

HITACHI

Mission-Critical AI in Extreme Environments

An Open-Source Stack for On-Prem

October 2026


© Hitachi, Ltd. 2026. All rights reserved.

HITACHI

Mission-Critical AI in Extreme Environments

An Open-Source Stack for On-Prem

October 2026





Georgy Okrovertskhov

VP of Engineering, AI Solutions

Executive Overview

- **VP of Engineering, AI Solutions:** Driving advanced artificial intelligence architectural strategies and high-impact industrial solutions.
- **20+ Years Experience:** Deep domain expertise across IT and Software Engineering leadership.
- **Academic Background:** Holds a PhD in Physics, bringing analytical rigor to complex software infrastructure.

A practical guide to calling me at work / parties

Gosha

Because if you try ordering a beer using my full legal name, you'll stay sober waiting.

George

The standard expectation when living in America.

Greg / Gregory

Because... "y"-not?

Jorge

Pronounced Mexican-style; courtesy of my local California friends who insisted.



Georgy Aleksandrovich...

Reserved exclusively for official documents, severe reprimands, or when introducing me to royalty.

Extreme Operational Environments

HITACHI

Mission-critical AI requires robust on-premises infrastructure across diverse industrial settings.




On-prem architecture guarantees operational autonomy, data sovereignty, and deterministic latency.

The Dependencies Behind one AI task

HITACHI

An AI task depends on more than its model.



Four responsibilities of an OSS operating stack

HITACHI



A reference architecture for operating industrial AI within the site.

Industrial AI application		LangGraph or ADK	PRODUCT INTEGRATION Agentic Runtime Onboard and deploy agents to the selected runtime environment. Application logic stays with the app team.
RESPONSIBILITY	FOUNDATION AND COMPONENTS	OPERATING OUTCOME	
Run locally	Linux OS + Kubernetes Containers and local model serving	Execute within the site	
Control access	Keycloak + APISIX Identity and gateway access policies	Govern access to services	
Recover work	Temporal + agent persistence Define workflow and agent-state ownership	Resume recorded progress	
Maintain locally	Argo CD + local repositories Approved configuration and software artifacts	Manage controlled updates	
Shared local services	Persistent storage	Observability	

Industrial AI Technical Stack

HITACHI


A layered reference architecture from physical infrastructure to agentic applications.



A maintenance workflow that survives interruption

HITACHI

Illustrative on-prem design: a compressor inspection coordinated locally with Temporal.



Self-hosted Temporal

History and durable approval waits


Policy-based Activity retries

Temporal preserves progress. Application design makes recovery correct.

OSS across cloud, data center, and edge

HITACHI

Place each AI capability where its data, latency, and operating needs are best served.



HITACHI