



Beyond experimentation: Production AI on a modern data foundation.

Kaycee Lai

Vice President, AI & Analytics



The Data Platform Enterprises Already Run On

Everpure™ delivers the industry's most advanced data storage platform to **store, manage, and protect** the world's data at any scale.

From AI to archive, Everpure provides a cloud experience with one, **unified storage-as-a-service platform across on premises**, cloud, and hosted environments—all built on the proven, non-disruptive Evergreen® architecture. By enabling organizations to **save time, money, and energy**, it's easy to see why Everpure is chosen by leading brands and has one of the highest Net Promoter Scores in the industry.

Everpure 2026 Fast Facts

Customer-First Innovation

14.5K

Global Customers

Organizations Across a Rapidly Growing Set of Industries Trust Everpure with their Data

40%

of U.S. Federal Agencies, Including Three Branches of the DOD

~64%

of Fortune 500 Companies

A FEW OF OUR CUSTOMERS

DENSO

MARS

Domino's

RED LOBSTER

fiserv.

SiriusXM

Recognition

6X

Great Places to Work

Recognized in Fortune's Best Workplaces in Technology List for Six Consecutive Years

12X

Gartner® Magic Quadrant™ Positioned as a Leader for Eleven Consecutive Years

84

Net Promoter Score One of the most consistently exceptional in enterprise tech

Environmental impact

BUSINESS VALUE

Up to

5X

Less Energy

Up to

5X

Less Rackspace

At least

90%

Less E-waste Generated

ENVIRONMENTAL

100%

of data center waste (by weight) from landfill guaranteed by new waste service vendor

77%

of the square footage of our global office real estate portfolio is certified LEED or equivalent

© 2026 Everpure

 2

14,500+ Customers and Counting



The World's AI Projects Run Here Too

QIXIO

AutoNation



GRAPHCORE

INTUIT



MARS

MEDIAZEN
GLOBAL NO.1 VOICE SERVICE CREATOR



OLYMPUS



ROSEN

servicenow



THG

...and over 100 other AI Customers

Highest in Execution, Furthest in Vision

Everpure is a Leader in the 2026 Gartner® Magic Quadrant™ for Enterprise Storage Platforms

Gartner and Magic Quadrant are trademarks of Gartner, Inc. and/or its affiliates.

Gartner does not endorse any company, vendor, product or service depicted in its publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner publications consist of the opinions of Gartner's business and technology insights organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this publication, including any warranties of merchantability or fitness for a particular purpose.

This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. [The Gartner document is available upon request from Everpure.](#)

Gartner, Magic Quadrant for Enterprise Storage Platforms, Jeff Vogel, Joseph Unsworth, Julia Palmer, Chandra Mukhyala, 19 August 2026.

Everpure was previously Pure Storage

Figure 1: Magic Quadrant for Enterprise Storage Platforms



© Gartner, Inc.


Gartner.

AI Has Moved From Experiment to Production


Inference workloads could make up more than 40 percent of data center demand in 2030, growing 35 percent CAGR until 2030.

Global data center demand by workload, 2025–30

Global data center demand, gigawatts



Global data center demand, %



“McKinsey projects inference demand to grow ~35% annually through 2030 and overtake training as the dominant AI workload.”

40%+

of data centre demand is inference by 2030

35%

annual growth in inference demand to 2030

Source: McKinsey Data Center Demand Model, 2025–30.

Agentic Apps Are Being Adopted Even Faster

79%

say AI agents are already adopted
(PwC, 2025)

62%

experimenting with agents; 23%
scaling (McKinsey, 2025)

61%

of CEOs adopting agents and
preparing to scale (IBM, 2025)

40%

of enterprise apps to embed task-
specific agents by 2026 (Gartner)

Agentic apps are the default, not the experiment

4/5 enterprises already run AI agents. The gap is scaling them into production, not getting started.

Sources: [pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html](https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html) | [mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai](https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai)
[ibm.com/thought-leadership/institute-business-value/c-suite-study/ceo](https://www.ibm.com/thought-leadership/institute-business-value/c-suite-study/ceo) | [gartner.com/en/newsroom](https://www.gartner.com/en/newsroom)



THE REALITY

> 50%

of enterprise AI proof-of-concepts never reach production.

~ 95%

of enterprise AI initiatives produced no measurable impact.

Source: IDC — "Why AI Projects Fail: Revealing the Data Infrastructure Gap"

Source: MIT — "The GenAI Divide: State of AI in Business 2025"

Data Readiness Is the #1 Reason AI Stalls

Most often cited by IT leaders in a 2026 IDC survey as the top issue blocking AI progress

Source: 2026 IDC Global AI Readiness Report



AI-Ready Data Means Four Things

Four criteria. Most enterprise data fails all four.

01 Trustworthy



Curated

Duplicates, boilerplate, corrupted files, and sensitive content filtered out before they reach the model. Safety labels applied.

02 Governed



Classified & access-controlled

Permissions match human access. An analyst reads a doc — that doesn't mean AI surfaces it to everyone. Without this, you lock everything or compromise compliance.

03 Traceable



Indexed & vectorized

Converted to vector embeddings that capture meaning — stored so AI can retrieve. Until this happens, your data doesn't exist from AI's perspective.

04 Findable



Accessible & searchable

Available to AI models via APIs and MCP tools. Searchable and easily integrated into existing tools and systems.

And DIY approaches are complex and resource-heavy.

Getting There Takes Months Agents Won't Wait

Agents & Agentic Apps Can't Wait....

Identify & Discover

1–2 weeks

Data engineer

Define business requirements and identify relevant data sources.

Ingest

2–4 weeks

Data engineer

Load data using LangChain, Fivetran, Airbyte, AWS Glue, Azure Data Factory.

Curate

3–8 wks

Data Scientist / Analyst

Clean and prepare data using Python, Pandas, Databricks.

Transform

4–10 wks

ML engineer

Experiment with chunking strategies using LangChain, Prefect, Kubeflow.

Vectorize

5–12 wks

ML engineer

Generate vector representations using Hugging Face, OpenAI, Cohere.

Index

6–14 wks

ML / Data engineer

Load embeddings into vector DBs like Pinecone, Weaviate, Milvus.

Serve

7–16 wks

ML engineer

Build the LLM application using LangChain, SageMaker, Google Vertex AI.

2–4 months of pipeline work before your AI can read a single document.

Data Stream Makes What You Store AI-Ready

An **integrated AI data solution** that turns raw data you already store into **AI-ready, governed data** — fully automated.

NVIDIA AI Data Platform

GPU Accelerated

API Accessible

Backed by FlashBlade

AI-Ready Data

Turns the data you already store from a passive archive into an active AI delivery vehicle. Reads the formats your data lives in — PDF, DOCX, PPTX, images & more.

Active, Not Passive

Consume AI, automate pipelines

Collapses months of pipeline work into a packaged service. Free your ML platform team for higher-value work.

Months → Minutes

Sovereign & Secure

Runs only on infrastructure you own, as part of the secure Everpure platform. Content never leaves your environment.

Your Data Never Leaves

Inside the Service

INGEST



CURATE



CLASSIFY



EMBED



INDEX



RETRIEVE




GENERATE

Simplified "AI Ready Data-in-a-Box" Everpure™



Automate and orchestrate the entire journey from raw data to AI-ready data



Make all data AI-ready



In-place data processing



Continuous processing



Enterprise security & governance



Scale on data needs



But Agentic Apps Need More Than AI-Ready Data

1

Access to data

Live, governed access to source systems

2



GPU integration

Accelerated compute for training and inference

3



AI models

Frontier and fine-tuned models, swappable

4

High-speed storage

Low-latency I/O that keeps GPUs saturated

5



RAG pipeline

Chunk, embed, retrieve and ground responses

6

Semantic search APIs

Vector search exposed as a shared service

They All Want the Same Asset: Your Data

Your data.

Every AI model, every copilot, every intelligent app and workflow runs on one fuel: your proprietary data. The more sensitive, the more valuable — and the more at risk. Enterprises are discovering that AI ambition and data control are in direct tension.

>50%

of AI pilots

never reach production — governance gaps are the #1 cause

Source: IDC — "Why AI Projects Fail: Revealing the Data Infrastructure Gap"



Each App Takes Its Own Copy of It

Same six needs — but each app fulfills them differently

Microsoft 365 Copilot

TO WORK, IT NEEDS

Your content indexed in Microsoft Graph

WHERE YOUR DATA LANDS

Microsoft's semantic index, plus connector copies of outside sources

WHO ENFORCES ACCESS

Microsoft 365 item permissions

Salesforce Agentforce

TO WORK, IT NEEDS

Documents and transcripts ingested into Data Cloud

WHERE YOUR DATA LANDS

Data Cloud objects and its vector index

WHO ENFORCES ACCESS

Salesforce sharing rules and trust layer

Glean

TO WORK, IT NEEDS

A crawl of every connected SaaS app

WHERE YOUR DATA LANDS

Glean's own unified index of your content

WHO ENFORCES ACCESS

Mirrored source permissions, rechecked at query

Source: Microsoft, Salesforce and Glean product documentation, 2026.


Every app you adopt means **another copy of your data** in someone else's index, under someone else's rules and how data is handled varies from app to app.



Why Give Up Control of Your Data?

Old World: Application-centric

The application was king, flanked by copies of data



Every workload demands another copy and the ETL, MDM, and BI stack that grew up around the application only compensates for the resulting fragmentation

New World: Data-Primacy

Data at the center. Apps and AI consume it, not own it



AI is the new consumer – it needs one canonical picture across every systems in real time, not a copy per application.

What If You Didn't Have To: Data Stream

DATA STREAM PLATFORM (1 product + 1 workflow + API enabled)

Access to data

GPU integration

AI models

High-speed
storage

RAG pipeline

Semantic
search APIs



FASTER TO BUILD

Weeks, not quarters

Teams wire agents to governed data and retrieval on day one instead of building platform plumbing.

LOWER COST TO RUN

One platform, not six

Consolidated data, GPU, storage and search removes duplicated tooling, integration and headcount.

IN PRODUCTION

Deployments that stick

Governance and observability are built in, so pilots survive real load.



So Bring AI to the Data, Not Data to AI

The instinct is to extract, copy, and move data into AI environments. That instinct is expensive, insecure, and unscalable. Instead, deploy AI inference and retrieval at the data's location — whether on-premises, in a private cloud, or within a governed data platform. The data never moves. The intelligence does.

- Eliminates uncontrolled data egress and exposure
- Reduces latency, cost, and compliance surface area
- Keeps sensitive data under existing access controls

Old Model

Copy data → move to AI environment → lose control

New Model

AI queries in-place → data stays governed → control preserved

Govern Once, at the Source



Classify sensitive data — PII, PHI, financial records — once, at the point of creation or ingestion. Attach policy to the data itself so that governance travels with it automatically, regardless of where it flows or which AI system consumes it.

→ **Classify Once**

Tag PII, PHI, and confidential attributes at ingestion using automated discovery and cataloging tools.

→ **Policy Travels**

Access controls, masking, and retention rules are embedded in the data asset — not the pipeline.


→ **AI Inherits Controls**

Any AI system querying the data automatically respects the attached governance policy. No exceptions.

Your Data Never Leaves. It Never Trains Someone Else's Model.

- **Sovereign & on-prem** — deployable into air-gapped, regulated environments. Nothing exposed to public model training.
- **Model-portable** — Claude, on-box Nemotron, or a local NVIDIA VLM. Swap the model; the governance doesn't move.
- **Governed & auditable** — answers scoped to real access controls, every agent call stamped against its contract.

Documents In, Answers Out



Your data stays where it is

read from the customer's own FlashBlade over NFS or S3

One secured boundary

OIDC at the door; each stream reachable by the readers granted it

Three ways to consume it

chat, semantic search, and an OpenAI-compatible API for your developers

Delivered on Storage You Already Own

A certified service that turns the data you already store into AI-ready, governed context — reachable from any app.



Certified NVIDIA AI Data Platform

DELIVERY - PARTNER-SHIPED & INSTALLED

Customer has FlashBlade + GPUs

Partner sells Data Stream only

Customer has FlashBlade

Partner sells GPU server, and Data Stream

Customer has GPUs

Partner sells FlashBlade, and Data Stream

Customer has nothing

Partner sells GPU server, FlashBlade, and Data Stream

SUPPORT- EVERPURE

Standard

8x5, included across tiers

Premium (Optional)

BCS 24x7, separate SKU

SUPPORT- THIRD-PARTY

NVIDIA

AI Enterprise (NIM, NeMo)

Compute Vendor

GPU Server

What It Looks Like When It Works

Transforming raw data into AI-ready intelligence with a full-stack, secure solution.

What it does

AI-Ready Data Stream

Turns passive storage archives into active AI delivery vehicles. Automatically ingests, curates, and indexes data into governed vector embeddings reachable from any source.

Drives 2-20x storage increase and drags FlashBlade adoption

Why it's different

Full-Stack Efficiency

Automates months of pipeline work into minutes. Seamless and certified Nvidia integration with AI Enterprise and NIMS. Built on FlashBlade for high throughput, ensuring sovereign security where data never leaves your control.

Secure data pipelines that never leave corporate control



Success Story: County Judicial System

"Hours of case research collapsed to seconds — with answers clerks can verify and records that never leave the county." Data Stream enables plain-language search across decades of files on sovereign infrastructure.

TAKEAWAYS

AI Success Isn't a Faster Box — It's a Governed Data Foundation

- **Win on the data foundation** — not just compute. AI-ready data is what moves a PoC to production.
- **Keep it sovereign** — real outcomes on your governed data, without exposing it to anyone's model.
- **Measure time to value** — the AI factory turns experiments into shipped applications.

Everpure x NVIDIA — the AI factory that turns your data into outcomes.

Thank you — Kaycee Lai · VP of AI, Everpure



Thank You