PROTOPIA

www.protopia.ai

Data Privacy That Scales With Your AI Infrastructure

Deploy AI in Production While Protecting Sensitive Data

Organizations deploying AI with sensitive data face critical security challenges:

Inference Endpoints Create New Security Vulnerabilities

Traditional security protects data at rest and in transit, but not during AI inference. As AI apps, data remains exposed in plaintext at endpoints.

Multi-tenant AI Applications Risk Data Leakage

Inference-stage vulnerabilities expose sensitive data, raising privacy risks for regulated industries, which slows adoption of AI use cases.

Traditional Data Protection Is No Longer Sufficient

Alternatives like redaction and data masking reduces data utility and AI use-case accuracy, impacting performance and cost.

Loss of Data Ownership with Third-Party Inference

When using third-party inference, enterprises lose control over where data is processed, undermining confidentiality and ownership.

Protopia Stained Glass: The Lightweight Solution for AI Inference Security

Protopia's Stained Glass Transform (SGT) converts sensitive data into irreversible, AI-compatible stochastic representations, eliminating raw data exposure risks while maintaining full model accuracy and near-zero performance impact.

Expand Data Utilization

Leverage your full data spectrum across the entire AI lifecycle while mitigating data leakage and exposure risks.

Improve Infra Utilization

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

Retain Data Ownership

Stochastic representations preserve accuracy for target models but are unintelligible to humans or other 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

Unprotected Prompt (Without Protopia)

Context & Instruction

Summarize the below document:
COUNTRY: Warsaw Pact/USSR
SUBJECT: Armaments Planning Within the Framework of the Warsaw Pact
DOI: Late 1981.The following report is a detailed account of the...

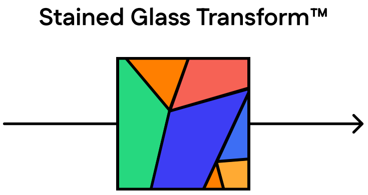

Protected Prompt (With Protopia)

Context & Instruction

```
}, // ', '); ){ '); }); a>?[<}); "> }); ]; ]; >?[<}); )); >?
[< } s};    '); >?[< */  }, ]); }; >?[<   '); KS}); }); '); ',
', }); */ // ',   for'); our', ">    s'); ', */ */    >?[< */
G     }; >?[<   ');    }; ]);  }; */ ', ', ){ '; }); //     ){ ', */
>?[<}; ', }); ){ ', /**     // ){ >?[< */ ', "; ]); */    ());
  }, } '; }; ',  ]);    }); ()); // }; }, ', (); } ', ();    }; ){ };...
```

Maintain High Performance While Eliminating Plain-Text Data Exposure

Model	Average Tokens Transformed	Sentence Completion (HellaSwag)	Language Understanding (MMLU)	Model Truthfulness	Abstraction Reasoning (ARC)
Llama 3.1 70B W/ SGT	98.44%	77.97%	77.88%	62.33%	51.02%
Llama 3.1 70B without SGT	0% (ie., plain-text)	77.61%	80.52%	66.9%	50.94%

Secure AI Inference Without Performance Trade-Offs

<1%
Added to inference time

~25 milliseconds
Latency for Llama 70B SGT for ~200 token

Data Protection Without Disruption
Model weights, accuracy, and AI workload performance rates remain intact

Roundtrip Data Protection: Preserving Data Ownership in Managed Inference

Protopia's Stained Glass Transform (SGT) integrates with high-performance inference endpoints to protect sensitive data throughout the AI inference lifecycle. Data stays secure from input to output, even in multi-tenant environments. Only the client sees the full prompt and response in plaintext—preserving ownership without tradeoffs.

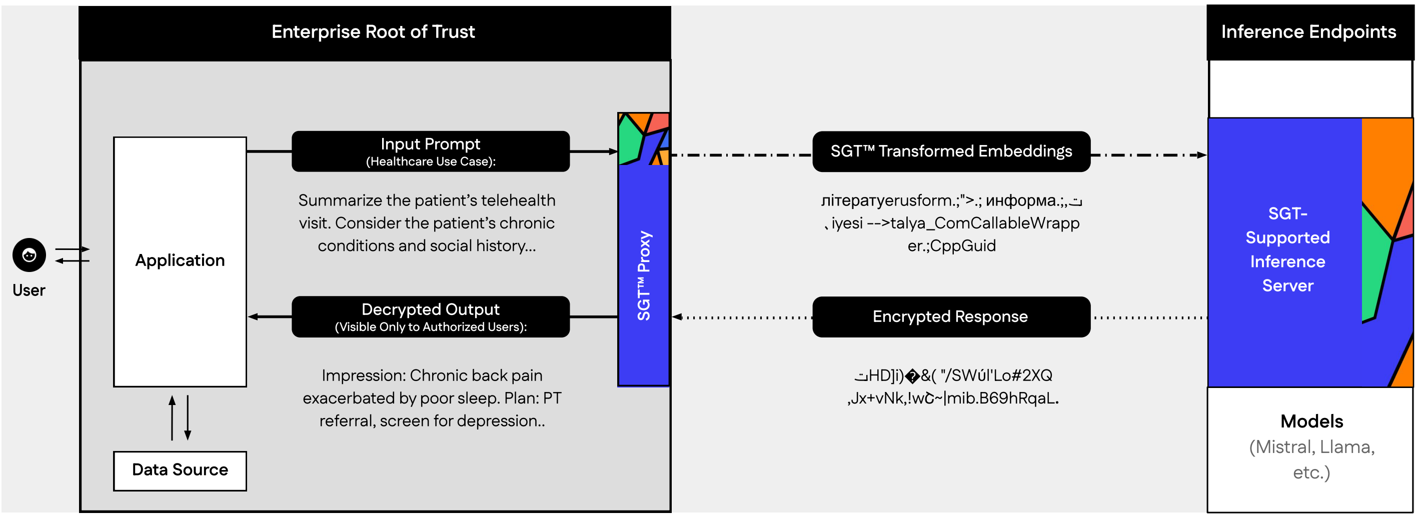

Summarize the patient's telehealth visit. Consider the patient's chronic conditions and social history...

Impression: Chronic back pain exacerbated by poor sleep. Plan: PT referral, screen for depression..

літературerusform.;">.; информа.;,ت
، iyesi -->talya_ComCallableWrapper.;CppGuid

تHD]i)�&("/SWúl'Lo#2XQ
,Jx+vNk,!wƻ~|mib.B69hRqaL.

Deploy Secure Enterprise Use Cases

Transition from POC to production by unlocking proprietary data assets for secure enterprise workloads.

Maximize AI Infrastructure ROI

Securely scale up enterprise users and data sources to maximize utilization of cost-effective inference endpoints.

Quickly Adopt Cutting-Edge Models

Access state-of-the-art models or your custom-deployed variants via serverless APIs, without exposing data in plaintext.

Secure AI Workloads Across the Full Inference Lifecycle

Protopia Stained Glass Transform (SGT) integrates seamlessly into inference endpoints, AI pipelines, and LLM applications without modifying underlying models. SGT unlocks secure AI adoption without compromising accuracy or scalability.

RAG Case Study: Advancing Medical Research While Protecting Data

Protopia partnered with Meta and a major U.S. health system (MHS) to securely implement a RAG application for oncology research. By using Protopia SGT, the institution overcame strict privacy constraints, enabling secure AI-powered collaboration with third-party researchers.

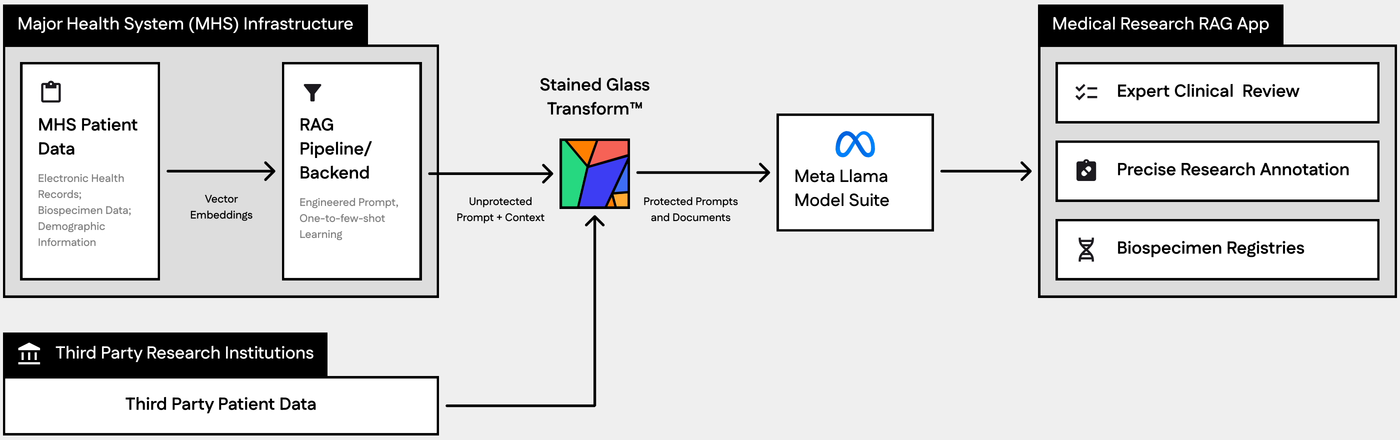
Major Health System (MHS) Infrastructure
MHS Patient Data
Electronic Health Records; Biospecimen Data; Demographic Information
Vector Embeddings
RAG Pipeline/ Backend
Engineered Prompt, One-to-few-shot Learning
Unprotected Prompt + Context
Stained Glass Transform™
Protected Prompts and Documents
Meta Llama Model Suite
Medical Research RAG App
Expert Clinical Review
Precise Research Annotation
Biospecimen Registries
Third Party Research Institutions
Third Party Patient Data

Gartner.

- Gartner AI TRiSM Market Guide
- Emerging Tech: Edge AI
- Emerging Tech: Data Privacy
- Gartner Hype Cycle x9

U.S. ARMY AMERICA'S NAVY

- U.S. Army XTech Scalable AI 1st Place '24
- U.S. Navy Trident Warrior '22

Partner to leading model, compute, and infra providers:

NVIDIA INCEPTION PROGRAM · aws · ORACLE PartnerNetwork · Meta · Lambda · NetApp · VAST

PROTOPIA

Learn More at Protopia.AI ↗