



# HITACHI

---

## **LATAM Partner Summit FY26**


Built to Grow. Powered by Data.




# HITACHI

# Hitachi VSP 360 and AI Insights e Observability

Italo Cocentino  
Hybrid Cloud Sales Leader



## VSP 360








# VSP 360

Transforming the VSP One Experience

*Control*

*Observe*

*Govern*




## Ecosystem Integration




## VSP 360 Platform



## VSP One



# VSP 360 and AI Insight and Observability Service




# Hitachi is delivering integrations since 2010...



# Hybrid AI Platform Drives Real Outcomes

Observability built for humans doesn't scale to AI





**Agentic AI** ➡ **Faster MTTR** ➡ **Lower Operational Costs**


# Map your IT Infrastructure to your Business Services

## Auto Discovery and Service Views


- On-Prem or cloud Deployment Architecture
- End to End Topology Mapping (Infrastructure to Application)
- Broader Library of 3<sup>rd</sup> party heterogenous coverage
- Agentless Asset Discovery
- Enhanced application layer observability capabilities
- Strong integrations with multiple vendor technologies.  
Specifically, Brocade/Cisco Fabric/Network integrations that underpin all other connected devices.



# Trend Matcher





# Trend Matcher



# The Deepest & Broadest Observability Platform


## Client Environment



## Client Business Outcomes



# Data Collection





# Intelligent Resource Optimization




Ready-to-use dashboards for immediate cost optimization of on-prem & cloud costs

See who is using the most resources

Enhance communication and transparency, and break down silos



See storage capacity and cost



Reallocate under-utilized resources to reduce infrastructure spend

## Most Utilized Clusters

### Top Clusters by CPU

[Show More >](#)

| Cluster     | Total CPU | CPU Utilization % |
|-------------|-----------|-------------------|
| staging-new | 4 Cores   | 18.82 %           |
| main-new    | 19 Cores  | 11.43 %           |

### Top Clusters by Memory

[Show More >](#)

| Cluster     | Total Memory | Memory Utilization % |
|-------------|--------------|----------------------|
| staging-new | 16.57 GB     | 56.23 %              |
| main-new    | 75.78 GB     | 26.15 %              |


### Top Clusters by Disk

[Show More >](#)


| Cluster     | Total Disk | Disk Utilization % |
|-------------|------------|--------------------|
| staging-new | 0.05 TiB   | 60.51 %            |
| main-new    | 0.57 TiB   | 44.81 %            |

Dashboards for storage, VM compute, and containers

# Hybrid Cloud Cost management (FinOps)




# On Premises Cost management




# Automations

Policy-driven alert response for self-healing and integration with existing IT service management systems


Reduce manual, repetitive tasks



Alert intelligence baked into the platform for remediation and self-healing



Define policies using natural language with Virtana Copilot




Self-healing through integration with orchestration platforms

# Application Observability Overview

## Key Features

- End-to-End Transaction Tracing
- Automated Root Cause Analysis
- Intelligent Log Correlation
- Synthetic & Availability Monitoring
- Kubernetes Visibility
- Hybrid Infrastructure Observability
- Agentic AI Automation
- MCP-Driven Context



# TracePath (Transactions) & Trace Observability

## Follow the transaction, not the noise




Trace business transactions end-to-end across services, clusters, nodes, and network paths.




Identify where time is actually spent and which dependency drives user impact.



Prioritize fixes by tying trace behavior to business impact and service health.



# Automated Root Cause for Business-Critical Apps



## Reduce MTTR from hours to minutes

Automatically pinpoint the true root cause across services and infrastructure.



## Protect SLAs and customer experience

Quickly identify which transactions, users, and endpoints are impacted and why



## Lower operational risk and burnout

Cut through alert noise and focus teams on the few issues that actually matter.

# Log Analysis and Correlation

## Transform logs into answers



Correlate logs with traces, alerts, and dependency context to pinpoint the failure chain.



Reduce false positives by clustering related issues into one coherent incident story.




Accelerate resolution by identifying the most relevant log signals for the impacted transaction.



# Synthetic and Availability Observability

## Know before customers do



Continuously test user paths to detect degradation early.



Separate true downtime from localized failures across DNS, network paths, and shared services.




Track availability and performance against business SLAs with clear evidence.

# Kubernetes Observability

## Make Kubernetes predictable

- ✓ Map services to the Kubernetes resources and runtime dependencies that govern performance.
- ✓ Identify contention and instability that traces alone cannot explain, including noisy neighbors and throttling.
- ✓ Connect pod and node conditions to the transactions they disrupt.



# CoPilot Agentic AI




## AI Agents

- ✓ Root Cause Analysis
- ✓ Querying Alert and Telemetry Data
- ✓ Writing Event Correlation and Enrichment Policies
- ✓ Creating Automated Response Policies
- ✓ Interacting with other AI Agents using MCP Server
- ✓ Documentation Questions / How to?



Pictured: Agent for creating automated response policies

# AI Factory Observability



- DYNAMIC TOPOLOGY**  
Automated discovery & dependency mapping eliminates blind spots
- INCREASE SERVICE RESILIENCE**  
Fast, accurate root cause analysis & automated remediation
- DRIVE COST EFFICIENCY**  
Visibility into GPU capacity, energy usage, and application-level expenses
- ENSURE OPTIMAL AI WORKLOAD PERFORMANCE**  
Insight into AI workload timing, contention & production impact
- FULL-STACK AI INFRASTRUCTURE OBSERVABILITY**  
Observability for all layers of your AI stack

# HITACHI

## Everflex AI Data Hub as a Service



Italo Cocentino  
Hybrid Cloud Sales Leader

# What is AI DataHub as a Service?

HITACHI

Hitachi Vantara

Why Hitachi Vantara

Products ▾

Solutions ▾

Services ▾

Partners ▾



Contact Sales



**EverFlex AI Data Hub as a Service  
Explainer Video**

Watch video →

*EverFlex AI Data Hub as a Service*

## Empower self-service AI data preparation

Enhance data operations and speed AI innovation while reducing cost and complexity of your distributed data landscape with Hitachi EverFlex.

Calculate TCO →


<https://www.hitachivantara.com/en-us/services/infrastructure-as-a-service/service-catalog/everflex-ai-data-hub-as-a-service>

# ... an Offer from our Everflex Catalog


| Domain                              | Purpose                                                                     | Key Products / Building Blocks                                                                  | Typical Use Cases                                                                                        |
|-------------------------------------|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Storage Services (STaaS)</b>     | On-premises block, file, object as well as software defined storage (SDS)   | VSP One (block/file/object), HCP - Hitachi Content Platform (object)                            | Databases, structured & unstructured data, application data, archival/object, multi-vendor observability |
| <b>Compute as a Service (CaaS)</b>  | Elastic compute for bare metal, VMs, containers with integrated networking  | Supermicro Compute, Cisco UCS Compute, Cisco Nexus/MDS networking, Brocade SAN                  | VM farms, container platforms, hybrid cloud                                                              |
| <b>Data Protection (DPaaS)</b>      | Managed backup, DR and cyber-resilient protection with defined RPO/RTO      | Hitachi Data Protection Suite + Commvault/Veeam; immutable copies                               | Operational recovery, compliance, ransomware recovery                                                    |
| <b>Ransomware Recovery Services</b> | Detect, protect and rapidly recover from attacks with immutable data copies | Immutable snapshots; HCP; automated recovery playbooks, Cybersense                              | Critical app recovery, air-gapped protection                                                             |
| <b>AI Data Hub as a Service</b>     | AI-ready lakehouse with data virtualization and governed access             | VSP One Block, Object, Cisco UCS, NVIDIA GPUs & Networking; Hitachi IQ Series, Zetaris Data Hub | Virtual Data Warehouse, Streaming Analytics, Data Governance, AI/ML Enablement                           |

# ... a Part of Hitachi DataLake House...

## Hitachi Data Lakehouse Reference Architecture



# ... a Replacement for some Pentaho use cases...



# ... a Way to reduce Customer TCO for AI projects...

|                                                    | Assumptions                                                    | Without AI DataHub | With AI DataHub     |
|----------------------------------------------------|----------------------------------------------------------------|--------------------|---------------------|
| DBU License Costs*                                 | An average of \$300,000 Databricks DBU per month               | \$3.6M             | <b>\$1.2M</b>       |
| EC2 Compute Costs*                                 | An average of \$600,000 AWS EC2 per month                      | \$7.2M             | <b>\$3.6M</b>       |
| Storage Costs**                                    | Amazon S3+MinIO current                                        | \$3.5M (est)       | <b>\$2.0M (est)</b> |
| Labor Cost**                                       | Current manual vs. auto pipelines, no data movement            | \$2.0M (est)       | <b>\$0.8M (est)</b> |
| Other Operational, Support and Maintenance Costs** | Improve on current significant data/app sprawl, + efficiencies | \$3.0M (est)       | <b>\$2.0M (est)</b> |

In total, \$9.7M estimated cost savings annually.

Value Engineering/TCO team to validate the Storage, Labor and Other Operational, Support and Maint costs

# What is DataHub as a Service? (ChatGPT)

Good question 🤖 — it sounds like marketing buzz until you peel it apart.


Here's the **plain-English** version:

---

## AI Data Hub as a Service = A Managed Data Environment for AI & Analytics

- Think of it as a **ready-to-use data platform** that a vendor (like Hitachi Vantara) hosts and runs for you.
- Instead of you building your own data lake, pipelines, governance, compute clusters, GPUs, storage, etc., they bundle all of it and deliver it as a service (like renting, not buying).
- It's designed to **make data "AI-ready"**:
  - collect data from different sources (databases, files, SaaS apps, sensors, cloud, on-prem),
  - clean and organize it (governance, security, quality checks),
  - and then let you **analyze it or feed it to AI/ML models** — without heavy manual plumbing.

# The AI Data Hub Journey



AI Data Foundation

Platform for Scalable AI



Data Readiness & Preparation

Activate Distributed Data



Governance & Enterprise Outcomes

De-Risk & Scale Efficiently




Financial  
Outcomes

Reduced Data Movement | Lower Infrastructure Waste | Platform Rationalization | Predictable ROI

TCO Reduction is not a phase -> It's a Continuous Outcome

# The AI Data Hub Advantage


## Traditional Data Pipelines



- Data Silos
- Complex Pipelines
- Variable High Cost

VS.

## AI Data Hub



- Unified & Governed Data
- Fewer Pipelines
- Cost Efficient

# Industry Specific AI Data Preparation

| <br>Healthcare | <br>Financial Services | <br>Telecom | <br>Retail | <br>Information Technology |
|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| Early Disease Detection                                                                         | Credit Scoring                                                                                          | Predictive Maintenance                                                                         | Market Basket Analysis                                                                        | Incident Management                                                                                           |
| Digital Care                                                                                    | Loan Analysis                                                                                           | Industry Data Exchange                                                                         | Next Best Offer                                                                               | Capacity Planning                                                                                             |
| Cardiac Pathways                                                                                | Anti-Money Laundering                                                                                   | Workforce Management                                                                           | Customer Lifetime Value                                                                       | Configuration Drift                                                                                           |
| Specialists Referrals                                                                           | Credit Fraud                                                                                            | Truck Roll Reduction                                                                           | Supply Chain Optimization                                                                     | Process Automation                                                                                            |


AI Data Hub industry models shift enterprises from data engineering-heavy architectures to outcome-driven AI platforms—delivering governed, AI-ready data across the hybrid cloud, with support for 18 industries. Additional industries are life sciences, insurance, manufacturing, government and public sector, transport and logistics, rail and water utilities, energy, construction, education, defense, media and content, entertainment and gaming, and human resources.

# Workshop Methodology

## Outcomes



## Activation



## Foundation



### Align on Outcomes

- Define target business outcomes, cost, risk, efficiency
- Align stakeholders on success metrics and KPIs
- Identify measurable outcomes

### Outputs

- Top 3 to 5 outcome-aligned use cases
- Defined success KPIs and value targets

### Activate the Data

- Assess data requirements across enterprise sources
- Map data requirements to prioritized use cases
- Identify data gaps, quality issues, governance needs

### Outputs

- Data readiness snapshot
- Use case to data mapping
- Identified gaps and remediation actions

### Align to the Foundation

- Align requirements to target architecture and platform capabilities
- Align infrastructure, security, and governance model
- Plan deployment, operations, and lifecycle management

### Outputs

- High-level architecture blueprint
- Platform and service alignment
- Initial execution scope

# Workbench

AI Data Hub

NDP Fabric Builder

Schema Store

Data Catalog

Query Builder

Virtual Data Mart

Data Lineage

NDP File System

Data Pipeline

Access Control

User Management

Roles & Permissions

Validate

Save

Clear Model

BEN.TZANNES@ZETARIS.COM

Data Sources

Data Pipelines

loan\_crm\_analysis

+ SQL

+ Sources

+ Sink

+ RAG

+ Others

- semantic\_layer\_definitions (6)

- client\_final

- client\_master\_validation

- customer\_master

- employee\_hours\_worked

- loan\_crm\_analysis

- patient\_observations\_report

BLOB\_BANKING.Crm\_events\_all

| Name           | Type   |
|----------------|--------|
| client_id      | string |
| account_id     | string |
| type           | string |
| row_id         | string |
| date_received  | string |
| product        | string |
| sub_product    | string |
| issue          | string |
| sub_issue      | string |
| consumer_compl | string |

Chunking

Embedding

LLM Tokenizer

LLM

VS Sink

VS Search

combined\_ta

projection\_0

AZURE\_MSSQL\_BANK.Loan

Columns

Filter

Preview (projection\_0)

| client_id | account_id | type  | row_id | date_received | product                 | sub_product              | issue              |
|-----------|------------|-------|--------|---------------|-------------------------|--------------------------|--------------------|
| C00000404 | A00000330  | User  | 720    | 20/2/2019     | Credit card             | NULL                     | Delinquent accoun  |
| C00000404 | A00000330  | User  | 1863   | 2/11/2020     | Bank account or service | Checking account         | Account opening, c |
| C00000404 | A00000330  | User  | 17332  | 15/12/2017    | Bank account or service | Savings account          | Account opening, c |
| C00000404 | A00000330  | User  | 23012  | 24/5/2020     | Bank account or service | (CD) Certificate of depc | Account opening, c |
| C00000403 | A00000330  | Owner | 1340   | 17/5/2016     | Credit card             | NULL                     | Credit rep         |

Total: 3130, Elapsed: 2.671 s | Latest result: Success

# Data Discovery and Catalog Management

File Sources

Virtual File Sources - (8)

- AWS\_S3 (2)

- AZURE\_BLOB (2)

- BLOB\_BANKING (14)

- BLOB\_MEDICAL (11)

- HCP\_FILES (2)

- NETWORK\_FILES (14)

- RAW\_STORAGE (20)

- SALESFORCE\_CRM (7)

Database Sources

Virtual Database Sources - (17)

- SNOWFLAKE\_BANKING (8)

- SNOWFLAKE\_MEDICAL (9)

- TERADATA (2)

- TERADATA\_CLOUD (0)

- UNIVERSITY\_DATAWAREHOUSE (11)

- zetaris\_auditlog (10)

- zetaris\_metastore (128)

Table Metadata

Tags

Metadata Change History

Metadata Search

|                       |           |           |            |   |   |  |
|-----------------------|-----------|-----------|------------|---|---|--|
| client                | 0         | 0 Bytes   | -          | - | + |  |
| disposition           | 0         | 0 Bytes   | -          | - | + |  |
| disposition_update    | 0         | 0 Bytes   | -          | - | + |  |
| district              | 0         | 0 Bytes   | -          | - | + |  |
| loan                  | 0         | 0 Bytes   | -          | - | + |  |
| luxury_loan_portfolio | 1,678     | 407.25 KB | 21/11/2024 | - | + |  |
| orders                | 6,471     | 286.27 KB | 21/11/2024 | - | + |  |
| student_performances  | 0         | 0 Bytes   | -          | - | + |  |
| transactions          | 1,056,320 | 103.27 MB | 21/11/2024 | - | + |  |

UNIVERSITY\_DATAWAREHOUSE.student\_performances Column Metadata

| Column Name                 | Data Type | Constraints | Statistics | Tags | Description                                                     |
|-----------------------------|-----------|-------------|------------|------|-----------------------------------------------------------------|
| gender                      | string    | NULL        | -          | pii  | The gender of the student                                       |
| race/ethnicity              | string    | NULL        | -          |      | The raceethnicity group of the student                          |
| parental level of education | string    | NULL        | -          |      | The highest level of education attained by the students parents |
| lunch                       | string    | NULL        | -          |      | Type of lunch the student receives standard or freereduced      |

Tags and Descriptions have been auto generated

# Simplify Data Security Management

The screenshot displays the Hitachi Data Security Management interface. The top navigation bar includes icons for AI Data Hub, NDP Fabric Builder, Schema Store, Data Catalog, Query Builder, Virtual Data Mart, Data Lineage, NDP File System, Data Pipeline, Access Control, User Management, and Roles & Permissions. The 'Access Control' tab is highlighted with a blue box and an arrow pointing to the handwritten text 'Data Security'.

The main content area shows a list of Data Marts on the left and details for selected Data Marts on the right. The selected Data Mart is 'student\_360'.

**Data Marts - (5)**

- customer360 (7)
- HR\_USL (4)
- patient\_360 (8)
- student\_360 (4)
- supplier\_360 (5)

**UNIVERSITY\_DATAWAREHOUSE.Student\_performances**

Virtual Table: student\_performances

| Virtual Column  | Real Column                 | Data Type |
|-----------------|-----------------------------|-----------|
| gender          | gender                      | string    |
| race_ethnicity  | race/ethnicity              | string    |
| parent_educatic | parental level of education | string    |
| lunch           | lunch                       | string    |
| test_preparator | test preparation course     | string    |
| math_score      | math score                  | int       |
| reading_score   | reading score               | int       |
| writing_score   | writing score               | int       |

**UNIVERSITY\_DATAWAREHOUSE.Account**

Virtual Table: account

| Virtual Column | Real Column | Data Type |
|----------------|-------------|-----------|
| account_id     | account_id  | string    |
| district_id    | district_id | int       |
| frequency      | frequency   | string    |
| parseddate     | parseddate  | timestamp |
| year           | year        | int       |
| month          | month       | int       |
| day            | day         | int       |

**SNOWFLAKE.BALANCES\_INPUT**

Virtual Table: BALANCES\_INPUT

| Virtual Column | Real Column | Data Type |
|----------------|-------------|-----------|
| date           | date        | timestamp |
| balance        | balance     | double    |
| product_id     | product_id  | double    |
| customer_id    | customer_id | bigint    |

**Semantic\_layer\_definitions.Loan\_crm\_analysis**

Virtual Table: loan\_crm\_analysis

| Virtual Column | Real Column   | Data Type |
|----------------|---------------|-----------|
| client_id      | client_id     | string    |
| account_id     | account_id    | string    |
| type           | type          | string    |
| row_id         | row_id        | string    |
| date_received  | date_received | string    |

# Supported Data Sources

## Relational Databases (RDBMS):

- Connects to any database with a JDBC API (e.g., Oracle, SQL Server, PostgreSQL, MySQL, Snowflake, Redshift)
- Supports direct query pushdown for optimal performance

## NoSQL Databases:

- Integrates with NoSQL sources that offer JDBC connectivity (e.g., MongoDB, Cassandra)

## Cloud Data Warehouses:

- Native support for cloud platforms like Snowflake, Databricks, AWS Redshift, Google BigQuery, Azure Synapse

## File-Based Data Sources:

- CSV, Excel, XML, JSON, Parquet, ORC, Delta, Avro
- Connects to files on local/network drives, AWS S3, Azure Blob, and other object stores

## Streaming Data Sources:

- Connects to Apache Kafka and other streaming platforms for real-time analytics

## REST API Connections:

- Supports integration with RESTful APIs for custom data sources


## External Metastore Connection:

- Hive Metastore, AWS Glue Metastore for schema management



# From Power to Productivity: Optimizing Tokens per Megawatt


- **AI performance** is now limited by power and infrastructure—not just models—making tokens per megawatt a key efficiency metric.
- **Poor data management** wastes energy and compute by moving and reprocessing unnecessary data.
- Providing only **filtered, virtualized, high-value data** to AI models reduces waste and increases productivity per unit of power.
- This **improves AI ROI** and lowers total cost without adding power, cooling, or infrastructure.
- The result is **measurable cost and energy savings**, especially in hybrid cloud environments where efficiency drives scale.



**Simplifies adoption of GenAI and Agentic AI Use Cases**

# Business Outcomes Driven by Real Hitachi Data

AI Data Hub unifies live IT system data with generated document intelligence to optimize incident resolution, enable proactive IT operations, and enforce governance—without creating more data sprawl.



# AI Insights e Observability

## AI DataHub as a Service



# Obrigado!!! Gracias!!!