

Powering Private Inference in AWS Environments

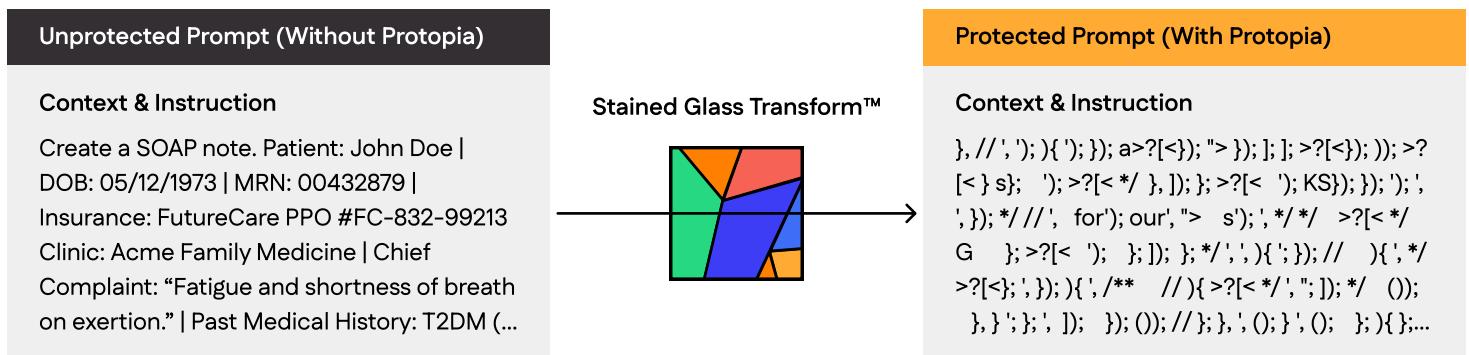
Unlock your sensitive enterprise data for AI

Available through Amazon Bedrock, SageMaker and AWS Marketplace, Protopia Stained Glass Transform (SGT) adds a drop-in privacy layer to protect critical data in enterprise AI applications. SGT unlocks cost-efficient, highly available AI compute, including on-demand GPUs and fully serverless inference by never exposing plain-text data to the target infrastructure.

 Unlock new AI use cases	 Lower total cost of ownership	 Accelerate time-to-impact
Enable AI applications to process even the most sensitive data, increasing the value of your enterprise AI while preserving privacy.	Run sensitive AI workloads on shared infrastructure. No need for isolated GPUs, enclaves, compute-intensive obfuscation techniques.	Use sensitive data safely in RAG and document workflows, without brittle token filtering or format-specific preprocessing.

How Stained Glass Transforms work

Stained Glass Transform (SGT) converts plaintext inputs into target-model-compatible stochastic embeddings through an OpenAI-compatible API, eliminating raw data exposure risks while maintaining full model accuracy and near-zero performance impact.



Maintain high performance while eliminating plain-text data exposure

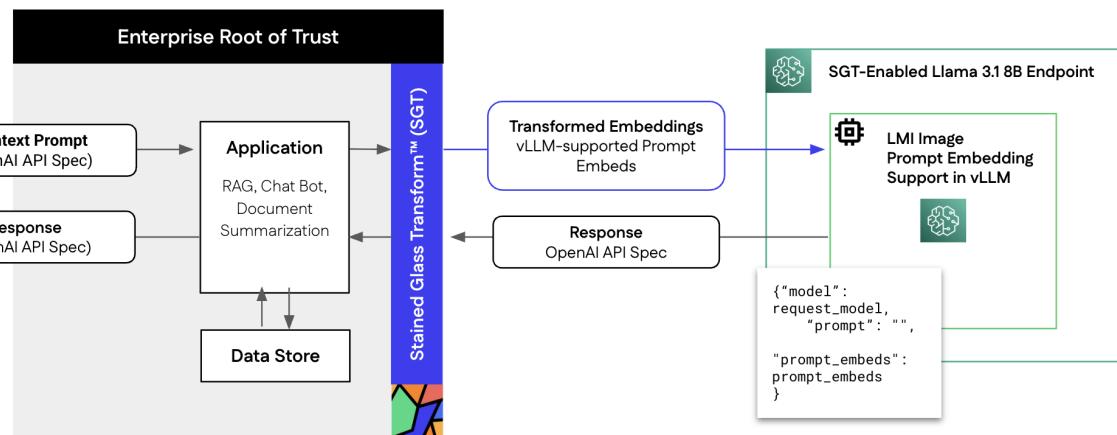
Model	Sentence Completion (HellaSwag)	Language Understanding (MMLU)	Model Truthfulness	Abstraction Reasoning (ARC)	Model weights, accuracy, performance remains intact
Llama 3.1 8B W/ SGT*	64.38%	50.13%	49.02%	67.63%	<1% Added to inference time
Llama 3.1 8B without SGT	67.2%	56.06%	52.99%	67.72%	~25 milliseconds latency for Llama 70B SGT for ~200 tokens

*98.44% average tokens transformed to stochastic representations.

Powering Private Inference in AWS Environments

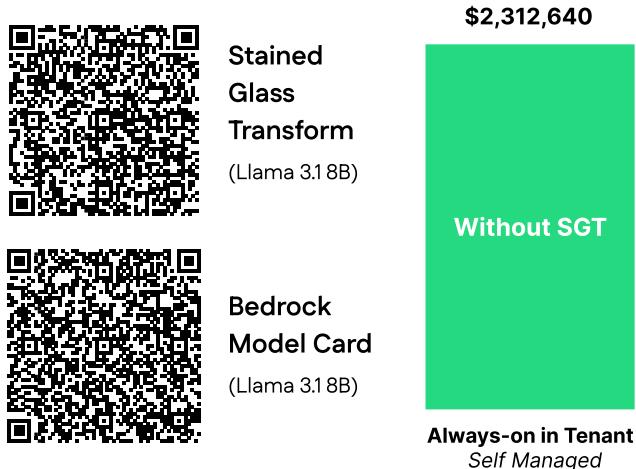
Streamlined & secure AI deployment on AWS

Stained Glass Transform ensures that sensitive information never exists in plaintext outside of the enterprise zone of trust. Customers maintain full data ownership without sacrificing accuracy, performance, or cost. SGT works seamlessly with AWS-native services like EKS, S3, and IAM as well as AWS's built-in procurement, billing, and governance tools.



Protopia Stained Glass unlocks managed AI infrastructure on AWS

Dedicated deployments are often the default choice for AI data privacy, minimizing the risk of data leakage or exposure to unauthorized users. But isolating workloads on dedicated compute carries a significant cost penalty. Protopia's Stained Glass Transform eliminates this tradeoff by keeping sensitive data inside your root of trust while still enabling fast, accurate inference on shared AWS infrastructure. By making it safe to run even your most sensitive workloads on managed inference services, Stained Glass reduces total cost of ownership and cuts operational overhead.



Example: Three 70B models serving 100 users in a global enterprise

