Brookhaven National Laboratory

U.S. DEPARTMENT of ENERGY

FY27 NPP LDRD Type B Pre-Proposal

Empowering ePIC with AI-Enabled Collaborative Tools and Infrastructure

PIs: Thomas Ullrich (presenter), Dmitry Arkhipkin

January 9, 2025

@BrookhavenLab

FY27 NPP LDRD Type B Pre-Proposal

Proposal title: Empowering ePIC with AI-Enabled Collaborative Tools and Infrastructure

Primary Investigator: Thomas Ullrich, Dmitry Arkhipkin

Other Investigators: Uma Ganapathy, Brian Page, Peter Steinberg

Indicate if this is a cross-directorate proposal: Yes
If yes, identify other directorates/organizations: SCDF

Proposal Term: October 1, 2026 - September 30, 2028

Program: NP

Abstract

This project will develop and deploy locally hosted artificial intelligence (AI) tools to enhance collaboration, communication, and efficiency within the ePIC experiment at the future EIC.

ePIC brings together hundreds of scientists and engineers worldwide whose coordination relies on conventional digital platforms. The effort will integrate medium-sized AI models into this environment to automate documentation and meeting summaries, assist software, data-analysis, and publication workflows, and enable rapid retrieval of technical and design information.

All AI processing will run on dedicated local hardware, ensuring privacy, data security, and independence from commercial cloud systems. Leveraging recent advances in compact, high-performance AI hardware, the project introduces a cost-effective and trustworthy approach that aligns with DOE's mission to accelerate scientific discovery through advanced computing and data-driven innovation.

Return of Investment & Broader Impact

- Modest, time-bounded investment to develop a reusable capability for BNL and the user community
- Near term: validate methods/software/instrumentation/workflow to reduce technical risk and speed iteration
- Medium term: transition to routine use with potential for higher data quality, faster turnaround, and more reliable operations
- Efficiency upside: possible savings in staff time and/or improved facility utilization (less rework, fewer failed runs, streamlined processes)
- Strategic impact: improved competitiveness for follow-on funding through demonstrated feasibility and broader adoption
- Broader impact: results, tools, and lessons learned can be shared and reused across programs, helping strengthen collaboration in other fields

Funding

- Budget Justification
  - Salary: 100% postdoc
  - Salary: RS7 T. Ullrich 5%
  - Salary: RS7 P. Steinberg 5%
  - Salary: RS5 B. Page 5%
  - Salary: SCI1 D. Arkhipkin 10%
  - Salary: SCI1 U. Ganapaty 10%
  - Materials and Services: modest computing and software support for AI model development and integration with collaboration platforms - $\mathcal{O}$($20k) in first year
  - Travel: support for the postdoc and team to participate in collaboration meetings (EIC, NPPS, AI for Science), and to coordinate with BNL's AI/ASCR initiatives

Description of LDRD - Motivation

Empowering ePIC with AI-enabled collaborative tools and infrastructure

- ePIC is a large, distributed collaboration (≈1,000+ members; 180+ institutions; 26 countries) supporting EIC detector construction and physics.
- Coordination overhead (meetings, documentation, software workflows, design reviews) consumes significant expert time.
- Physics has a history of collaboration-driven innovation deploying the latest technologies (e.g., the World Wide Web) which today means intelligent assistance.
- Project goal: deploy privacy-preserving, locally hosted AI assistants integrated with ePIC collaboration tools to reduce friction and risk.

Key constraint: all AI inference runs on dedicated on-prem hardware (no commercial cloud) to protect sensitive internal data.

Description of LDRD - Concept and Approach

- Build on the CRISP* collaboration platform to add an “AI module” that connects to documents, events, tasks, and workflows.
- Run medium-sized language models locally on AI-enabled mini-PCs for secure, cost-effective deployment.
- Implement secure ingestion + indexing + retrieval (RAG) over internal ePIC documentation, code, and notes.
- Provide user-facing functions: meeting/document summarization, action-item extraction, Q&A over the knowledge base, and workflow assistance.
- Two-year plan: Year 1 pilot infrastructure + prototype assistant; Year 2 expand features, fine-tune for ePIC terminology, evaluate and train users.

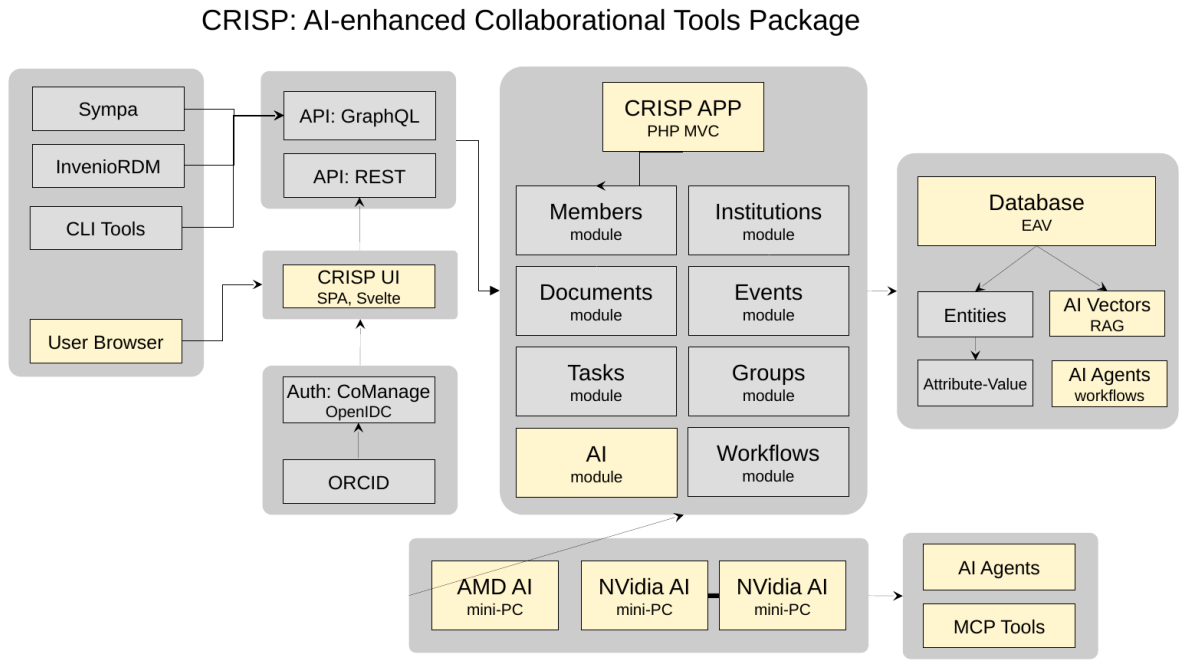

Design emphasis: privacy, model governance, and minimal disruption to existing collaboration workflows.

* Collaborative Research Information Sharing Platform

Deliverables, Team, and Resources

Deliverables	Personnel	Procurements & Milestones
• Operational AI-enhanced prototype integrated with ePIC tools • Deployable software stack for on-prem AI assistants • Documentation + training materials for users and maintainers • Evaluation report with metrics and recommendations	• New postdoc (100%): integration, model adaptation, pilot rollout • D. Arkhipkin (10%): infrastructure + AI deployment • U. Ganapathy (10%): systems integration • T. Ullrich / P. Steinberg / B. Page (5% each): ePIC coordination + user engagement	• Equipment: ~$20k in Year 1 (AI-enabled mini-PCs for pilot users) • M6: hardware + base model operational • M12: pilot assistant tested with user feedback • M18: expanded features integrated • M24: evaluation (targeted time savings and user satisfaction)

Success metrics (targets): ~30% reduction in time spent on documentation/coordination (survey-based) and ≥80% positive feedback from pilot users.

Summary

Secure, BNL-hosted AI assistant to streamline ePIC collaboration

Problem: collaboration knowledge is spread across documents, code, working groups, subsystems, and meetings — essentials at times hard to find and reuse

Concept: on-prem AI assistant (CRISP-based) for fast Q&A, context retrieval, and meeting/notes support

Approach: deploy an AI model, implement secure ingestion+indexing and AI workflow/reporting agents

Deliverables: deployable AI mini-PC/software stack; ePIC-tuned assistant; documentation/training; impact evaluation

Outcome: reduced technical risk and faster iteration for ePIC, with a reusable pattern for other large collaborations

Key constraint: secure, BNL hosted AI-implementation (no commercial cloud)

Backup Slides

CRISP: AI-enhanced Collaborational Tools Package

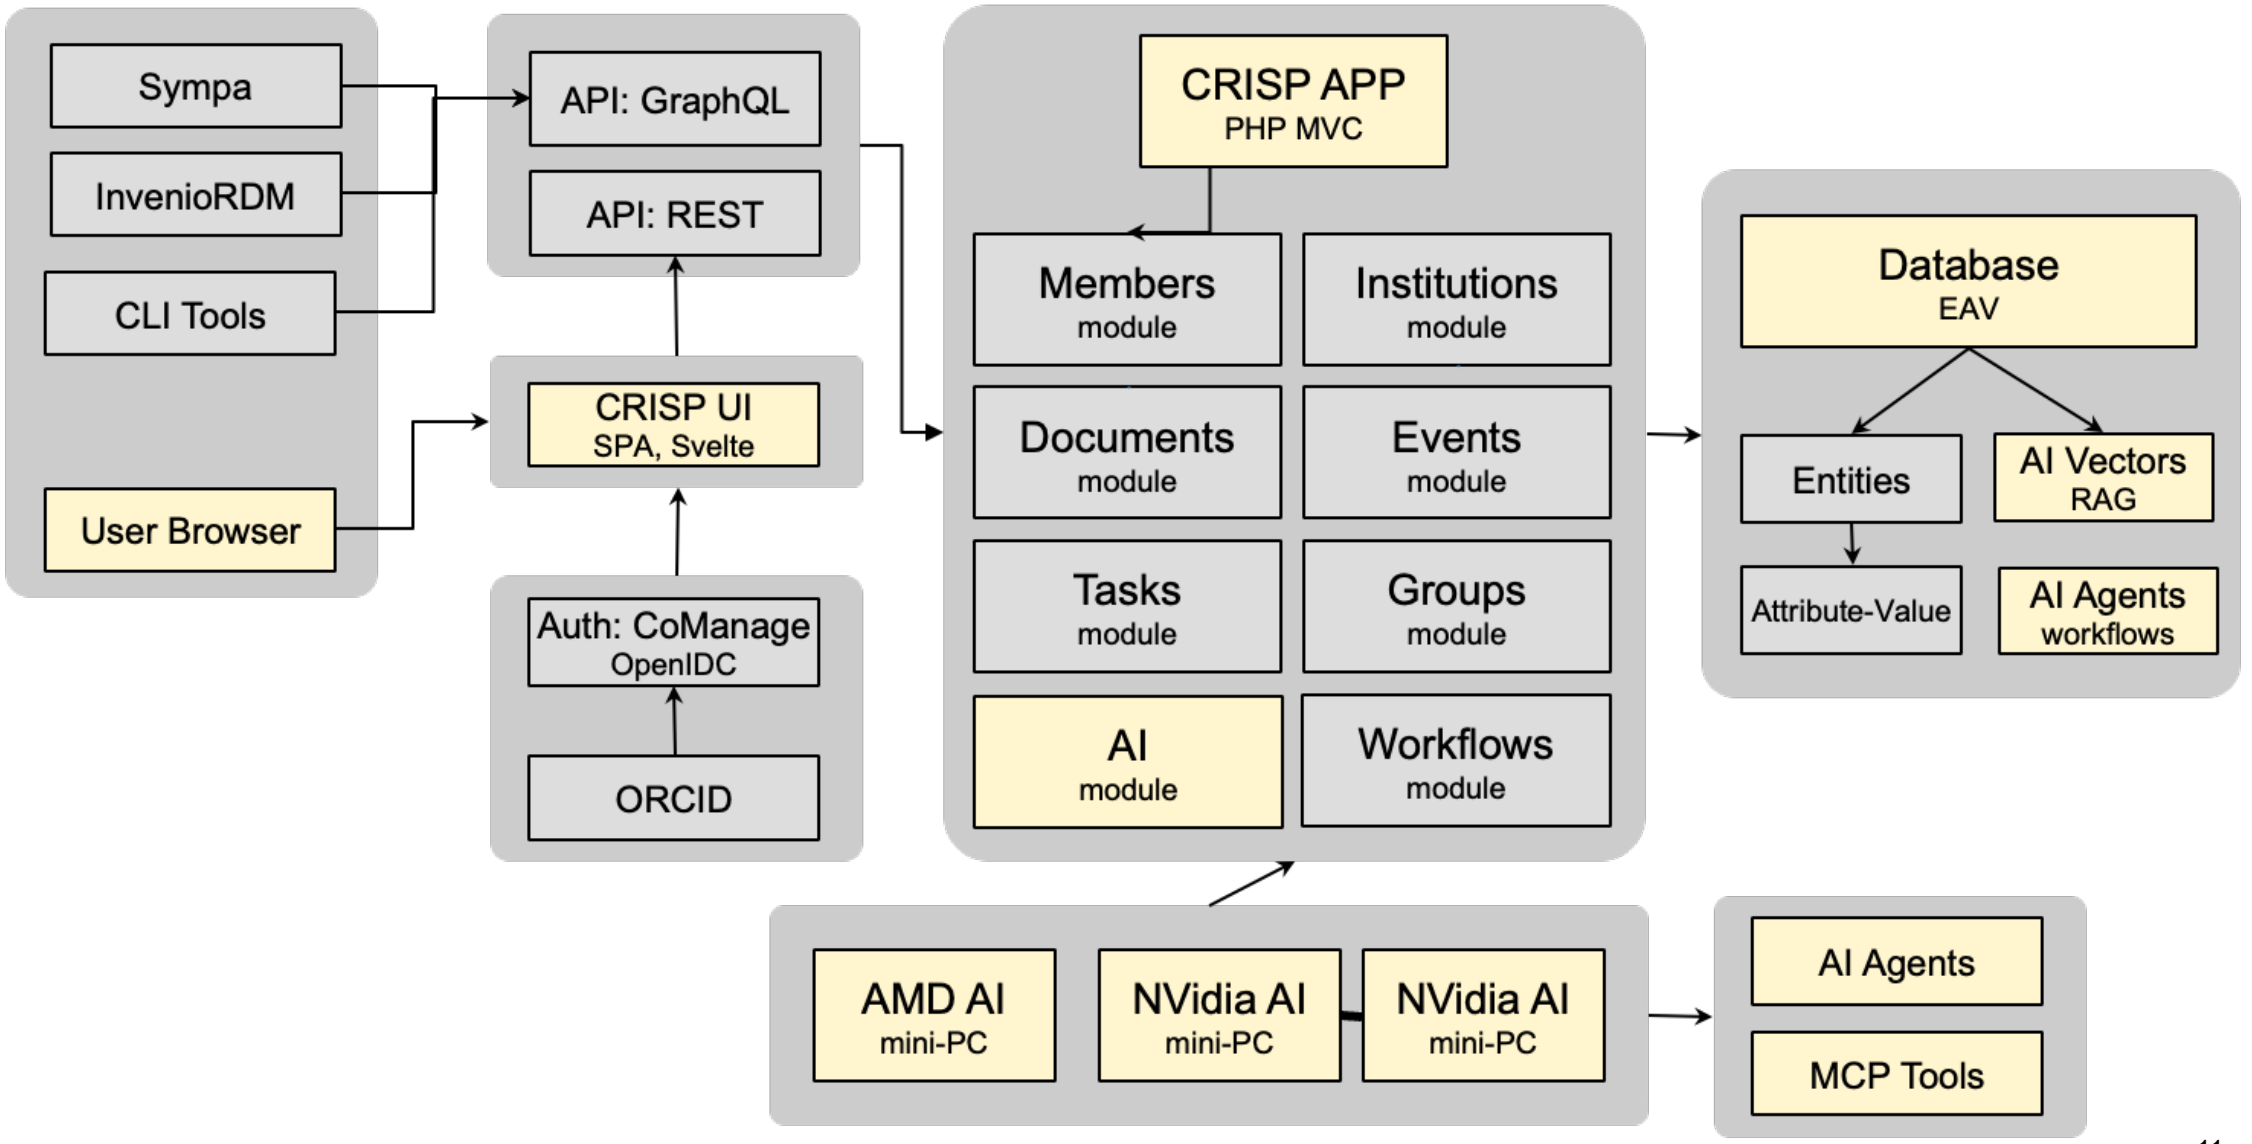

CRISP: AI topologies to be implemented and tested