PROTOPIA

Unlock Your Sensitive Data for AI

Defense organizations leveraging AI for critical missions face key security challenges:

1. Protecting data during AI employment and training.
2. Securing access to external data.
3. Expanding LLM use while safeguarding information.
4. Utilizing allied forces' sensor data without exposure.

Protopia AI's Stained Glass Transform (SGT) converts sensitive data into irreversible, AI-compatible stochastic representations, eliminating raw data exposure risks.

KEY VALUE PROPS

Securely Expand Data Utilization

Make available details of the entire breadth of your data across the AI lifecycle (model training, fine-tuning, inferencing) while neutralizing exposure risks.

Improve Infrastructure Utilization

Securely use data across on-prem, hybrid, and cloud environments for efficiency and faster value realization.

Retain Data Ownership

Create stochastic representations that preserve utility for target models but are unintelligible to humans or other AI models.

Deploy Anywhere

SGTs add minimal latency (typically ms) and can even run on CPUs or embedded devices, including sensors.

How Stained Glass Transforms Work

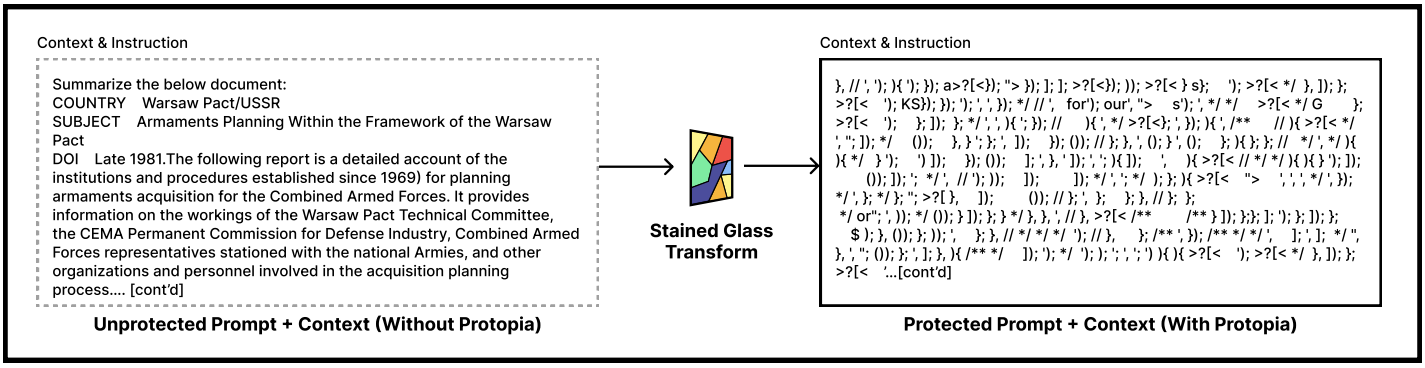

Response quality is maintained with obfuscated data

Model	Average Tokens Transformed	Sentence Completion (HellaSwag)	Language Understanding (MMLU)	Model Truthfulness	Abstraction Reasoning
Mistral 7B V 0.2 Base Model	0% (ie., plain-text)	76.53%	57.21%	68.27%	50.94%
Mistral 7B V 0.2 w/Stained Glass	98.44%	76.67%	55.39%	67.90%	51.02%

Note: Mistral 7B model is used as example for validation

Trivial Latency for Stained Glass Transform Layer

Stained Glass for Mistral 7B Latency:
~15 millieseconds

Generation Latency of Mistral 7B:
~few seconds

PROTOPIA

Seamless Integration with Existing AI Pipelines

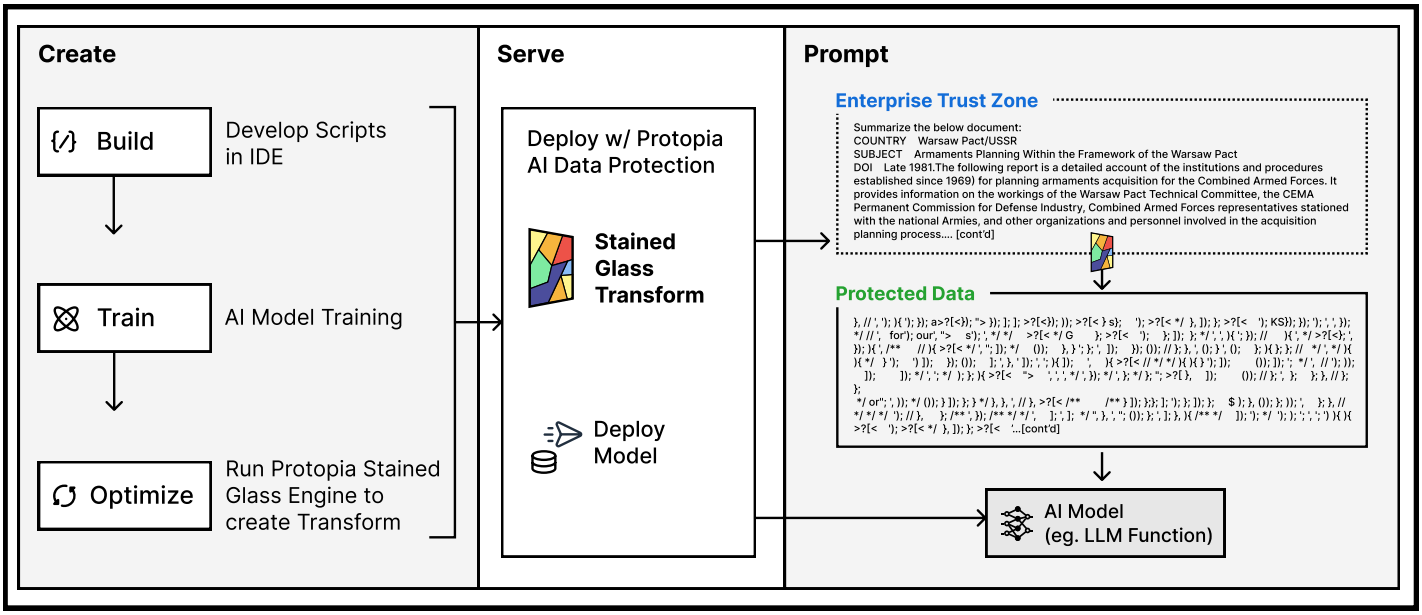

No deviation from existing AI/ML deployment pipelines

- No change to the weights of trained AI model being targeted
- Stained Glass Engine introduces minimal latency compared to target model training
- Stained Glass Transform introduces next to no latency in inference pipeline

No dependency on non-proven or scarce technology (HW or SW)
Built on widely used standard SW packages:

- Python 3.9, 3.10
- Pytorch >= 2.1.0
- *Hugging Face Transformers
- *Hugging Face Tokenizers

*Hugging Face Libraries are not required for non LLM use cases

Spotlight: Securing Personnel Information for the US Navy

SGT transformed images of VIP personnel in live video for the U.S. Navy as a part of the Trident Warrior Exercise 2022, preserving AI accuracy while protecting identities.

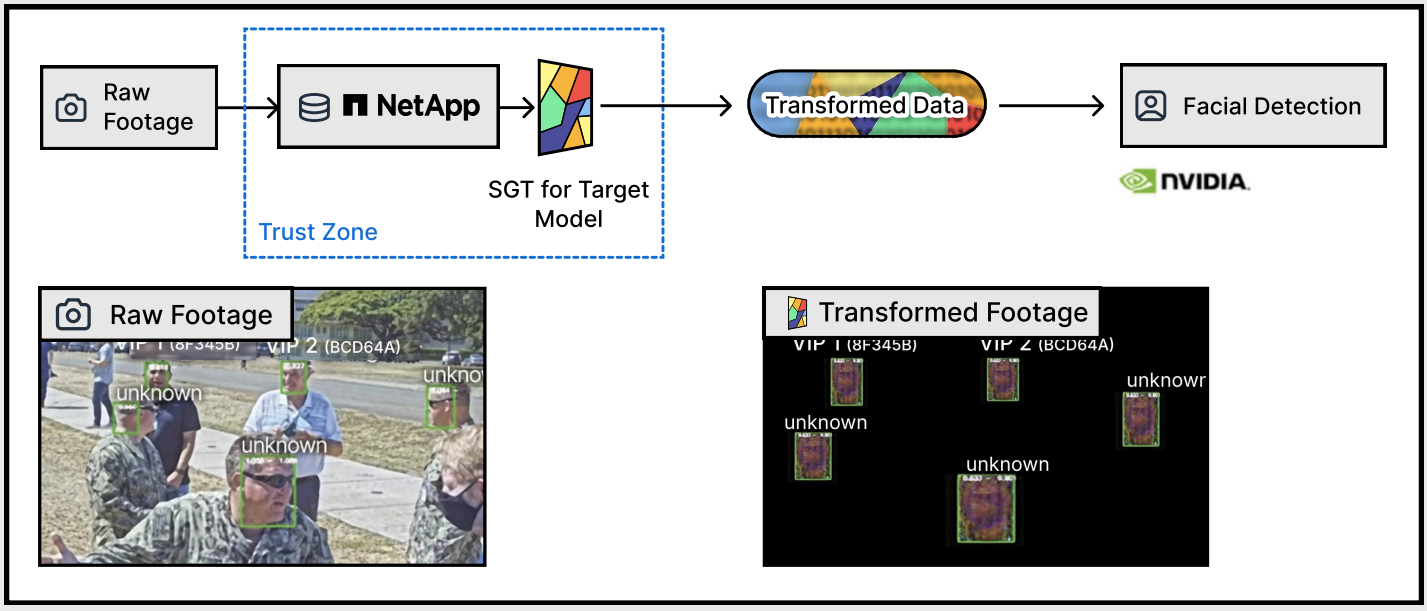

Partner to leading compute and data infra providers

NVIDIA NetApp ORACLE PartnerNetwork

Gartner COOL VENDOR 2022

Hype Cycle 2024 x 9
AI TRiSM 2023 x 6
Emerging Tech Radar 2024
Emerging Tech Radar Edge AI 2024

aws

One of the 21 startups selected from 1,200 for AWS Gen AI Accelerator