Contract Name Here

Proposal Title Here

Volume Title Here

Solicitation Number: XXX-XXXXXX

Submitted on: DD Month YYYY

Submitted to: Office
Point of Contact
Contact Information

1. Provide Company Name, Title, Date, Point of Contact, Email Address/Phone/Address

Company
PhantomShift Principal Investigator White Paper Submission
Date: December 4th 2025
Point of Contact: Principal Investigator

1.1. EXECUTIVE SUMMARY: Provide a detailed description/summary of your solution.Proposals should include the description of the solution, techniques, applicability to the intended use. Submissions should define what capability desirements the solution is addressing, any desirements we should be aware of and any current deployment of capability in support of government or commercial systems.

PhantomShift is an advanced research prototype developed by Company to transform the landscape of offensive security assessment through AI-augmented penetration testing. The solution integrates large language models (LLMs) directly with a live Unix-like terminal environment—accessible through a secure Streamlit-based browser dashboard—to serve as an intelligent co-pilot for hands-on red teams and penetration testers. The platform operates by continuously observing the user's actions and terminal commands, leveraging contextual awareness to provide testers with actionable recommendations that include copy-pasteable commands, explanations of intent, potential impacts, risk assessments, and key monitoring indicators. This approach significantly reduces cognitive load, mitigates context switching between documentation and the terminal, and enhances decision-making by keeping testers focused and informed in real time.

Company's solution specifically addresses capability desirements for reducing manual effort, increasing workflow efficiency, and maintaining high standards of risk awareness in offensive security operations. The system's architecture employs a Python session engine, stateful history management, an LLM orchestration layer for generating suggestions, and an optional retrieval-augmented generation (RAG) knowledge base grounded in established penetration testing best practices. PhantomShift's design prioritizes responsible use: it is optimized for isolated lab environments or authorized engagements, ensuring human testers retain full control through editable and rejectable AI recommendations, comprehensive risk flagging, and session persistence for context continuity. Current deployments are focused on collaborative partnerships with security researchers, red teams, and training organizations seeking to experiment with advanced, AI-enabled testing workflows. Company brings extensive experience supporting national security and commercial clients, validated by 18 years of mission-critical engineering for defense and intelligence agencies, and leverages this operational insight to ensure PhantomShift aligns with the stringent standards and operational realities of government and enterprise environments.

1.2. TECHNOLOGY CONCEPT: Describe the unique aspects of your solution, any detailed technical specifications, and the proposed work as it relates to our known/ unknown capability desirements.

PhantomShift represents a novel approach in AI-augmented penetration testing, distinguished by its tight integration of large language models with a live Unix-like shell accessible via a secure browser interface. The core innovation lies in the seamless feedback loop between the user's terminal activity and the contextual suggestions generated by the system. The Streamlit-based dashboard embeds an interactive terminal, underpinned by a Python session engine that meticulously manages command history and state, ensuring suggestions remain relevant to the ongoing assessment. An LLM interface layer interprets session context to synthesize actionable, copy-pasteable commands with detailed intent, impact analysis, indicators to monitor, and risk assessments. To further enhance recommendation fidelity, PhantomShift incorporates a retrieval-augmented generation engine connecting to a curated knowledge base of vetted penetration testing methodologies, attack playbooks, and cheat sheets. Optional connectors enable targeted web search when internal resources are insufficient, heightening the flexibility and coverage of PhantomShift's AI co-pilot.

The proposed work seeks to address known capability desirements such as reducing cognitive workload, accelerating repetitive or research-intensive tasks, and enhancing the analytical depth available to human testers—all without diminishing human oversight or introducing uncontrolled risk into test scenarios. PhantomShift's architecture enables controlled, session-persistent testing that supports pausing, resumption, and full review of prior actions, a critical requirement for regulated or complex engagements. Human-centered controls permit editing or discarding AI-generated suggestions before execution, ensuring compliance with engagement rules and ethical guidelines. For unknown or emerging capability requirements, PhantomShift's modular API layer, built atop a lightweight FastAPI service, creates opportunities for rapid customization, team collaboration integration, and future expansion to tool-specific fine-tuning and policy-aware guardrails. The Company team has implemented automated output normalization as part of ongoing stability improvements and is prepared to support additional operationalizations driven by partner feedback, drawing from deep experience in offensive security, cyber engineering, and mission-driven IT innovation.

1.3. APPROACH: Describe how your solution solves the designated Focus Area. If applicable,identify any past or current efforts in which your solution has successfully demonstrated results that have, partially or in its entirety, addressed known/unknown capability desirements.

PhantomShift solves the designated Focus Area by delivering an AI-assisted penetration testing environment that augments the technical workflows of red teams and security researchers. The solution achieves this by embedding an interactive Unix-like shell into a secure, browser-based interface and orchestrating large language models to analyze session history, interpret context, and proactively surface actionable recommendations that align with established best practices. This approach dramatically reduces the time spent referencing external documentation, automates routine yet error-prone commands, and enables testers to maintain concentrated focus during complex operations. Human operators retain full control of every step, leveraging the system's contextual suggestions for increased speed and precision while rigorously reviewing proposed actions against engagement rules and risk assessments. By integrating session persistence, comprehensive logging, and editable AI responses, PhantomShift ensures that critical decision points and operational continuity are maintained across both short and extended engagements.

Company's previous efforts in offensive security and mission-critical engineering underpin PhantomShift's ongoing development and operational viability. The company's support for government and defense entities has included deployment of advanced cyber testing frameworks, incident response automation, and custom workflow tools for high-security environments. Lessons learned from these engagements directly influenced PhantomShift's emphasis on session integrity, secure handling of sensitive workflows, and adaptability to evolving user requirements. Within current research partnerships and training labs, Company has demonstrated PhantomShift's capability to increase penetration tester productivity, foster safer experimentation, and surface higher-quality risk analysis without introducing undue automation risk. The system's modular architecture and responsive knowledge management pipeline continue to adapt in line with feedback from real-world collaborators, validating its suitability for both known and emergent capability needs in dynamic security operations.

1.4. ROM COST/SCHEDULE: Lay out your Rough Order of Magnitude (ROM) unit cost and minimum delivery lead time for your company's solution(s) to support a 1-6 month user assessment. You may want to list the variables that could influence the cost and schedule parameters, and you may want to include any assumptions made in calculating the ROM cost and schedule.

Company estimates the Rough Order of Magnitude (ROM) unit cost for deploying PhantomShift to support a 1-6 month user assessment falls within the range of $90,000 to $225,000, scalable according to assessment duration, number of concurrent users, deployment environment requirements, and level of customization or technical support requested. This baseline includes initial configuration of the PhantomShift platform in a controlled, isolated test environment; user onboarding and operational training for red team personnel; continuous access to the core AI co-pilot features; and remote support and maintenance throughout the assessment period. The delivery lead time for initial deployment and user readiness is four weeks from receipt of order and finalized requirements, contingent on prompt access to integration endpoints and designated test infrastructure. Shorter lead times are negotiable when deploying unmodified prototype builds to pre-existing, lab-standard environments.

Key variables impacting cost and schedule parameters include the need for site-specific security hardening, integration with customer-managed infrastructure, incorporation of advanced team collaboration features, or customization servicing specific toolchains or reporting requirements. Assumptions underlying this ROM include the use of Company's existing hosting and support infrastructure, no handling of production or classified data, and user assessments being conducted strictly in accordance with recommended safe usage guidelines. Additional charges may apply for extensive environment hardening, regulatory policy adaptation, participation in extended after-action reviews, or substantial modifications to the research prototype aligned with unique customer workflow needs.

1.5.

Company ensures all submission content, including the technical, management, and cost sections, is presented in a concise manner that does not exceed three pages in Times New Roman 12-point font, exclusive of the title page. The document layout is meticulously structured with clear section demarcations, aligned margins, and single-spacing to maximize information density while maintaining readability. All graphics, tables, or figures are purposefully minimized or formatted to fit within these constraints, ensuring that critical details regarding PhantomShift’s features, architecture, prior performance, cost, and delivery timelines are fully addressed within the established page limit. Company’s disciplined approach to proposal compliance guarantees that evaluators receive a comprehensive yet succinct overview, as required by submission guidelines, and supports smooth review and evaluation procedures.