

# Financial Analyst Meeting

2026

# Forward Looking Statements and Non-GAAP Information

This presentation contains forward-looking statements regarding our products, business, operations and financial performance, including but not limited to our views relating to our future period financial and business results, including our expectations regarding rule of 40; our capital allocation priorities, including research and development investment, strategic mergers and acquisitions and the return of excess capital; our total addressable and serviceable addressable market estimates and our expectations regarding market growth and continued market share gains; the expected mix of core and new product revenue; demand for our products and subscription and consumption offerings, including Evergreen//One, and the related sales mix; our sales pipeline and the timing and magnitude of large orders, including sales to hyperscalers and large enterprises; our technology and product strategy and roadmap, including DirectFlash, FlashBlade//EXA, the Enterprise Data Cloud and our expansion into data intelligence; our opportunity in modern data software, scale AI, and hyperscale environments, our ability to meet hyperscalers' performance, price and other requirements, to expand with existing and land new hyperscale customers, and the timing and amount of hyperscale revenue; the expected contribution to our revenue from modern data software, scale AI, and hyperscale products; our ability to manage supply chain disruptions and procure sufficient flash and other components; the impact of component cost increases and NAND pricing cycles on our pricing, gross margin and share gains; the anticipated effects of our acquisition of 1touch; our expectations regarding product and technology differentiation, sustainability and energy savings for customers, new investments and partnerships; and the impact of inflation, currency fluctuations, tariffs or other adverse economic conditions.

Forward-looking statements include all statements that are not historical facts and can be identified by terms such as “anticipate,” “believe,” “expect,” “could,” “seek,” “estimate,” “intend,” “may,” “plan,” “potential,” “predict,” “project,” “should,” “target,” “will,” “would” or similar expressions and the negatives of those terms. The forward-looking statements in this presentation are based on information currently available to us and represent our beliefs and assumptions only as of the date of this presentation. Except as required by law, we assume no obligation to update these statements publicly, or to update the reasons actual results could differ materially from those anticipated, even if new information becomes available in the future.

Forward-looking statements are subject to known and unknown risks and uncertainties and are based on potentially inaccurate assumptions that could cause actual results to differ materially from those expected or implied. Additional risks, uncertainties and other factors are included under the captions “Risk Factors” and “Management’s Discussion and Analysis of Financial Condition and Results of Operations,” in our reports previously filed with the SEC, which are available on our website at [investor.everpuredata.com](http://investor.everpuredata.com) and on the SEC’s website at [www.sec.gov](http://www.sec.gov).

This presentation includes non-GAAP financial measures as defined by SEC rules. As required by Regulation G, we have provided reconciliations of those measures to the most directly comparable GAAP measures, which are available in the Appendix. We have not reconciled our forward-looking non-GAAP measures to the most directly comparable GAAP measures because items such as stock-based compensation expense and payroll tax expense related to stock-based activities that impact these measures are not within our control and/or cannot be reasonably predicted. Accordingly, reconciliations of these non-GAAP financial measures are not available without unreasonable effort.

Note about 52/53-week fiscal year: We use a 52/53-week fiscal year ending on the first Sunday after January 30.

# Opening and introductions

**Paul Ziots** Vice President Investor Relations

# Agenda

## TECHNOLOGY DEEP-DIVES

Live webcast · 9:00–11:05am

- |                |                                                                                                                                                          |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>9:00am</b>  | <b>Doors open for a light breakfast</b>                                                                                                                  |
| <b>9:30am</b>  | <b>Opening &amp; introductions</b><br>Paul Ziots, Vice President Investor Relations                                                                      |
| <b>9:35am</b>  | <b>Data primacy: Data management strategy for the AI era</b><br>Prakash Darji, General Manager DX, and<br>Ashish Gupta, General Manager, Data Management |
| <b>9:55am</b>  | <b>Building the foundation for Enterprise AI</b><br>Chadd Kenney, Vice President Product Management                                                      |
| <b>10:15am</b> | <b>Meet FlashBlade//EXA. More AI. Less waiting.</b><br>Rob Lee, Chief Technology and Growth Officer                                                      |
| <b>10:35am</b> | <b>Combined Q&amp;A</b>                                                                                                                                  |
| <b>11:05am</b> | <b>End session</b>                                                                                                                                       |

## MAIN SESSION

Live webcast · 12:00–2:40pm

- |                |                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>12:00pm</b> | <b>Everpure strategy overview &amp; expanding our opportunity</b><br>Charles Giancarlo, Chairman and CEO                                                                  |
| <b>12:30pm</b> | <b>Everpure's fundamental advantages &amp; right to win</b><br>Rob Lee, Chief Technology and Growth Officer                                                               |
| <b>1:10pm</b>  | <b>Foundation for further growth &amp; long-term financial framework</b><br>Tarek Robbiati, Chief Financial Officer                                                       |
| <b>1:55pm</b>  | <b>Q&amp;A</b><br>Charles Giancarlo, Tarek Robbiati, Rob Lee,<br>John “Coz” Colgrove, Founder & Chief Visionary Officer,<br>Bill Cerreta, GM of Hyperscale, Prakash Darji |
| <b>2:40pm</b>  | <b>End session</b>                                                                                                                                                        |

# Welcome

**Charles Giancarlo** Chairman and Chief Executive Officer

# Data primacy: Data management strategy for the AI era

**Prakash Darji** General Manager, DDX

**Ashish Gupta** General Manager, Data Management

## KEY TAKEAWAYS

# Everpure is increasing our value & leverage within enterprises

- 01 Everpure is focused on helping customers **maximize the value of data** in the world of AI while solving data integration costs and challenges
- 02 Everpure's right to play comes from the proven NPS and trust of **serving mission critical enterprise data** for all structured and unstructured applications
- 03 Everpure's right to win comes from a unique advantage of **having both storage and data context** at the source

# What is the biggest data management challenge?

## Prior to AI



Data Integration associated costs are a bulk of data management spend

\$

## Post AI



Agents need shared context across systems and typically, this is high data copy and integration


\$\$\$

# Environments weren't architected for AI



## Fragmented context

Multiple platforms with partial views of data mean multiple teams, tools, and support contracts



## Data movement tax

Copying data between systems wastes time and introduces risk and bypasses security and governance



## Integration complexity

Each connection point becomes a potential failure mode and creates multiple representations of data



# Data primacy makes your architecture ready for AI



## Shared context

Data understood universally across all applications and silos



## Data-reuse

Universal data objects act as source of truth and used without copy where data is decoupled from apps



## Plug-and-play workflows

Applications, analytics, and agents consume and contribute data, but don't own the data

# Data primacy...

**Reduces  
your data  
integrations**

**Reduces  
data  
copies**

**Increases  
data  
coherence**

**Reduces  
cyber-attack  
surface**


**Optimizes  
token  
spend**

# Four fundamental long-term beliefs about data

Lead to storage & data management becoming more tightly integrated

01

Data will be distributed




02

Business logic as data




03

Context will be stored with data



04

Shared context needed for AI



A storage operating system will need to store data context and provide shared context

# Enterprises must own their own data

Data sovereignty is about ownership, not location

## A STORAGE PROBLEM

### Data Residency

Where does the data sit?

Which performance tier / array is it placed in?

Which protection group or policy is it placed in?

What cost tier is it placed in?



## A DATA PROBLEM

### Data Sovereignty

Who owns it?

Who is allowed to train on it?

Who can use it?


Does it still differentiate us?



## THE MORE VALUABLE AI BECOMES

**The more valuable control of the underlying data becomes**

# Traditional competitors are working against data sovereignty



## Everpure's strategy

- Data Sovereignty is a north star
- Broad interop with ecosystem
- Roadmap focuses on expanding capability for data in place

## Integration vendors

- Mainly chasing AI Agent Studios & Orchestration vs Data Contextualization


# What is data intelligence?

Understand  
the Data

## Universal Data Intelligence

Govern the Data

Store the Data



### Discover

Finding and  
connecting to data



### Classify

Identifying and  
labeling metadata



### Contextualize

Adding information about  
business purpose,  
meaning and relationships

# Data intelligence already works where the data already sits

Expansion of accessible budget as relevance of  
data intelligence exists for multiple personas



**Key customer  
use-cases:  
Data intelligence  
already adopted at  
large enterprises in  
regulated industries  
across all personas**

**14,000**

**databases scanned in under two weeks**

Global 100 card and payment leader. DSAR time cut from days to minutes.

**91M**

**card records found unprotected, then fixed**

National payment processor. A liability nobody knew was open.

**2 months**

**early to GDPR compliance**

European banking leader. Four weeks from proof of concept to production.

# Everpure's EDC is the key leverage to serve the data that matters most

Most companies come at the AI data problem from unstructured content, analytics or models. We start from the mission-critical systems of record.

**01**



**Enterprise trust,  
measured**

**84**

NPS, audited

**~64%**

of the Fortune 500

**02**



**Mission-critical,  
today**

**~15,000**

global customers

**03**



**Unified data  
plane**

**All workloads**

archive to AI

+Net Promoter Score Per Audited NPS customer score for 2025

# Better together: Storage and data compound


A better together platform solving both data residency and data sovereignty

## DATA LEVERAGE

Total visibility gives Everpure a view of Share of Wallet and allows us to move workloads to Everpure more seamlessly


## STORAGE LEVERAGE

Our platform becomes the turn-key way to show what context with data is worth, and to deliver it into AI use cases



# Data primacy roadmap starts with data intelligence

...and storage management benefits from data intelligence to provide data residency



# Key outcomes from Everpure

## SOVEREIGNTY USE CASES

**Secure data  
enclaves**

## RANSOMWARE USE CASES

**Detect anomalies  
& verify data  
integrity**

## AGENT POLICY USE CASES

**Deliver high  
availability data  
to agents**

## AI USE CASES

**AI quality &  
GPU efficiency**

## KEY TAKEAWAYS

# Everpure is increasing our value & leverage within enterprises

- 01 Everpure is focused on helping customers **maximize the value of data** in the world of AI while solving data integration costs and challenges
- 02 Everpure's right to play comes from the proven NPS and trust of **serving mission critical enterprise data** for all structured and unstructured applications
- 03 Everpure's right to win comes from a unique advantage of **having both storage and data context** at the source

# Building the foundation for Enterprise AI

**Chadd Kenney** Vice President of Product Management

**UNIFY**




**One platform for  
Enterprise Core + AI**

**CONTROL**



**Autonomy across the  
data estate**

**OPTIMIZE**



**Optimize AI data  
and workflows**

**Expand platform value → Scale through autonomy → Improve AI economics**




# One platform for Enterprise Core + AI

AI expands the value of a platform already deployed across the enterprise

# AI makes a fragmented data estate increasingly unsustainable


Every system is a silo: its own console, policies, SLAs, and tools. No centralized management, compliance, or governance



**AI amplifies the cost and complexity of fragmentation**

# The opportunity is to manage data, not storage

Remove the need for customers to manage storage, so they can start managing their data



# Enterprise data cloud brings the entire data estate together

UNDERSTAND

## Data Intelligence

### Context and meaning

Enterprise-wide data context and understanding where governance follows the data

CONTROL

## Intelligent Control Plane

### Autonomous governance

Apply intent and governance to data, to stop managing storage and start managing data

STORE + SERVE

## Unified Data Plane



### All workloads

Enterprise core, core AI, backup, archive, analytics, file, object, block, and cloud


## Evergreen Architecture

### One foundation, always modern

Continuously modern, resilient, efficient and non-disruptive infrastructure

# One unified data plane extends our value to AI

Covering every enterprise use case with a wide range of performance and economics from Archive to AI



# Evergreen is a subscription to innovation

Systems never get old; they get the newest innovation in place. You cannot build a unified data plane on a shaky foundation

Data intelligence

Intelligent control plane

Unified data plane



**Evergreen Architecture**

Stateless, modular systems. Controllers, flash, and software evolve in place around the data.

## Always modern, never down


Hardware and software upgrades without downtime or migration keeps system new

## 15 years of innovation

9 hardware generations and hundreds of software releases, all stateless and modular by design

## 200 new features per year


Continued R&D Investment enables a subscription to innovation



# DirectFlash delivers industry-leading density, reliability and efficiency

Efficiency improvements compound across every terabyte on the platform and every generation of flash

DATA ESTATE GROWTH    Petabytes → Exabytes




150 TB → 300 TB+

DirectFlash modules

Annualized failure rate by quarter

— Off-the-shelf SSD    — DirectFlash



**>10×**

**More density**

More usable data in every rack

**>5×**

**Less power**

Less power per usable terabyte

**>10×**

**Better reliability**

0.13% vs 1.2% annualized failure rate

**85%**

**Lower operating cost**

Versus the disk estates replaced




# Autonomy across the data estate

As AI increases complexity, Fusion allows software value to scale faster than operational overhead

# Fusion turns intent into autonomous action

Intent, autonomy, and governance delivered through Fusion. Value scales with the fleet and software, rather than with customer administrative headcount



## Define Intent

Intent, policy, protection, placement, and service levels are declared once, for the whole estate

## The Platform Governs

Placement, protection, recovery, and remediation happen without a person watching a console

## Value Increases

Every additional system under the control plane is operated by the same software, not more headcount

# Context and governance follows the data

## FOLLOWS THE DATA

Service level

Metadata

Identity

Policy

Protection state

Business context

SEMANTIC DATA GROUPING

Context, policy, and access follow the data wherever it goes, so the infrastructure can manage it autonomously through intent

## INTENT, EXPRESSED ONCE

*"This dataset must remain in Germany."*

*"Only approved models may use this dataset."*

Datasets organized by meaning and use, not by the storage system

ENFORCED ACROSS COPIES · SYSTEMS · LOCATIONS · REGIONS · APPLICATIONS · MODELS · AGENTS

Intent becomes governed action across the entire data estate

# Data context + storage context creates more intelligent control

The control plane knows the policies. Data Intelligence knows the meaning. **Both accumulate in one platform, on data that is already there.**

## Storage Context

Governance and application awareness the Intelligent Control Plane already enforces across the fleet

Policy

Protection

Placement

Service level

Application owner

Residency

Performance

Recovery

+  
Super  
Context

## Data Context

Meaning, semantics, and relationships that Data Intelligence discovers in the data itself

Meaning

Classification

Business entities

Relationships

Sensitivity

Lineage

Freshness

Ownership

Better governance

More trusted AI

More accurate  
autonomous decisions

Policy applied to what  
the data actually is

Better data discovery  
and reuse

# Fusion adoption expands software value across our installed base

New capabilities arrive through the existing Evergreen software relationship. **Software distributes faster than hardware.**

**2,300+** customers enabled with Fusion / Enterprise Data Cloud

**2x** adoption doubled in 6 months

## options


Fusion has delivered **savings in the hundreds of hours per month** across our infrastructure team, **freeing engineers to focus on higher-value initiatives**...

**80%**

reduction in management overhead

# Evergreen//One turns growth into recurring consumption

Customers define the outcome. Everpure operates and evolves the infrastructure. Unpredictable data and AI growth becomes recurring consumption rather than a one-time infrastructure forecast.



Evergreen® //One

## Declares the outcome

Capacity · Performance · Availability  
Protection · Recovery · Efficiency

## Deliver the SLA

Operates · Upgrades · Maintains buffer  
Rebalances · Protects · Expands

## Scales with demand

More workloads · More data · More AI  
More performance · More locations

**Evergreen removes the need  
to predict the technology.**

**Evergreen//One removes the  
need to predict the demand.**




# Optimize AI data and workflows

Move beyond storing AI data to improving the performance and economics of AI itself

# Enterprise data becomes the foundation for AI

Core AI is another set of workloads on the same Enterprise Data Cloud. The proprietary data it depends on is already there.



**The data AI needs is already on the platform**

---

**No additional silos or copies**

---

**Both workloads need the same things**

# We move beyond storing AI data to improving AI economics

Data Stream activates the data. KVA activates the data platform inside the AI runtime.

## Data Stream

Activate the data

---

**AI ready data**

---

**Less data movement**

---

**Governance follows the data**

## Key Value Accelerator

Optimize the runtime

---

**Up to 20x faster inference**

---

**Lower GPU consumption**

---

**Better cost per token**

# A platform unlike anything else in the industry

Enterprise AI requires a data foundation that can **unify** the estate, **control** it autonomously, and **optimize** how AI uses data

## UNIFY

### Growth of the platform

Enterprise Core + AI on one architecture

---

**More workloads. More data.  
More consumption.**

## CONTROL

### Scale software value

Intent, policy and governance  
across the estate

---

**2,300+ customers enabled  
2× adoption in 6 months**

## OPTIMIZE

### Improve AI economics

Optimize data, inference and  
GPU utilization

---

**Up to 20× faster inference less data  
movement and duplication**



---

15 years of innovation brought together enables one platform for the next era of enterprise data + AI

# Meet FlashBlade//EXA. More AI. Less waiting.

**Rob Lee** Chief Technology and Growth Officer

# AI-attached storage demand is set to grow ~1.6x by 2030



## Agentic AI drives inference growth

Knowledge workflows, customer support and coding



## High-performance storage tracks GPUs

Every GPU deployment pulls storage demand with it



## Neoclouds fill the GPU demand gap

Racing to supply the capacity hyperscalers leave open

\*Includes Core, Core AI and Scale AI TAM, presented on a calendar year basis.

# Storage is a hidden bottleneck in the GPU race

Where the focus has been:

## GPU Deployment

- Securing GPU supply
- Securing surrounding equipment
- Power, space, facilities



Where the challenges really are:

## GPU Utilization

- Stability and reliability
- Data and tenant onboarding
- Scaling performance and operations
- Evolving workflows create new performance bottlenecks

**Most AI clusters run below 85% GPU utilization because storage can't keep up. Every percentage point is real money.**

# Performance and consistency, under load



## Data gravity

Multi-petabyte training sets must live close to GPUs without becoming siloed.



## Multimodal and concurrency

Legacy parallel file systems are complex and were designed for legacy HPC. Modern AI is dynamic, multimodal, highly concurrent. At-scale is where most systems collapse.



## Reliability and operational simplicity

Downtime = lost revenues and SLA penalties. Neoclouds run lean teams, no dedicated storage staff – storage needs to “just work”, at most demanding scales.



## Checkpoint economics


Frontier-scale training writes tens of terabytes per checkpoint. Synchronous. Training stops while it writes. A full run can produce petabytes of checkpoint data alone.



## Security and compliance

Regulated data, tenant isolation, auditability—non-negotiable in financial services, healthcare, and government.

# FlashBlade//EXA: metadata and data scale independently



## METADATA PERFORMANCE

### Key-value database on FlashBlade

- Highly consumable metadata
- Native NFS solution, no database
- Scale to trillions of files



## SCALE

### Grow each component as needed

- 1 TB/s to 50 TB/s
- 100 PB to multiple EB
- Near-unlimited namespaces



## PREDICTABILITY

### Consistency when it matters most

- No bottlenecks
- Consistency for thousands of servers
- Pricing predictability

# Performance driven by software

## KEY CHALLENGES

### Scaling metadata as performance and capacity scales


Much focus is placed on data transfer rates, but metadata performance is just as important and harder to scale.

### Organizing metadata for parallel access

Traditional methods for organizing metadata are too serialized – designed in the era of HDD. Parallel access to data starts with Parallel access to metadata.

### Optimizing coordination and performance scalability

Running in parallel requires coordination. Minimizing the coordination costs maximizes scalability.



## PURE'S APPROACH SINCE DAY 1

### KV-store based metadata

Optimized for maximum parallelism and properties of Flash. Key-Value based metadata allows for massive scalability and parallel performance, without introducing bottlenecks.


### Distributed architecture

Highly-optimized distributed transaction-engine allows highly concurrent access, unlocking parallelism and scalable performance by reducing coordination costs.

# Software performance, unlocked by unique hardware


## FlashBlade//S

High-performance storage for enterprise AI  
Proven AI leadership, serving 100's of AI customers



## FlashBlade//EXA

Extends FlashBlade to highest-scale AI training. Purity Software,  
on massive scale-out open architecture



# Delivering real results

Decreased TCO

## 58% lower

Including 25% lower power and cooling,  
75% lower administration

Resilient performance

## #1

In performance benchmarks versus the closest  
competitors, MLPerf and SPECstorage

Massive scale

## 10,000+

GPUs supported, from enterprises to neoclouds



### Lower total cost

Operational overhead  
reduced 40–60%



### Faster innovation

Launch and iterate AI services:  
high-throughput data prep,  
large-scale training,  
low-latency inference



### Reduced risk

One platform simplifies security,  
compliance and disaster recovery



### Future scale

Grow from terabyte to exabyte  
scale seamlessly

# FlashBlade//EXA

## The #1 AI Storage Platform in MLPerf® v3.0



# #1

**in EVERY checkpointing  
& KV-cache metric submitted**

Closed Division (Standardized)

# 877 GiB/s

**peak checkpoint write throughput**

1T model, 30 data nodes

# 85,736 tok/s

**KV-cache throughput**

Llama 3.1 8B, storage-only


MLPerf Storage v3.0, Closed division: standardized benchmark, identical workloads across all vendors. Checkpointing tested at 1T (1250B) parameter scale — the storage load of a 1,024-accelerator training cluster, driven by 32 client hosts. Full results: [mlcommons.org](https://mlcommons.org)

© 2026 Everpure | Financial Analyst Day

# AI-native Neocloud removes the storage ceiling



## BEFORE




Limited to 30–40 GB/s

Large engineering staff to operate

Complexity & 3rd-party software dependency

## WITH FLASHBLADE//EXA



100+ GB/s per node

Predictable linear scaling

2× faster than the nearest competitor in under half a rack

## 256 //EXA nodes

**Consistency at scale**

In production. Same performance, same latency

## +10-20% overall

**Faster training and inference**

Versus the prior configuration in head-to-head testing

## 90%+

**Sustained GPU utilization**

On B300 with FlashBlade//EXA. Storage is no longer the constraint

# Combined Q&A

**Prakash Darji** General Manager, DDX

**Ashish Gupta** General Manager, Data Management

**Chadd Kenney** Vice President of Product Management

**Rob Lee** Chief Technology and Growth Officer



**Please join us for lunch.**

**Programming will resume at 12:00pm PT.**

