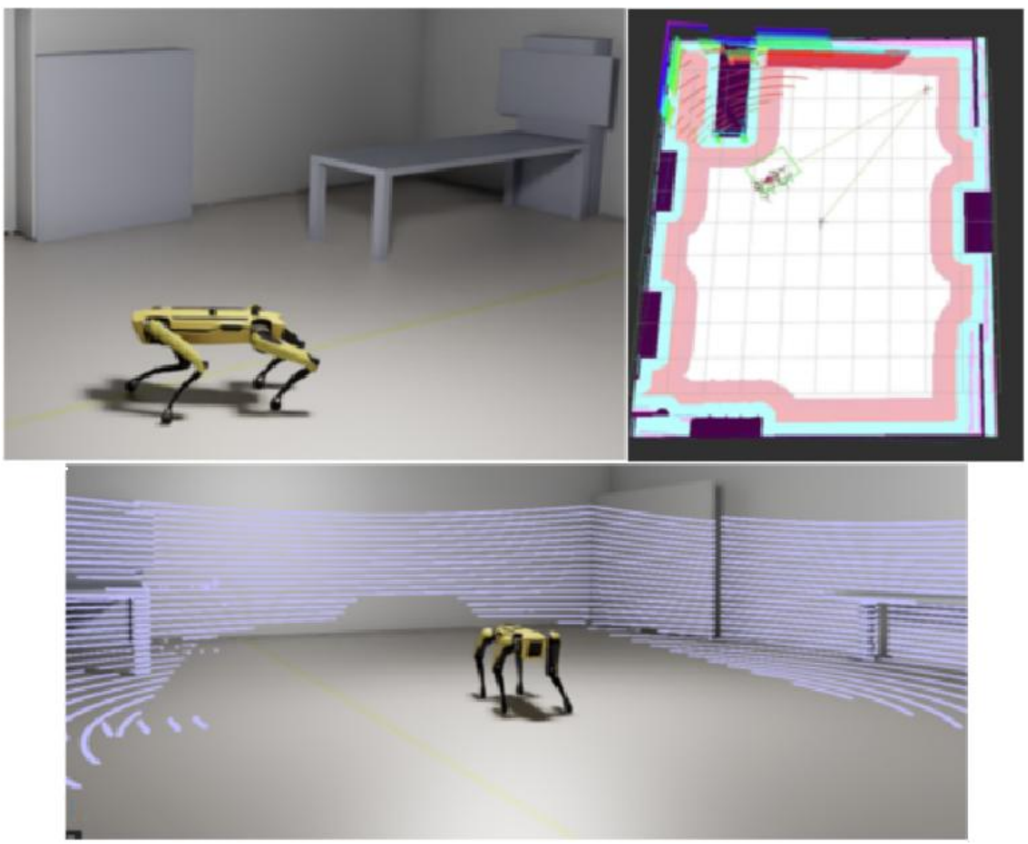

Figure 12. Initial Isaac Sim LiDAR room-scan test.

Isaac Lab machine learning is being improved for the robotic components of the system. The codebase is organized around the Isaac Lab project repository with separate asset repositories for the manipulator, the tank, and the PTZ camera models. The main workspace contains the environment configuration, custom manager terms, and test scripts. The three principal assets are a nuclear waste-tank model, the manipulator arm, and the PTZ camera. The tank is treated as a static scene asset with collision geometry while the arm and PTZ are modeled as articulated systems imported from universal robot description files (URDFs).

URDF files are generated from existing project CAD files from INL and FIU. To create a functional robot model usable in simulations, the base CAD geometry must first be converted into a URDF file. A Solidworks extension, called Solidworks to URDF Exporter, is used for this process. The exporter tool steps through the process of creating axes and coordinate systems for each joint of the robot. The tool configures what type of movement each joint can achieve as well as physical, inertial, visual, and collision options. Physical, material and inertial properties are automatically exported from the CAD solid models. When the export function is complete, a full folder is created with URDF options and generated meshes. This process was completed for FIU’s robotic arm, the Yahboom robot, PTZ camera, and Spot robot. These were then combined with static CAD file meshes of the representative pit and brought into Isaac ROS.

Reliable data collection underpins the entire inspection workflow, and each robotic platform carries a tailored sensor array to meet this need. Spatial mapping and high-resolution imaging are handled by the ZED X Mini RGB-D camera while the Livox Mid-360 LiDAR provides detailed environmental scanning for object detection and terrain mapping. Monocular cameras supplement these with continuous visual recordings to support navigation and operator awareness. Specialized instrument radiation detectors, pressure sensors, and temperature sensors round out the suite, capturing the environmental parameters essential for safe operation in potentially hazardous conditions. All sensor feeds are routed to the Nvidia Jetson over Ethernet and published across the ROS 2 network.

The ROS 2 publish-subscribe model governs how information flows between the system’s many components. Each sensor broadcasts its output to designated topics, which other nodes, including control systems and data pipelines, monitor and consume as needed. Nvidia Isaac can pre-simulate these signals for model development and test deployments (Figure 13) and then be fed directly to control the physical components. DeepLynx Stream ingests this ROS 2 data, passing it through DeepLynx Nexus before surfacing it in DeepLynx Visualize, a Unity-based interface that provides operators with a live, interactive view of the entire operation. This front end supports real-time data rendering, manual control inputs, collision awareness, and informed decision-making, all driven by continuous sensor feeds from the field.

Figure 13. (left) Robotic arm tracked via PTZ camera with RTX path-traced lighting and (right) multi-environment simulation with randomized initial configurations.

Collectively, the layered integration of sensing hardware, distributed processing, and interactive visualization creates a coherent and resilient system capable of operating across complex and hazardous environments. ROS2 serves as the connective tissue binding multiple robotic units into a unified operational framework, while the digital twin layer ensures that collected data is not only captured but actively used to drive smarter, safer inspection outcomes.

3.2 Cameras, Visualization, and Operator Interfaces

The operator interface is being jointly developed by INL and Brigham Young University (BYU), with an emphasis on creating a system that supports effective robotic operation, situational awareness, and post-inspection analysis. Several key design challenges have guided development, including managing diverse sensor inputs and their representations, ensuring operator visibility into and trust of robotic actions, and enabling efficient interpretation of collected inspection data.

BYU is developing a data pipeline from ROS 2 to Unity for the robotic-prototype interface, which will later integrate with DeepLynx Visualize. This interface is designed around a dual-screen layout: a control-oriented display (Figure 14) and a visualization-focused display (Figure 15). The control screen enables operators to actively manage robotic platforms, including both the FIU robotic arm and the Yahboom crawler, with an ability to switch between systems. It also provides real-time positional awareness of robots within their environment. The interface supports both autonomous and manual operation modes, allowing operators to intervene when autonomy is insufficient or when specific areas require closer inspection. Planned robotic movements during autonomous operation are also presented to improve operator understanding and trust.

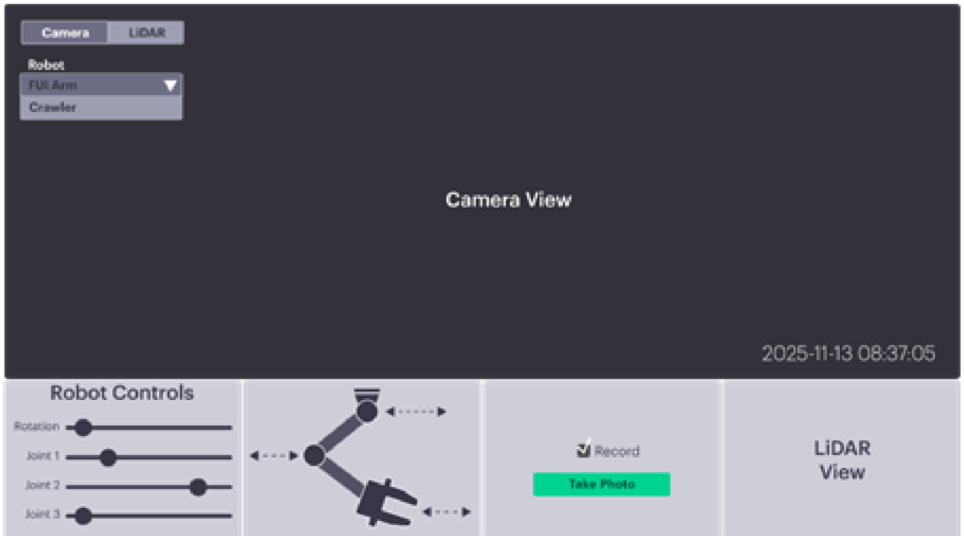

Figure 14. Mockup of operator interface.

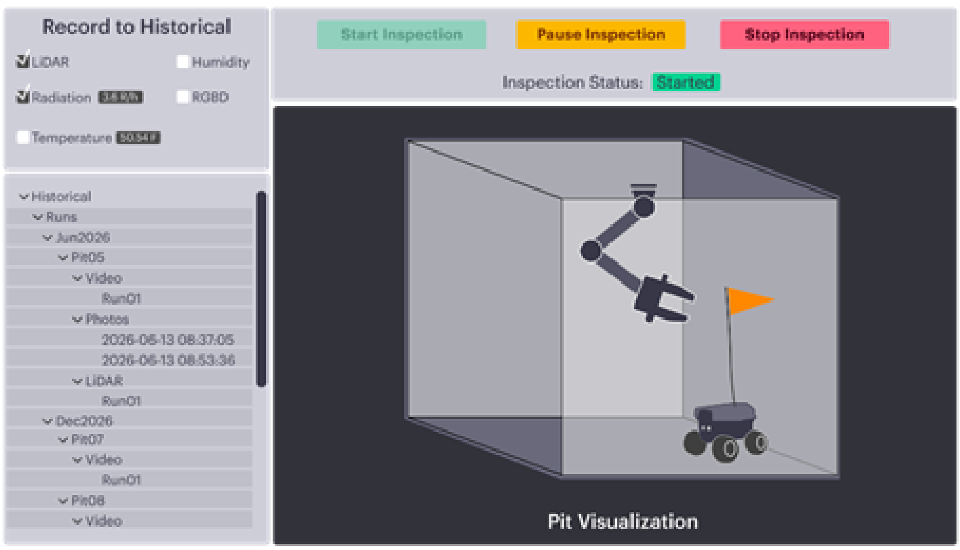

Figure 15. Mockup of scan interface.

The visualization screen provides a real-time view of the inspection environment through live camera feeds, including ongoing integration of the ZED X camera and LiDAR systems. Sensor data layers such as LiDAR, radiation, humidity, temperature, and RGB-D can be selectively enabled or disabled, allowing operators to interpret both visual and non-visual data (e.g., radiation signatures) in context. These overlays enhance situational awareness by combining multiple data sources into a unified display. The visualization interface will also include controls for recording, pausing, and stopping data collection, enabling the capture of images and video for later analysis, as well as playback functionality for reviewing recorded inspection runs.

To support spatial awareness, the interface incorporates a 3D representation of the inspection environment. This includes a model of the pit generated from RGB-D data, along with a dynamic representation of the robot's position within that space. While LiDAR data are collected during inspections and used for SLAM, they produce point clouds that are not immediately suitable for real-time mesh-based visualization in Unity. Converting these point clouds into usable 3D models requires post-processing. To address this limitation, INL is investigating the use of Nvidia NVBlox in combination with the ZED X Mini RGB-D camera to generate real-time 3D reconstructions. Because RGB-D sensing is passive and depends on external illumination, lighting must be treated as part of the sensing architecture, rather than as an accessory. The end-effector light provides close-range illumination at the inspection point, while the PTZ camera contributes a second adjustable light source from a separate viewpoint that can reduce harsh shadows and improve visibility on surfaces that would otherwise remain underlit. This cooperative lighting approach improves the quality of images, video, and reconstruction data for operators and is also important for autonomy development. This is because the robot system must be trained and evaluated on visual inputs that reflect realistic operating conditions in the dark pit environment.

Following inspections, data are made available within DeepLynx Visualize for detailed analysis. This post-inspection interface allows analysts to examine both visual and quantitative sensor data mapped onto the 3D model of the environment. Sensor readings can be overlaid spatially to identify patterns or anomalies, while captured images and video are associated with directional cones in the 3D space to indicate their origin and field of view. A playback feature reconstructs the robot's path through the environment alongside the data collected at each point in time, enabling analysts to better understand when and where specific signatures were detected.

The teams currently have data pipelines, visualizations, basic controls and streaming capabilities functional, and should be completing video playback shortly. While camera streaming was accomplished, issues with the Yahboom testing robot have prevented the BYU team from extensively testing and fully implementing sensor visualizations. BYU coordinated with FIU to gain access to the sensors on the prototype robot arm through virtual private network (Figure 16). Tasks to create a visualization prototype project have been met, and further refinement and features will be developed next year.

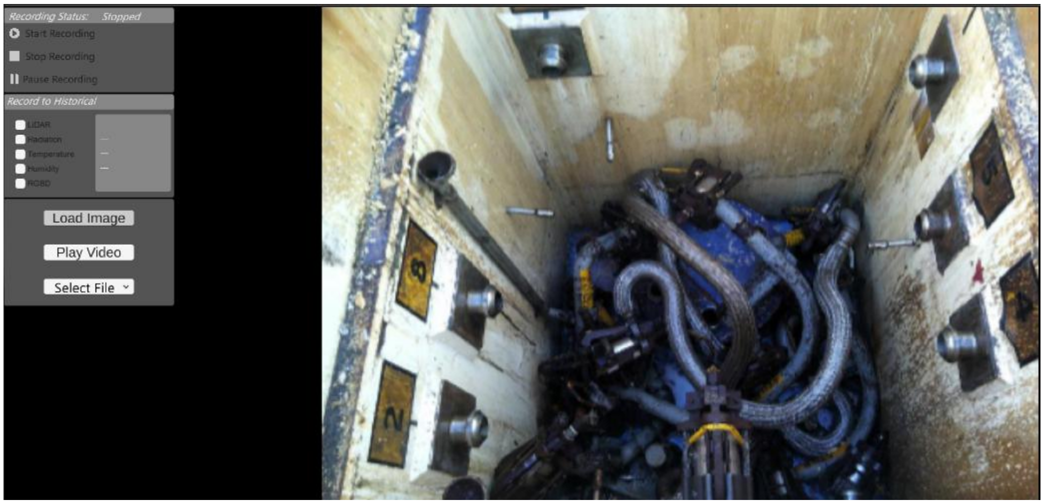

Figure 16. Video viewer interface.

3.3 Effectiveness of Multi-Robot Approach

The underground pits at the Hanford Site present a uniquely complex inspection challenge because each contains a dense arrangement of pumps, valves, piping, hoses, and miscellaneous equipment that collectively obstruct lines of sight and limit physical access. No single entry point offers a complete view of the interior, and threading one robot through this maze of obstacles while capturing full coverage would demand an exceptionally capable machine. Such a platform would need an extensive range of motion, considerable reach, and the structural strength to carry a meaningful sensor payload simultaneously. While engineering such a system is technically achievable, it comes at a steep cost. The financial burden of development alone is substantial, compounded by the need for specialized materials and custom components that are neither readily available nor inexpensive. Managing the motion of a highly articulated robot also introduces significant control complexity, requiring sophisticated monitoring to confirm that every segment of the mechanism is behaving as intended. Layered on top of this are the human factors; designing an interface through which operators can meaningfully interact with such a system is its own challenge, and building toward autonomy on that foundation becomes exponentially harder.

A variety of simpler, purpose-built robots offer a far more-practical path forward. For pit inspection, specifically, the straightforward jointed arm can be positioned to sweep across the upper portion of the pit, orienting its camera to capture overhead perspectives, while the compact crawler operates independently along the pit floor, gathering close-range readings and upward-facing imagery. Together, these two platforms can collectively achieve what a single highly complex robot would struggle to do alone, and they can do it with components that are either less expensive to develop or available directly from existing commercial suppliers. Each approach carries its own tradeoffs, but for this application, a coordinated multi-robot strategy represents a more-realistic and achievable route to success. The PTZ camera provides additional space lighting and a 3rd-person view of the robotic arm. With a coordinated operating system, the different robots can map the space, recognize where there may be gaps, and identify where the other robots will need to fill those gaps. Work is already underway with machine learning and route planning for the robotic arm to autonomously move around the pit with the PTZ camera down a secondary viewport automatically following the motion of the arm and providing extra light for it.

This line of thinking is not unique to robotics. Throughout human history, the most effective and widely adopted tools have tended to be those purpose-built for a narrow set of tasks. A carpenter reaches for a hammer, a drill, or a saw depending on the job at hand. The idea of merging all three into one device is more novelty than utility. The tradeoffs simply do not hold up. However, when robotic solutions are proposed, the persistent tendency is to default to a single all-encompassing machine. As autonomous systems grow more sophisticated and inter-robot communication and coordination become more reliable, the case for collaborative multi-robot architectures is gaining traction within the broader robotics-research community [4].

3.4 Challenges and Solutions

The advantages of deploying multiple robots in tandem do not come without corresponding complications. When the operating environment involves nuclear-waste inspection, and the platforms in question differ significantly in design and function, the engineering hurdles become particularly pronounced [5]. Achieving seamless collaboration between robots that rely on different hardware configurations and communication frameworks is a non-trivial problem, especially when each unit is simultaneously generating its own spatial data that must be reconciled into a single, coherent environmental model. Beyond inter-robot coordination, sustaining continuous field operations in a setting with restricted infrastructure access demands a dependable and lasting power strategy. Equally important is the need for onboard processing powerful enough to support autonomous decision-making in real time, all while presenting a steady stream of complex sensor data to human operators in a format that is immediate, interpretable, and actionable [5].

One challenge during Year 2 was the need to continue advancing simulation-to-real validation, mapping workflows and multi-robot coordination before the final APES robotic platforms were fully available for integrated testing. To address this, the project used commercial-off-the-shelf robotic systems, including the Boston Dynamics Spot and the Yahboom ROSmaster R2, as practical surrogate platforms for testing LiDAR simulation, SLAM, point-cloud registration, and coordinated inspection concepts. This approach allowed the team to continue validating core autonomy and digital-twin workflows on physical robots, rather than as candidate replacements for the final APES system.

Using these surrogate robotic platforms provided an effective bridge between software development and the eventual APES deployment configuration. The resulting demonstrations showed that robot-generated point clouds from different platforms and viewpoints could be aligned and merged into a unified map, supporting the broader project goal of using multiple specialized robotic assets to improve inspection completeness in cluttered Hanford pit environments. In this way, Spot and Yahboom served as enabling testbeds for sim-to-real progress, rather than as candidate replacements for the final APES system. The following case study provides a concrete example of how these surrogate platforms were used to validate multi-robot mapping and digital-twin workflows in a physical test environment.

4 Multi-Robot LiDAR Scan Case Study

As an initial demonstration of viability, a coordinated multi-robot mapping experiment, was carried out at INL's Visualization Lab. The exercise paired two fundamentally different platforms—a Spot quadruped and a Yahboom Rosmaster R2—each outfitted with a Livox Mid-360 LiDAR unit, with the shared objective of producing a complete 3D representation of the space. The experimental layout, illustrated in Figure 17, divided the lab into distinct operational zones: the Spot was confined to the area outside a blue-taped boundary but within a surrounding striped perimeter while the Yahboom Rosmaster R2 was assigned exclusively to the interior of the blue-taped zone, where it contended with a more cluttered obstacle arrangement. Both platforms executed their assigned mapping tasks concurrently. This zonal division mirrors the constraints expected in actual waste-tank pit inspections, where physical obstructions and spatial limitations will naturally restrict different robots to different portions of the environment, sometimes requiring sequential rather than simultaneous access to build a full picture.

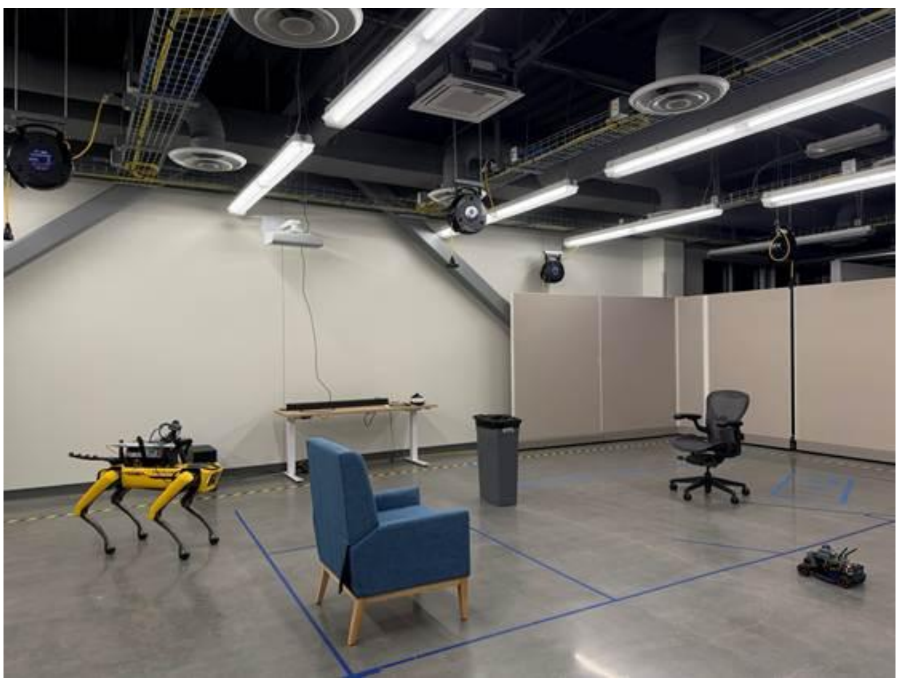

Figure 17. Multi-robot LiDAR mapping experiment setup.

Concurrent deployment of both platforms maximized spatial coverage across the lab. The Spot's agility made it well-suited to the more-open, unobstructed perimeter zone, while the Rosmaster R2's lower profile allowed it to handle the tighter, more-confined interior. This division of labor reduced the risk of coverage gaps that a single platform would likely encounter. Each robot performed continuous 360-degree scans using its Livox Mid-360 sensor, gathering high-density point-cloud data from perspectives unique to its zone. Map generation was powered by the Fast LiDAR-Inertial Odometry (FAST-LIO) algorithm [6,7], which enabled both robots to process incoming LiDAR data and construct accurate local maps in real time while in motion. The individual point cloud outputs prior to integration are shown in Figure 18.

Once data collection was complete, the two point clouds were brought together into a single unified map using CloudCompare [8] and the Point Cloud Library [9] to handle alignment and registration. The resulting combined dataset benefits from the overlapping coverage between zones, which reinforces map consistency and reduces the impact of any localized scanning gaps. Prior to merging, the Spot's dataset contained approximately 2.29 million points while the Rosmaster R2 contributed roughly 1.28 million. Following alignment, registration, and resampling, the consolidated point cloud was standardized to 3 million points. An isometric view of the final registered map is presented in Figure 19. This outcome demonstrates that a coordinated multi-robot approach can deliver both efficiency gains—through parallel operation—and improved map quality through complementary, overlapping coverage across distinct spatial zones.

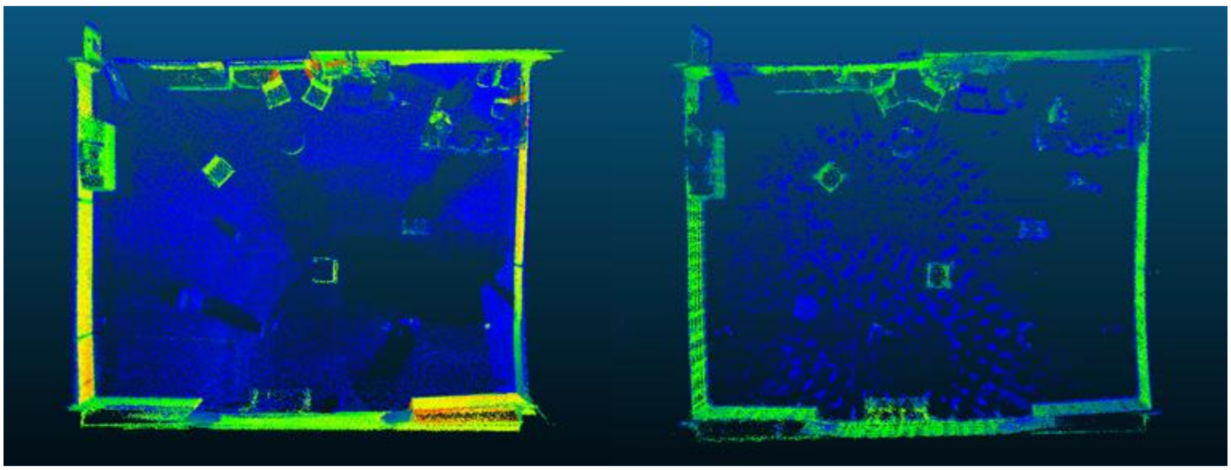

Figure 18. Spot (left) and Yahboom (right) point-cloud maps before aligning and merging.

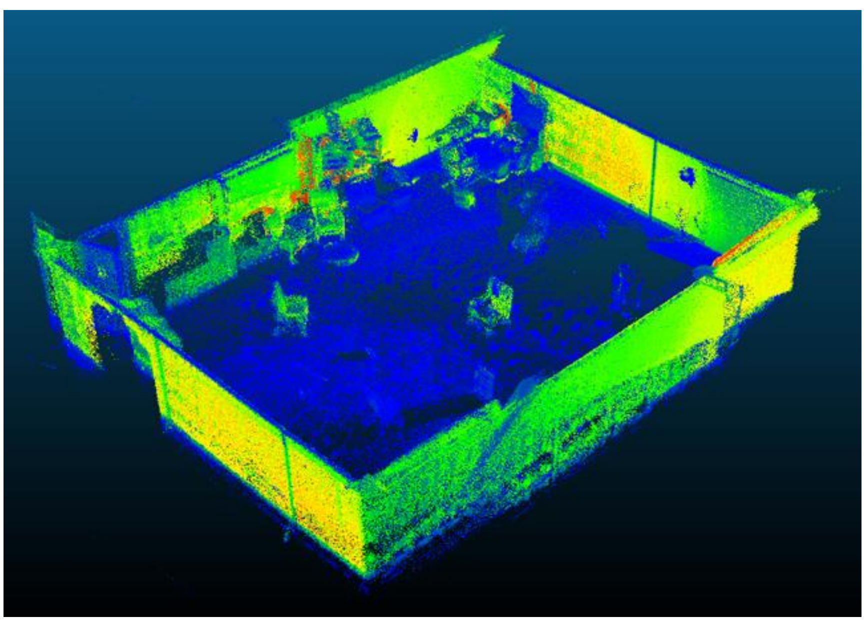

Figure 19. Isometric view of aligned and registered point cloud map.

5 Florida International University, Research Methodology and Technical Development of A Robotic Arm[1]

5.1 Track 1: Development of a Custom Small-Footprint Dexterous Manipulator

5.1.1 Architecture Down-Selection: Marionette Articulated vs. Scissor-Arm

Early in the program year, the team conducted a design trade study comparing two candidate robot architectures: a scissor-arm mechanism and a serial link-arm (Marionette) articulated manipulator. The scissor-arm concept offered simplicity and inherent structural stiffness, but it had difficulty maintaining meaningful reach inside the pit. It also offered limited workspace dexterity: the end-effector could not easily reorient to navigate around obstacles.

Marionette architecture was selected as the go-forward design. It consists of a base joint and four articulated revolute joints (Joints 1–4) plus a tool joint at the end-effector, connected by rigid aluminum links. When in zero position, the system collapses along the viewport axis and can be lowered vertically through the aperture. Once inside the pit, the joints articulate to sweep the camera through the interior space. Figure 20 shows the current physical prototype with component labels.

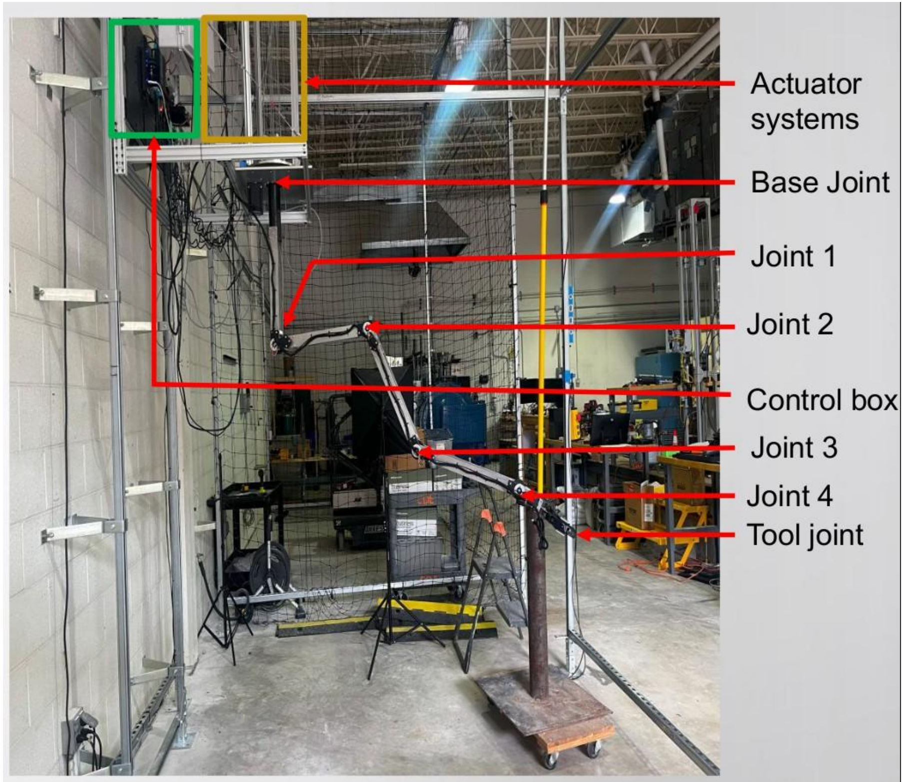

Figure 20. Current Marionette robot prototype. The system includes a base joint, Joints 1–4, a dedicated control enclosure, and a tool joint at the distal end. Linear actuator systems are mounted at the top of the deployment column.

5.1.2 Wire-Driven Linear Actuator Control

Each joint in the robot is actuated by a pair of antagonistic drive wires routed from the joint to a bank of **linear actuators** positioned at the top of the deployment column, external to the pit environment. This remote-actuation approach keeps all motorized and electronic components out of the radiation field, which is critical for system longevity and maintainability. A joint rotates as one wire is extended and its antagonist is retracted; the resulting motion is controlled by commanding differential displacement in the actuator pair.

This architecture introduces a coupling between wire tension and joint torque that is both a feature—it enables force-transparent compliance control—and a challenge, since wire elasticity and routing geometry produce nonlinear relationships between actuator position and joint angle. Managing these nonlinearities is a central motivation for the sensing infrastructure described next.

To provide reliable joint-state feedback, each revolute joint is instrumented with a custom-designed magnetic encoder board mounted directly at the joint pivot. The encoder uses a magnetic disc rotating against a fixed Hall-effect sensor array, giving a high-resolution, absolute-angle measurement that is immune to wire-stretch artifacts that would otherwise corrupt estimates derived purely from actuator position. Figure 21(a) shows the CAD rendering of the encoder integration at a representative joint, and Figure 21(b) shows a picture of the installed encoder on the physical robot.

5.1.3 Joint Encoders for Proprioceptive State Estimation

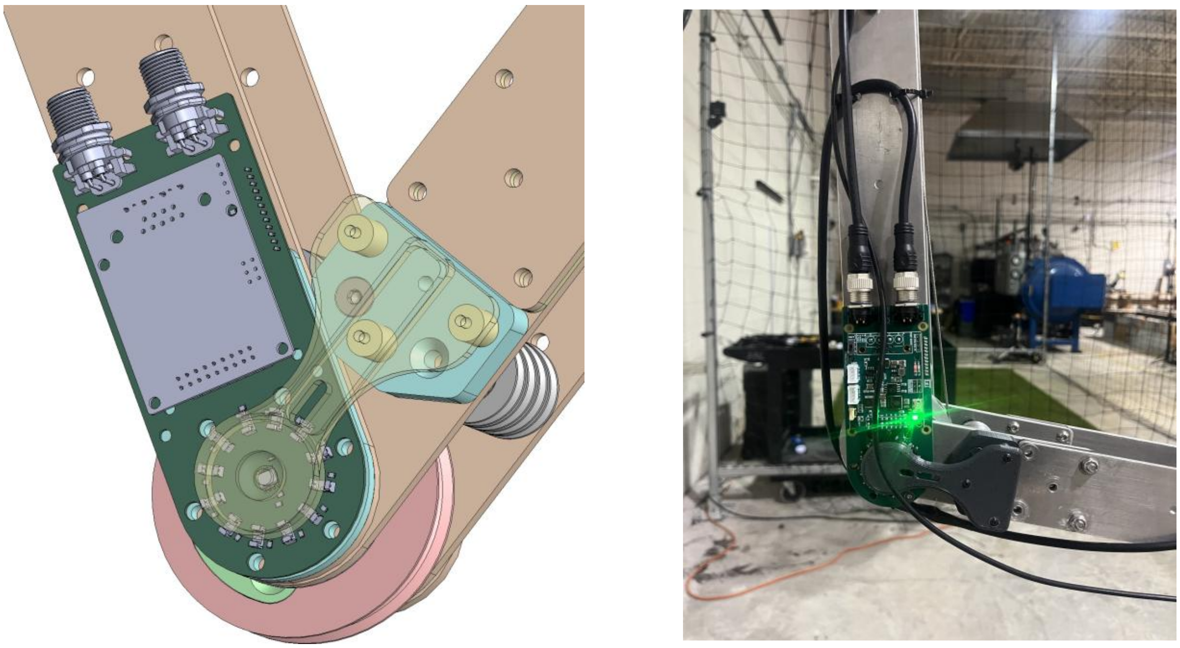

Figure 21. Joint encoder design and physical implementation. (a) CAD model of the joint encoder assembly integrated at the joint pivot. The magnetic disc (pink) rotates against the encoder PCB (green), which also carries the routing board for the drive-wire conduit termination. (b) Physical joint encoder installed on the robot arm. The green PCB with its indicator LED is visible, along with the conduit connectors that protect signal wiring routed along the arm.

5.1.4 Inline Load Cells for Wire Tension Sensing

In addition to joint angle, the team installed inline load cells into each drive wire at the actuator end of the routing path. Each load cell is a miniature S-beam type transducer rated for the expected wire-tension range, connected in series in the wire so that it measures the actual tensile load being transmitted to the joint. Figure 22 shows the load cell installation in the actuator bay.

Together, the joint encoders and load cells give the control system a complete proprioceptive picture of the robot state: the angular position of each joint (encoder angle), and how hard each wire is working to hold it there (load-cell tension). This dual-sensing architecture is the foundation for the closed-loop controller improvements planned for the next program year.

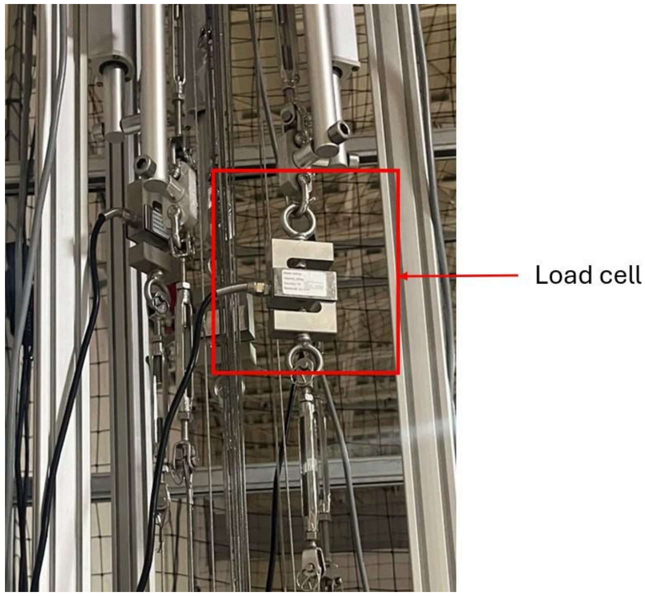

Figure 22. Inline load cell installed in the wire-driven actuation system. The S-beam transducer (highlighted in red) is mounted in series with the drive wire, providing real-time tension measurements for each joint's actuation channel.

5.1.5 Camera End-Effector and Live Scene Capture

A ZED X Mini stereo camera is currently mounted at the tool joint as the primary inspection end-effector. The ZED camera provides RGB video, depth maps, and a live pointcloud at up to 30 Hz. Figure 23 shows a representative live capture from the ZED camera during a laboratory inspection test, with the RGB view, color depth map, and the generated 3D point cloud all visible in the StereoLabs interface.

Figure 23. Live inspection scene captured by the ZED X Mini stereo camera at the robot's end-effector. The StereoLabs interface shows (left) the RGB view and color depth map, and (right) the live 3D point cloud with IMU orientation overlay. This data is streamed in real time to the ROS 2 planning stack.

The point cloud from the ZED camera feeds directly into the ROS 2 pipeline where it is registered into the planning scene coordinate frame and used both for real-time planning scene and for the inspection coverage analysis described in Track 2.

5.2 Track 2: Simulation Framework Development

5.2.1 Simulation-to-Planning Integration and Inspection Coverage Analysis in Nvidia Isaac Sim

A physics-accurate enhanced inspection framework of the robot was constructed in Nvidia Isaac Sim, incorporating the full kinematic description including joint limits, link geometry, and end-effector frame definitions. A representative Hanford waste-tank pit environment was modeled as the deployment scene: a concrete-walled rectangular enclosure populated with representative infrastructure elements—standpipes, fittings, structural brackets—to replicate the cluttered, occlusion-prone interior conditions that the physical robot will encounter. Figure 24 shows the simulated pit model with the Marionette articulated robot deployed through the view port.

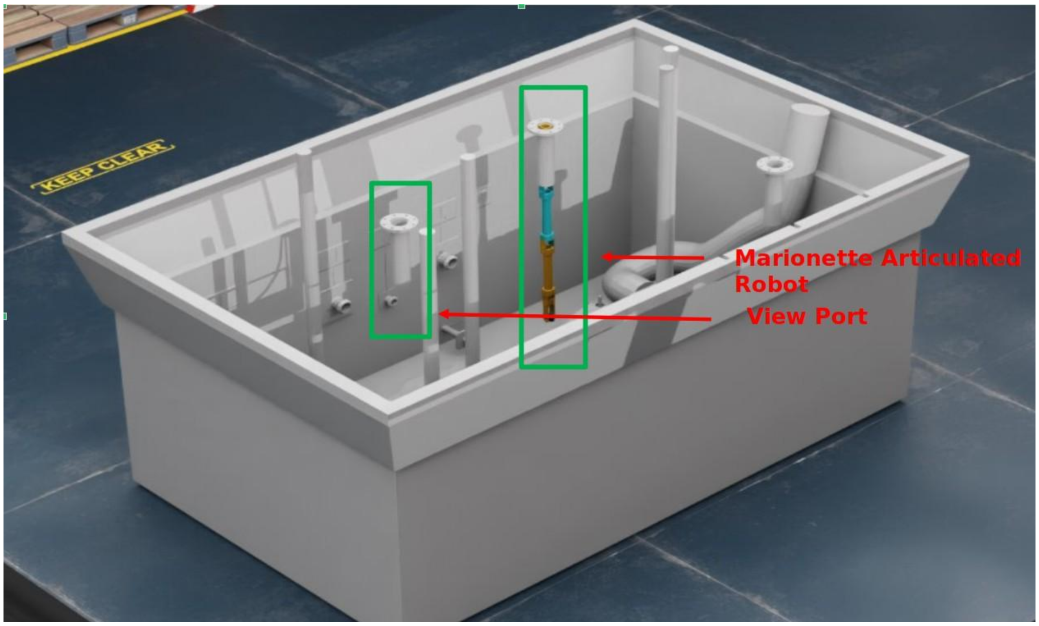

Figure 24. Nvidia Isaac Sim model of a representative Hanford waste-tank pit. The Marionette articulated robot (blue/gold) is deployed vertically through the view port, with standpipes and internal pit infrastructure visible. Green boxes highlight the two view port aperture locations available for deployment.

FIU created a simulation (Figure 24) that was synchronized with a MoveIt 2 motion-planning context running in ROS 2, ensuring that the planning stack operates with kinematic parameters—joint limits, collision geometry, and reference frames—that are identical to those used in the physical system. This coupling is essential for the eventual sim-to-real transfer of planned trajectories and, later, learned control policies. Figure 245 shows both the Isaac Sim rendering of the robot inside the pit environment and the corresponding MoveIt 2 planning scene representation.

Within the Year 2 simulation effort, FIU's primary contribution was the synchronization of the manipulator model between Isaac Sim and the MoveIt 2 planning stack in ROS 2. By keeping joint limits, collision geometry, reference frames, and end-effector kinematics consistent across planning and simulation, the team ensured that workspace characterization and inspection coverage analysis were performed on the same robot description that will support the physical system. This made the simulation environment directly useful for reachability assessment, collision-aware trajectory evaluation, and identification of coverage gaps rather than as a standalone visualization exercise.

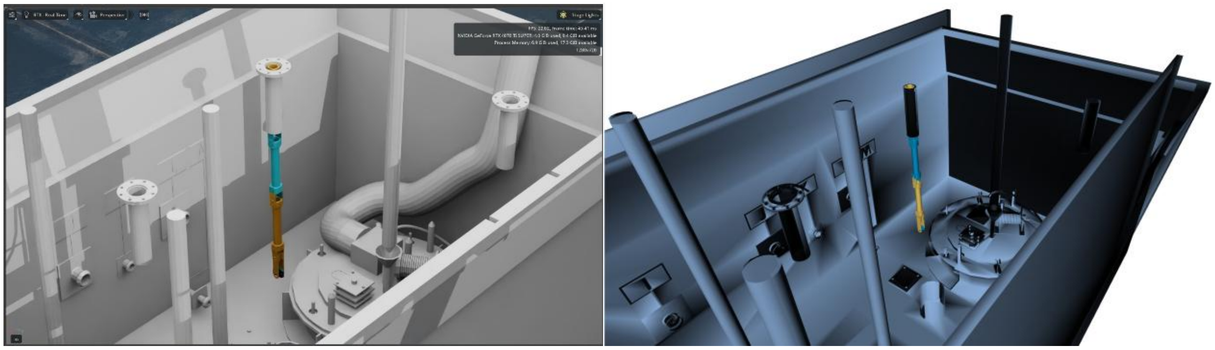

Figure 25. Simulation-to-planning integration. (a) Physics-accurate Isaac Sim environment with RTX real-time rendering. Nvidia Isaac Sim RTX rendering of the Marionette robot deployed inside the representative Hanford pit model. (b) MoveIt 2 planning scene using the same kinematic description and collision geometry. The robot's collision model and the pit geometry are represented identically in both environments, ensuring that trajectories planned in simulation are directly transferable to the physical system.

5.3 Workspace Characterization via Monte-Carlo Sampling

To quantify the kinematic reachability of the robot within the pit environment, the team implemented a Monte-Carlo forward kinematics (FK) sampling workflow. Joint configurations are randomly drawn from the full joint-limit range using MoveIt's built-in random state sampler; the FK solution maps each configuration to an end-effector position.

$\mathbf{p} = [x, y, z]^T$ in the planning frame. With $N = 3.2 \times 10^6$ samples, the resulting point cloud was voxelized at a resolution of $\Delta = 0.05$ m to estimate the reachable workspace volume:

$$V_{ws} = N_{occ} \cdot \Delta^3 \qquad (1)$$

where N_{occ} is the number of occupied voxels. The fraction of the pit volume physically reachable by the end-effector was then computed as:

$$\eta_{pit} = \frac{V_{ws\cap pit}}{V_{pit}} \qquad (2)$$

5.4 Inspection-Coverage Heatmap via Ray Casting

Beyond kinematic reachability, the team implemented an **inspection coverage analysis** to identify which surfaces of the pit interior can actually be seen by the camera along any feasible robot trajectory. For each camera pose along a candidate inspection path, surface points are classified as *seen* if they simultaneously satisfy three conditions: (1) they fall within the sensor's range limits $[r_{min}, r_{max}]$; (2) they lie within the camera's horizontal and vertical field of view; and (3) a ray cast from the camera origin to the surface point does not intersect any occupied voxel in the Octomap collision representation (i.e., no occlusion).

The resulting binary classification is aggregated into a color-coded 3D point cloud—green for observed surfaces, red for unobserved—and published as a ROS 2 PointCloud2 message, enabling direct overlay on the planning scene in RViz. This representation gives the operator and the planner immediate visual context for identifying coverage gaps.

6 Results and Discussion

6.1 Workspace Envelope

The Monte-Carlo FK analysis confirms that the robot's reachable workspace has an approximately axisymmetric lateral footprint about the deployment (view port) axis—a geometrically natural result for a serial manipulator deployed vertically through a central aperture.

The lateral reach extends to approximately ±1.03 m in both x and y, with a vertical depth range of $z \in [-1.51, 0.076]$ m referenced from the deployment origin (negative z corresponds to depth into the pit). Figure 26 shows the sampled end-effector positions under two conditions: the raw kinematic workspace with no collision checking (blue), and the collision-free workspace computed within the pit environment (green).

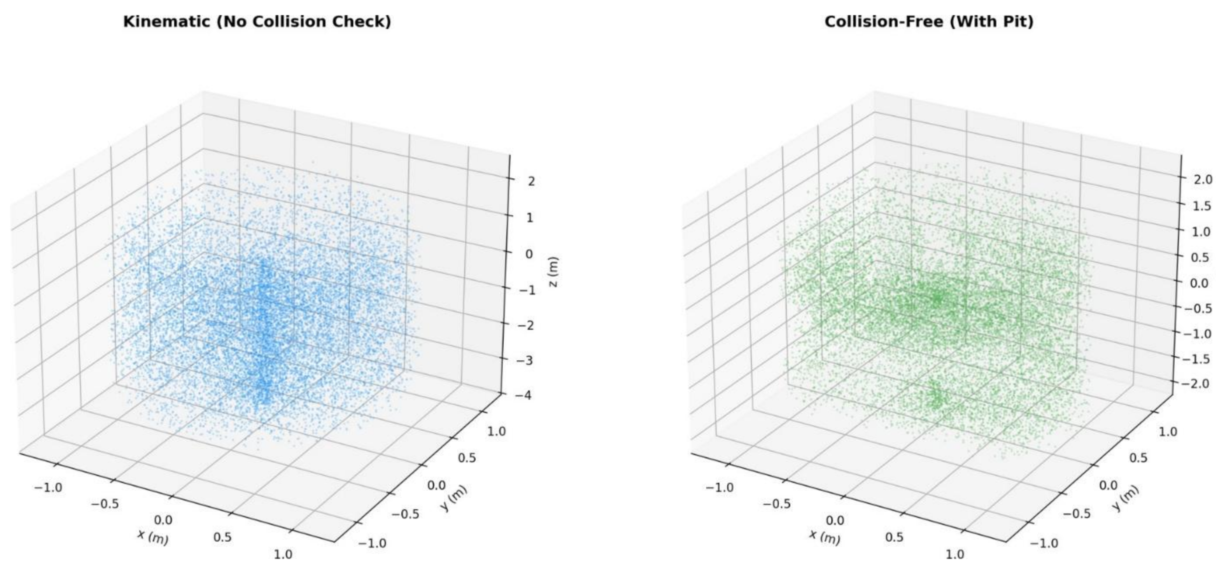

Figure 26. Monte-Carlo FK workspace sampling results. (Left) Kinematic workspace with no collision checking: the full reachable envelope of the end-effector sampled from $N = 3.2 \times 10^6$ random joint configurations. (Right) Collision-free workspace after filtering configurations that result in intersection with the pit geometry. The reduction in vertical extent reflects configurations eliminated by contact with pit walls, floor, and internal infrastructure.

For the full converged run for pure kinematic scenerio, the voxelized workspace volume was $V_{ws} = 21.96$ m^3. Intersecting this volume with the pit axis-aligned bounding box (V_{pit} = 39.50 m^3) yields a volumetric coverage fraction of 24.5%.

While 24.5% may appear modest, it is important to note that this figure reflects physical endpoint reachability: the end-effector actually occupying that volume. Inspection does not require physical presence at every point; the camera can observe structures well beyond the reachable envelope via line-of-sight across open volume. The effective visual-coverage fraction is therefore substantially higher and is what the inspection heatmap analysis directly measures.

6.2 Inspection Heatmap and Coverage Gaps

Figure 27(left) shows the inspection heatmap overlaid directly in the MoveIt planning scene, and Figure 27 (right) shows the standalone 3D view of the seen/unseen classification across the pit surface geometry.

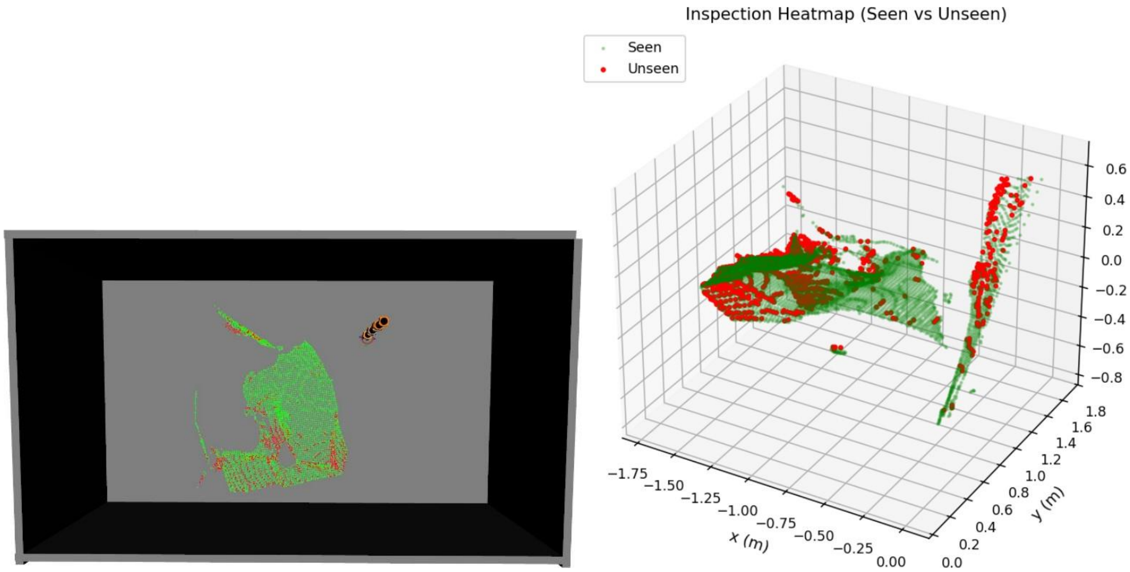

Figure 27. Inspection coverage analysis results. Green = observed; Red = unobserved. (Left) Coverage heatmap overlaid in the MoveIt 2 planning scene. Green points indicate surfaces observed by the simulated inspection trajectory; red points indicate unobserved regions. The robot end-effector is visible (orange) at its deployment position. (Right) Standalone 3D inspection heatmap. The dense green cluster represents the main floor area that is well-covered by the simulated trajectory; the red vertical stripe on the right corresponds to a structural occlusion boundary where the camera's line-of-sight is blocked.

Several important observations emerge from the heatmap. First, the majority of the accessible floor area and lower wall surfaces—the regions most likely to contain inspection targets such as valve bodies and pipe connections—are classified as observed (green) under the test trajectory. Second, persistent red clusters appear in two characteristic locations:

1. Directly beneath dense structural clusters where downward-looking angles are geometrically blocked.
2. Along sharp structural boundaries where shadows in the Octomap collision volume prevent ray penetration. These occlusion-driven gaps are not apparent from kinematic reachability analysis alone, which underscores the value of the combined workflow.

The heatmap overlaid directly in the planning scene is particularly useful operationally: a path planner, whether human-operated or automated, can immediately see which trajectory adjustments are needed to sweep the camera past the red regions without having to interpret abstract coverage metrics.

7 Future Plans

The results from this reporting period establish a strong technical foundation. Two major development thrusts are planned for the next program year, each building directly on the hardware and simulation infrastructure developed to date.

7.1 Thrust 1: Closed-Loop Low-Level Control Using Load Cell and Encoder Fusion

The current robot operates under open-loop joint position commands derived from actuator displacement models. While functional, this approach accumulates error over time due to wire stretch, hysteresis in the routing path, and manufacturing tolerances in the joint assembly. The availability of real-time data from both the joint encoders (joint angle) and the inline load cells (wire tension) creates an opportunity to implement closed-loop joint-level control.

The planned approach involves development of a model-based torque estimator that uses wire-tension measurements and joint geometry to estimate the torque being delivered at each joint, then using this estimate in a cascade control loop: an outer-position controller generates torque setpoints based on encoder error, and an inner-tension controller adjusts actuator velocity to track those setpoints. This architecture should reduce tracking errors, improve repeatability, and enable the robot to detect and respond to unexpected contact forces—important for safe operation in a cluttered pit environment.

Additionally, the team plans to characterize the wire-transmission nonlinearities (elasticity, pulley friction, length-tension coupling) through systematic identification experiments, and to incorporate the resulting model into a real-time feedforward compensation term. The combination of model-based feedforward and sensor-driven feedback is expected to yield substantial improvements in end-effector positioning accuracy.

7.2 Thrust 2: Embodied AI for Autonomous Inspection Coverage

The longer-term thrust (Table 2) is the development of an embodied AI agent capable of autonomously planning and executing inspection trajectories that achieve complete coverage of a pit interior, including the occlusion-driven gaps identified in the current heatmap analysis, without pre-programmed paths.

The simulation framework developed this year provides exactly the training environment needed for this effort. The agent will be formulated in the Isaac Sim/ROS 2 environment, where it can receive state observations (current joint configuration, accumulated coverage heatmap, camera depth data) and take actions (joint velocity or position commands) with the objective of maximizing the fraction of pit surface observed while minimizing redundant revisitation and avoiding collisions.

The reward signal will be shaped around the inspection heatmap metric seen—the fraction of pit surface classified as observed—incentivizing the agent to actively seek out and navigate toward unobserved (red) regions. Curriculum learning strategies will be used to progressively increase scene complexity during training, starting with simple open-volume pits and gradually introducing the cluttered infrastructure geometry representative of actual Hanford conditions.

Once a policy is trained to a satisfactory performance level in simulation, transfer-learning techniques will be employed to adapt the policy to the physical robot hardware. The key challenge in sim-to-real transfer for this system is the mismatch between the simulated joint dynamics (idealized, continuous) and the physical wire-driven system (nonlinear, with hysteresis and elasticity). The closed-loop controller improvements from Thrust 1 are expected to reduce this gap significantly by making the physical system behave more like the simulation model. Domain randomization—varying simulated wire parameters, actuator delays, and sensor noise levels during training—will further improve the robustness of the learned policy to hardware imperfections.

The ultimate goal is a system where the robot, once lowered through a viewport into an unknown pit, can autonomously explore and inspect the interior, building a live coverage map, identifying unobserved regions, and adapting its trajectory in real time until full coverage is achieved without direct human control of individual joints.

7.3 Summary of Planned Activities

Table 2. Planned Activities for the Next Program Year

Activity	Description	Timeline
Wire transmission characterization	Identify elasticity and friction model parameters from experimental data Cascade joint controller	Implement encoder + load cell fusion in closed-loop position/torque control
AI environment setup	Formalize observation space, action space, and reward for coverage	Embodied AI
AI agent training (sim)	Train autonomous inspection policy in Isaac Sim/Lab with curriculum	—
Sim-to-real transfer	Domain randomization + policy fine-tuning on physical hardware	—
Full system demonstration	End-to-end autonomous inspection test in representative pit mock-up	—

7.4 Design Review at FIU

In September 2025 INL, FIU, and representatives from H2C met in Florida at the FIU campus to review the current status of the project, evaluate the prototype designs (as seen in Figure 28), and discuss future work on the project. The main objective of these meetings was to review the progress of the prototype arms. The two functional designs that were showed off were the scissor arm design and the cable-driven arm. After the review, the teams compared the functionality of each prototype to the functional requirements to determine which would perform best in pit environments. During this time, input from the Hanford operations team was very valuable. The decision was made to stop development of the scissor arm and develop an updated final prototype of the cable-driven arm.

In addition to the physical prototype designs the teams reviewed the plans for an operator interface, discussing how site operators would prefer to interact with a system and what features would allow for easy adoption in the tank-farm sites. These discussions also included the Unity-based operator interface and the broader simulation and autonomy training workflow being developed across the project using Nvidia Isaac tools. While full autonomous inspection is the end goal of the project, a need for full manual control was emphasized.

Figure 28. Project Teams at FIU with prototypes.

7.5 Project Updates with Funding Pause

Roughly halfway through the project year, it became apparent that the government shutdown and other factors might interrupt funding for the project. With the second half of the year's funding at risk of delay, it was decided that the teams should pull back on certain activities to ensure that funds were in place for contract obligations, intern-support obligations, needed purchases, and top-priority tasks. Near the end of the Calander year, it was confirmed that the second half of Year 2 funding would be delayed, with no definitive date in place for release of additional funding. At the direction of DOE, the project looked at the task list and milestones for Year 2, decided which tasks should be placed on pause and what could be accomplished with the remaining funds. A plan was put into place for pausing these tasks, and documenting progress so that they could be picked up again when funding was received.

With the remaining funding, it was decided that the project could support three major tasks:

1. Publication support and conference presence
2. Demonstrate robot, digital twin, AI-based autonomy training, and visualization prototypes
3. Creating knowledge bases, repositories, and instructional documents for startup of work.

The project submitted three conference papers, two conference panels and is planning one journal publication, and participants felt it important to complete the presentations and the papers. The papers and presentations are as follows:

1. 26153 Multi-Robot collaboration for hazardous environments for Waste Management Symposia [10]
2. 26388 Simulation driven control for a riser deployed articulated robot in Hanford waste tank pits for Waste Management Symposia [11]
3. Calibration Fault detection and performance optimization for the Journal of Field Robotics
4. Learning time optimal inspection policies for IEEE conference
5. INL and FIU combined panel on digital engineering and robotic simulation for Waste Management Symposia
6. INL and FIU combined panel on national lab and university collaboration for Waste Management Symposia.

Various tasks from the Year 2 scope were paused due to funding and are expected to resume at a later date. They include construction of a final version of the robotic arm at FIU and shipping it to INL for integration into the full system, and building a pit mock up at INL. Additionally, finalization of the mechanical design of the above-pit robot support structure and introduction of extraction mechanisms, robot control system development, and an Isaac Sim and digital twin simulation to physical demonstration are planned. These tasks are in various stages of development, with progress documented. The project has hired interns through the summer to continue development of Isaac Lab machine learning and AI training for simulated autonomous inspection of the pits and development of the visualizations. Intern support continues simulation development, and the near-term objective is to convert the current validated environment into a training-ready autonomy framework. Planned work includes defining measurable inspection-coverage metrics, expanding reward design around coverage quality, collision avoidance, and efficiency, increasing domain randomization across pit layouts and clutter conditions, and incorporating ZED and LiDAR outputs more directly into the decision loop. Additional effort will also examine how the PTZ system can function, not only as an external observer, but as an active lighting and sensing asset that improves overall scene observability during autonomous inspection. At this time, it is expected that the full prototypes planned for the Year 2 milestone will be completed.

References

1. Cao, P. et al. 2025.“Power Quality and Load Capacity Evaluations of an Electric Vehicle for Multi-Robot System Applications.” *2025 IEEE Conference on Technologies for Sustainability (SusTech)*, Los Angeles, California: 1–6, https://doi.org/10.1109/SusTech63138.2025.11025685.
2. Nexxis. (n.d.). “Argus 3D SLAM system.” Retrieved November 5, 2025, from https://nexxis.com/product/argus/.
3. NVIDIA Corporation. (n.d.). “Isaac ROS (Robot Operating System).” NVIDIA Developer. Accessed November 5, 2025, from https://developer.nvidia.com/isaac/ros.
4. Bernus, B. et al. 2020. “Robotic Specialization in Autonomous Robotic Structural Assembly.” *2020 IEEE Aerospace Conference*, Big Sky, Montana: 1–10. https://doi.org/10.1109/AERO47225.2020.9172620.
5. Cao, P. et al. 2025. “Development of a Multi-Robot System for Autonomous Inspection of Nuclear Waste Tank Pits.” *Applied Sciences* 15(17): 9307. https://doi.org/10.3390/app15179307.
6. Xu, W. and F. Zhang. 2021. “FAST-LIO: A Fast, Robust LiDAR-Inertial Odometry Package by Tightly-Coupled Iterated Kalman Filter.” *IEEE Robotics and Automation Letters* 6(2): 3317–3324. https://doi.org/10.1109/LRA.2021.3064227.
7. Xu, W. et al. 2022. “FAST-LIO2: Fast Direct LiDAR-Inertial Odometry.” *IEEE Transactions on Robotics* 38(4): 2053–2073. https://doi.org/10.1109/TRO.2022.3141876.
8. Girardeau-Montaut, D. 2016. “CloudCompare.” *EDF R&D Telecom ParisTech* 11(5). https://www.eurosdr.net/sites/default/files/images/inline/04-cloudcompare_pcp_2019_public.pdf.
9. Rusu, R. B. and S. Cousins. 2011. “3D is here: Point Cloud Library (PCL).” *IEEE International Conference on Robotics and Automation*, Shanghai, China: 1–4. https://doi.org/10.1109/ICRA.2011.5980567.
10. D’Andrea, A., P. Cao, and K. Houck. 2026. “Multi-Robot Collaboration for Hazardous Environments.” *Waste Management Symposia*, Paper No. 26153.
11. Rayhan, M. et al. 2026. “Simulation driven control for a riser deployed articulated robot in Hanford waste tank pits for Waste Management Symposia.” *Waste Management Symposia*, Paper No. 26388.

U.S. DEPARTMENT *of* ENERGY

For more information, visit: energy.gov

DOE/Publication Number INL/RPT-26-93692
June 2026