




# Blueprint for Agentic AI in Industrial Operations

## From Real-Time Data to Autonomous Action

---



# Speakers



**Shashank Sharma**  
Dr. Product Marketing



**Kudzai Manditereza**  
Sr. Industry Solutions  
Advocate



# At a glance

---

## MISSION

Empower customers to get **value** from their data

## WE PROVIDE

### IoT Data Streaming Platform

Connect, stream, and govern data from edge to cloud in real time

## TRACTION

### 200+ Enterprise Mission Critical Customers

BMW, Lilly, Netflix, Audi, Air France-KLM, Mercedes Benz, Unilever, ...

## OUR DIFFERENCE

### Real-time distributed intelligence for Industrial Data

Reliable, scalable, & secure IoT data streaming across the enterprise

## FOUNDED

**2012**

Headquartered in Germany with offices around the globe

# Why Agentic AI in Industry Needs a New Operations Backbone

Industrial operations are under pressure. Three gaps block autonomous AI at scale:



## Infrastructure

Batch and point-to-point architectures cannot deliver the real-time response that autonomous agents require.



## Intelligence

Without a semantic layer, data streams carry no operational meaning—agents cannot reason about cause and effect.



## Trust

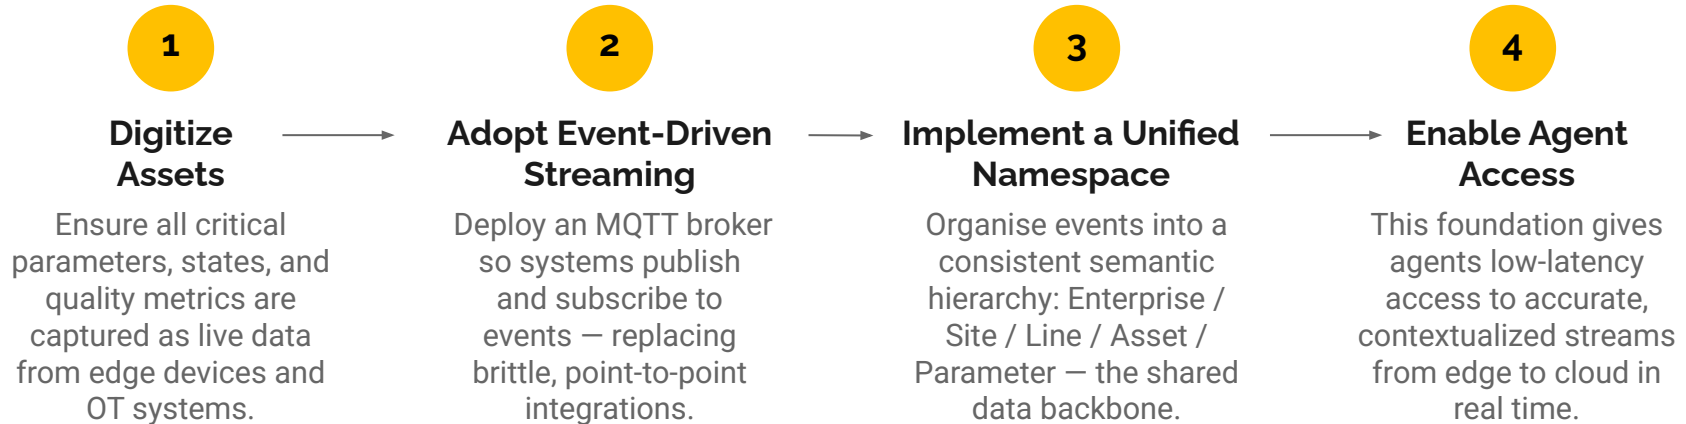
No governance, no autonomy thresholds, no safety guardrails. Agents that can't be trusted can't be deployed.

**Key takeaway:** Before deploying agents, fix the underlying data, semantic, and governance foundations that determine their success.



# Building the Real-Time Data Foundation with MQTT and UNS

## Step 1 of the Blueprint



**Key takeaway:** MQTT plus a Unified Namespace turns fragmented plant data into a single, real-time nervous system that agents can trust.

# Ontologies & Semantic Graphs: Making Data Agent-Ready


## Step 2 of the Blueprint

### The Problem

Hierarchical topic structures alone cannot express why events matter or how entities relate.

Siloed systems produce data that is structurally correct but semantically opaque – agents cannot reason about cause-effect relationships across lines, sites, and supply chains.

### The Solution



#### **Ontology-driven semantic graphs**

Capture equipment, processes, materials, and relationships – constraints and rules included.



#### **Standardised data models**

Reusable across plants for equipment, processes, materials, and quality metrics.



#### **Cross-domain causal reasoning**

Agents can infer: "a feeder issue affects quality, energy, and schedule across the line."

**Key takeaway:** A semantic layer turns raw streaming data into shared operational meaning that multiple agents can reason over simultaneously.



# Where Agentic AI Delivers Value First

Step 3 of the Blueprint – Four strategic domains



## Production Continuity

Coordinated predictive maintenance, real-time anomaly detection, autonomous scheduling adjustments.



## Quality Assurance

Continuous quality prediction, in-line adjustments, automated root-cause tracing across material genealogy.



## Resource Efficiency

Autonomous energy management, waste reduction, and dynamic resource allocation across production lines.



## Throughput Optimisation

Changeover optimisation, dynamic scheduling, multi-objective parameter tuning for maximum output.

## Autonomy Maturity

### 01 Diagnostic

Monitoring & alerting.  
Human decides.

### 02 Prescriptive

Recommendations.  
Human approves.

### 03 Autonomous

Closed-loop control.  
Governed action.

**Key takeaway:** Focus on high-value domains first and progress through staged autonomy – don't jump directly to full automation.

# A Practical Taxonomy of Industrial AI Agents

Design modular, reusable capabilities – not one-off bots



## Monitoring Agents

Continuously assess operational state, detect anomalies, and prioritize alerts across streaming telemetry from edge to cloud.



## Diagnostic Agents

Perform automated root-cause analysis across equipment, process parameters, and material genealogy data.



## Predictive & Optimisation Agents

Forecast failures and quality outcomes, and tune parameters for multi-objective performance targets.



## Coordination & Learning Agents

Orchestrate workflows across systems and continuously improve other agents based on operational outcomes.

**Key takeaway:** Agent roles – monitor, diagnose, predict, optimise, coordinate, learn – enable modular reusable capabilities across plants.

# Governance Frameworks for Safe Agentic Operations

Industrial agents directly impact equipment, product quality, and safety – demanding rigorous governance, not just model accuracy.

## Phase 1



### Design for Control

Define ownership, risk tiers, autonomy thresholds, and hard safety constraints up front. Non-negotiable before any agent touches operations.

## Phase 2



### Engineer Technical Controls

Emergency stops, circuit breakers, robust integration boundaries, and transactional integrity across all connected systems.

## Phase 3



### Operate with Oversight

Continuous monitoring, explainability, secure audit trails, and disciplined change management for evolving agents.

**Key takeaway:** Trustworthy agentic AI is engineered through explicit risk tiering, guardrails, and lifecycle governance – not prompt engineering.



# From Hierarchies to Distributed, Multi-Agent Intelligence

A shift in operating model, not just technology



## Shared Intent

Multi-agent operations replace siloed KPIs with shared objective functions that align all decisions to enterprise outcomes.



## Bounded Autonomy

Hard constraints and risk tiers keep agent actions within validated boundaries. Every agent knows its limits.



## Common Context

The UNS and ontology give every agent the same operational reality – a shared language across lines, sites, and systems.



## Human Co-Agency

Humans remain collaborators with clear escalation and handoff protocols – not bystanders to automated decisions.

**Key takeaway:** The shift to multi-agent operations is as much about operating model and shared intent as it is about technology.



# How HiveMQ Powers Streaming, Semantics & Agent Collaboration

One platform – the central nervous system for multi-agent industrial intelligence.

## HiveMQ Broker

Resilient, scalable MQTT event streaming at the core – connecting every edge device to the enterprise.

## HiveMQ Edge

Protocol translation from legacy OT into MQTT topics aligned with your Unified Namespace.

## HiveMQ Data Hub

Data normalisation, semantic modelling, and policy enforcement for agentic workloads.

## HiveMQ Pulse

Distributed intelligence layer – real-time observability across agents, streams, and operations.

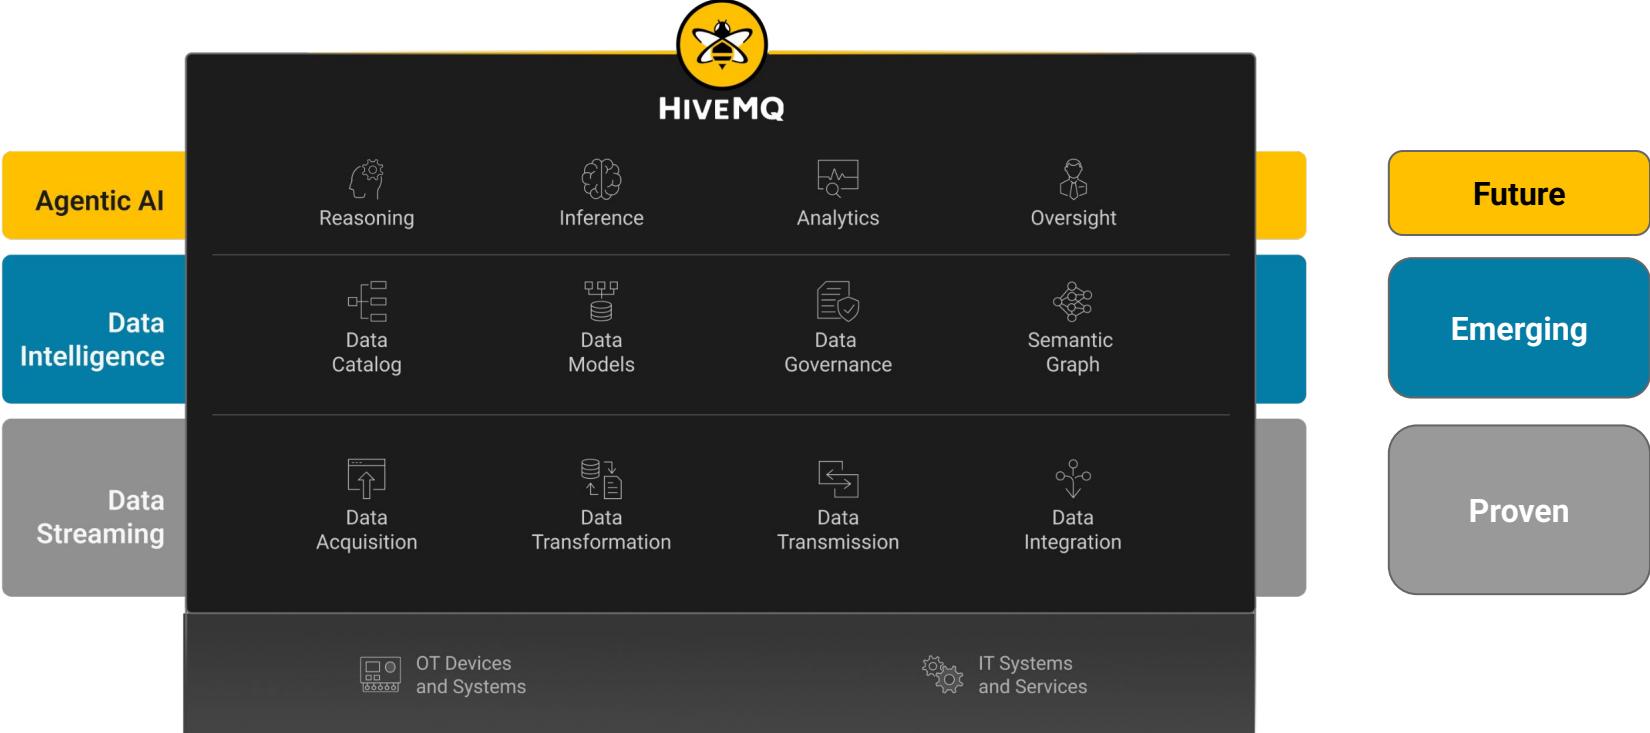
Edge → Broker → Data Hub → Pulse → Enterprise AI

**Key takeaway:** HiveMQ turns your MQTT and UNS investments into a scalable foundation for coordinated multi-agent intelligence.



# How HiveMQ Powers Streaming, Semantics & Agent Collaboration

One platform – the central nervous system for multi-agent industrial intelligence.



# Getting Started with Agentic AI in Your Operations

Where can I learn more?



**Assess & Prioritise:**



**Implement the Backbone:**



**Deploy Staged Agents:**



**Scale with Governance:**



To learn more, [download the whitepaper](#).





Thank you!

