

# VANTIQ



Platform technology

2025

Paul Butterworth

# Architectural Philosophy

# Problem

Opportunity to increase value by analyzing data as soon as it is available

Most organizations fail to exploit the real-time nature of this data electing to store it and analyze it in the background

Much value can be realized by analyzing and responding to the data in real-time.

AI is critical to situational analysis and situation response

But building these real-time, event-driven, AI-centric applications is difficult

# Difficulty: Intelligent Systems Depend on Properly Integrated Infrastructure

- Developer IDE
- Persistent Storage
- Data Streaming (input and output)
- Stream Analytics (CEP or processing data streams)
- Pub/Sub Style Event-Bus with Catalog and Ledger
- API Management
- Serverless functions (Scheduled, Triggered, Called Remotely)
- Web and Mobile Client Development
- Distributed Application Deployment (any geographical distance)
- Distributed and Federated Queries
- User/Security Management
- Real-Time Testing
- Resource Management and Reporting
- Auditing



# Difficulty: Intelligent Systems Require Novel Implementation Strategies

- Asynchronous messaging
- Reliable Messaging
- Event Brokers
- Asynchronous execution
- Distributed execution
- Databases
- Document storage
- Transformation systems
- Integration with external systems
- Access to real-time data sources
- Predictive and generative AI systems
- Fault tolerant infrastructure
- Authentication system integration
- Document storage

# Vantiq's Approach

Pre-integrate infrastructure critical for real-time applications

Do the hard system integration that supports performance, reliability, scalability, security throughout the integrated infrastructure.

Integrate flexible security model into the infrastructure

Build a set of tools that seamlessly integrates with the infrastructure across the SDLC

Use multiple levels of abstraction to simplify development, deployment and operations

Listen to customer needs to continue to enhance the platform

# Platform Evolution

Incorporated client builder to speed prototyping and the development of internal applications

Integrated with enterprise IT systems and technologies. Now roughly 180 integrations as well as an open source toolkit

Introduced novel componentization models for increased reuse in suites of systems

Incorporated Requirements and Design tools for even higher levels of productivity and simplified onboarding

Incorporated Generative AI in a modular fashion both for use in applications and increasing developer productivity.

# Platform technical overview

# Vantiq Platform

1. Design, compose/assemble/orchestrate, operate and evolve systems that:

- Identify situations in real-time and orchestrate the responses to the situations
- Distributed
- Mission critical
- Integrate with external systems and technologies
- Collaborate across systems and users
- Exploit predictive, generative and KR AI models

2. Deploy across cloud, on premise and edge affording very high levels of productivity and agility across the SDLC:

- Design
- Develop
- Test
- Deploy & Operate

# Architecturally Unique

## Infrastructure Architecture

- Fully asynchronous execution for high efficiency, low latency and high scalability.
- GenAI and Predictive AI deeply integrated into platform.
- Declarative fault tolerance for availability & reliability
- Native support for distributed applications
- Security embedded in platform and applied to all applications hosted on the platform
- Integrations isolated from servers for security.

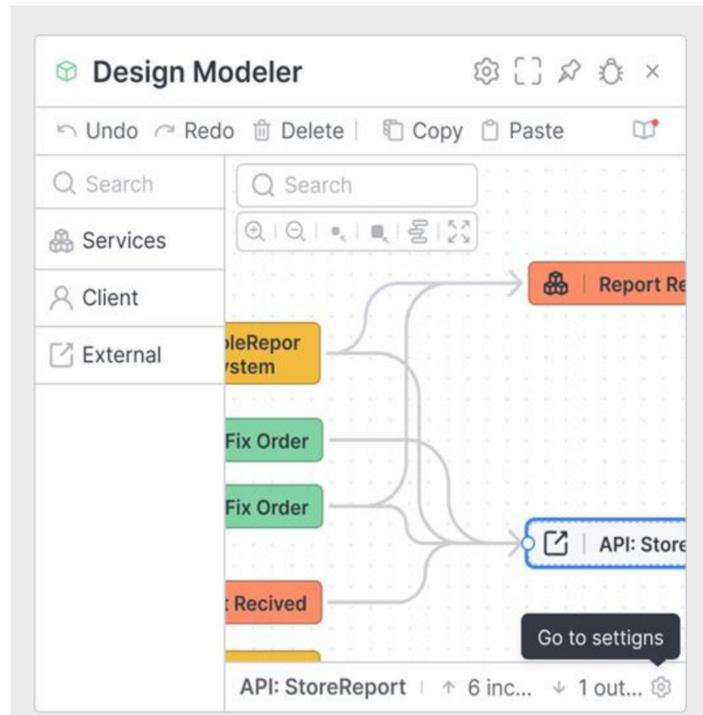
## Productivity Architecture

- AI assistance
- Multiple specification abstractions
- Components
- Assemblies
- Full SDLC support
- Full stack development - services AND clients
- Automated partitioning & deployment

# Productivity Tools

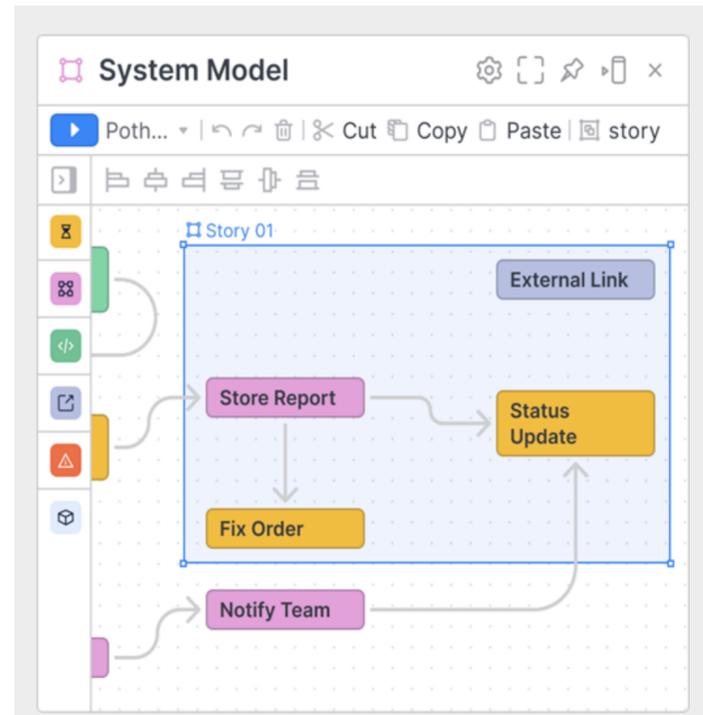
## 1. Design

EVENT MODELER



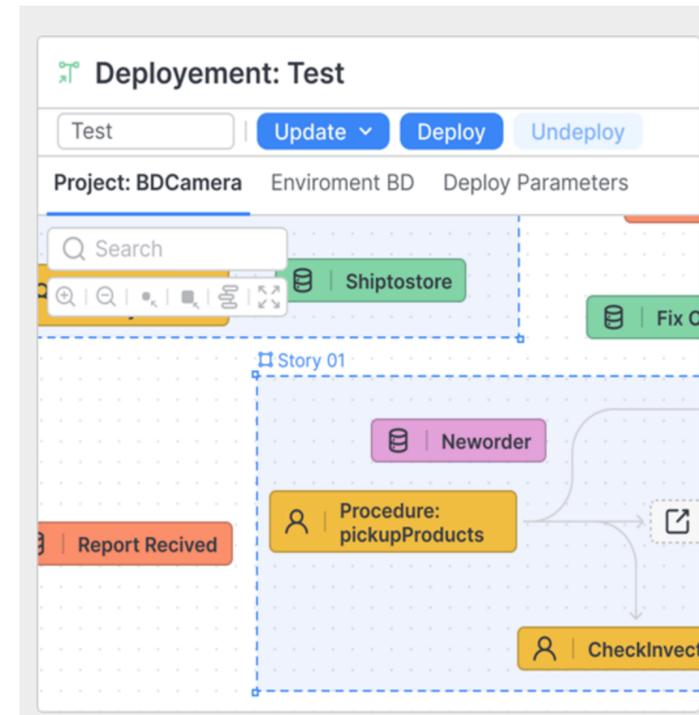
## 2. Develop & Test

VISUAL APPLICATION BUILDER



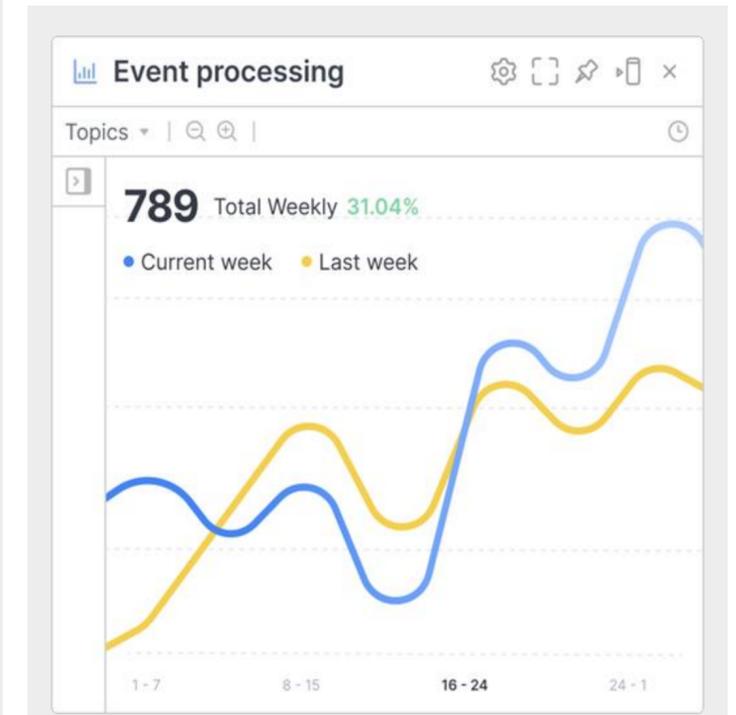
## 3. Deploy

VISUAL PARTITIONING



## 4. Operate

PERFORMANCE MANAGEMENT



# Logical server architecture

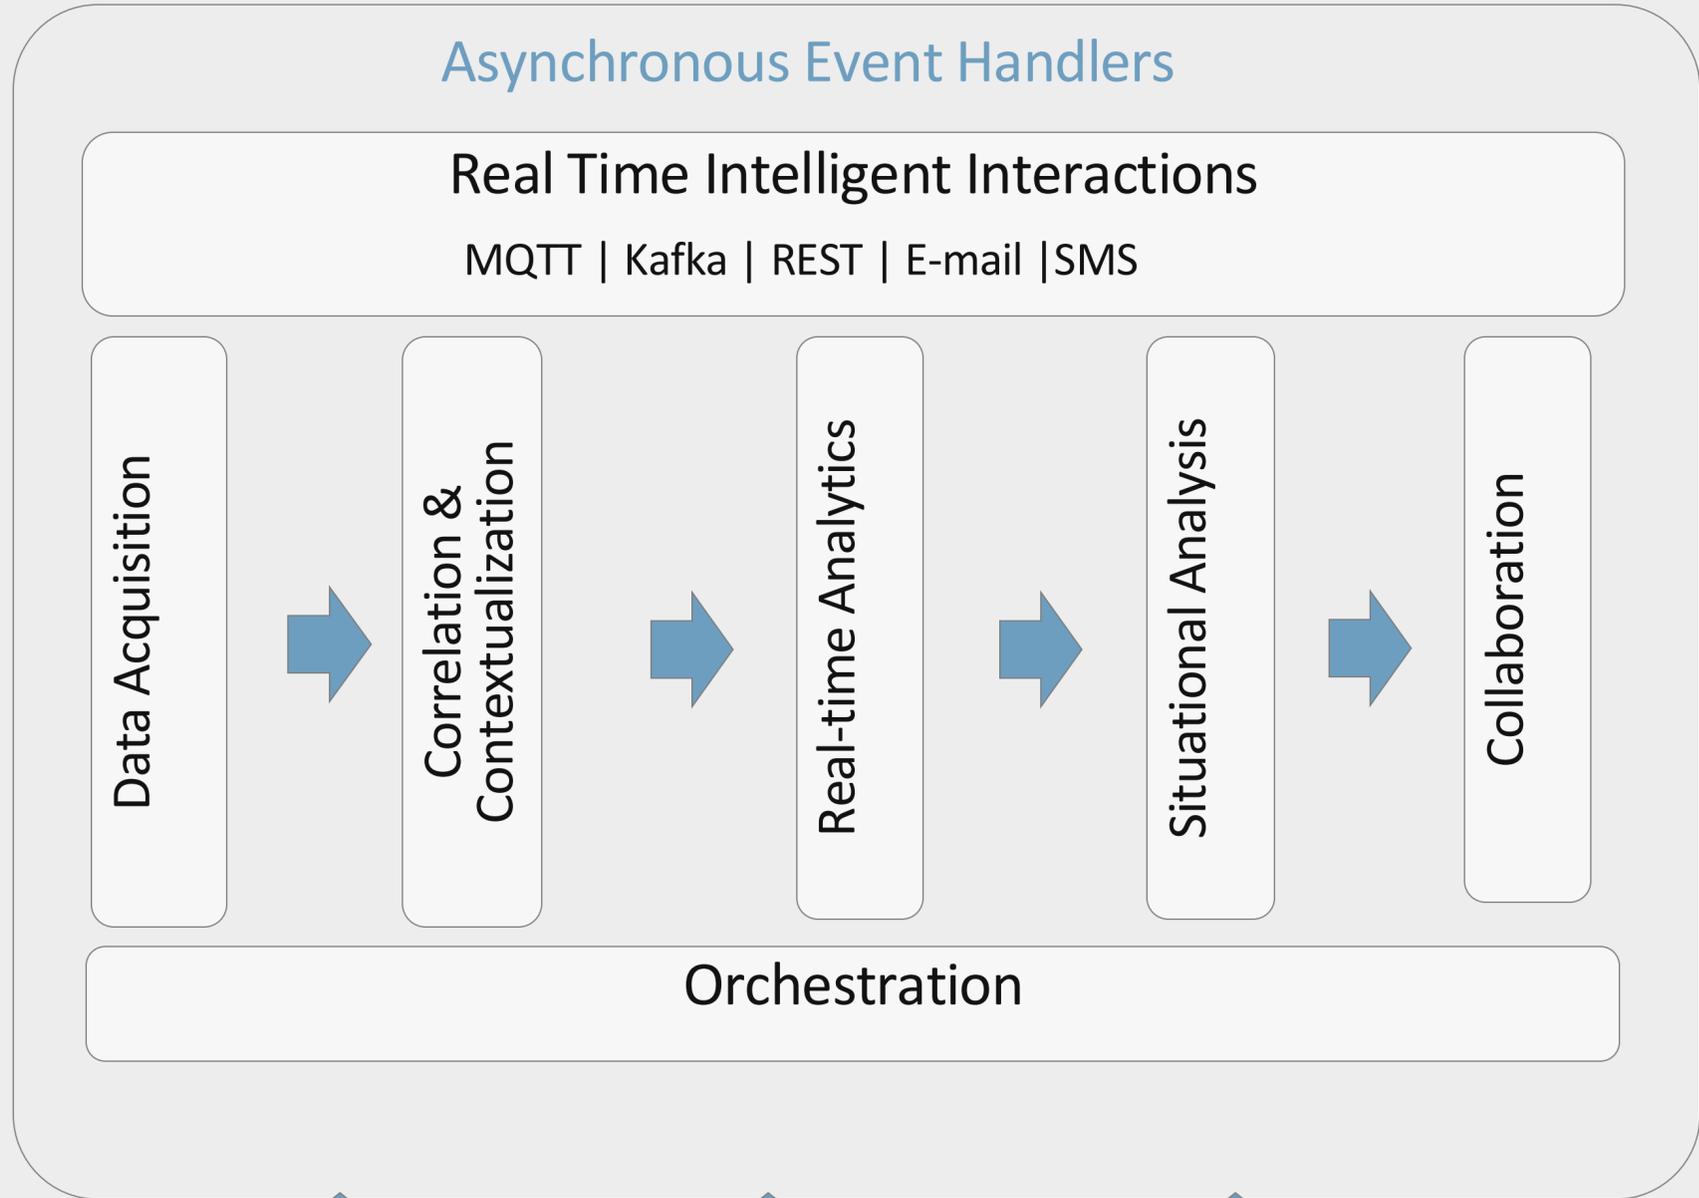
**UI Tools**  
Requirements Model  
Design Model  
Service Builder  
Client Builder  
Collaboration Builder  
System Modeler  
Test Builder  
Assembly Manager  
Partitioning  
Deployment  
Performance Monitor

→  
**Devices**  
*Beacons/IoT/...*

→  
**Streams**  
*Machine/Logs/...*

→  
**Social**  
*Twitter/Yammer/...*

Real Time Data Sources



Notification

Actions request

↑ **AI Models**  
*OpenAI/AWS/watsonx/...*

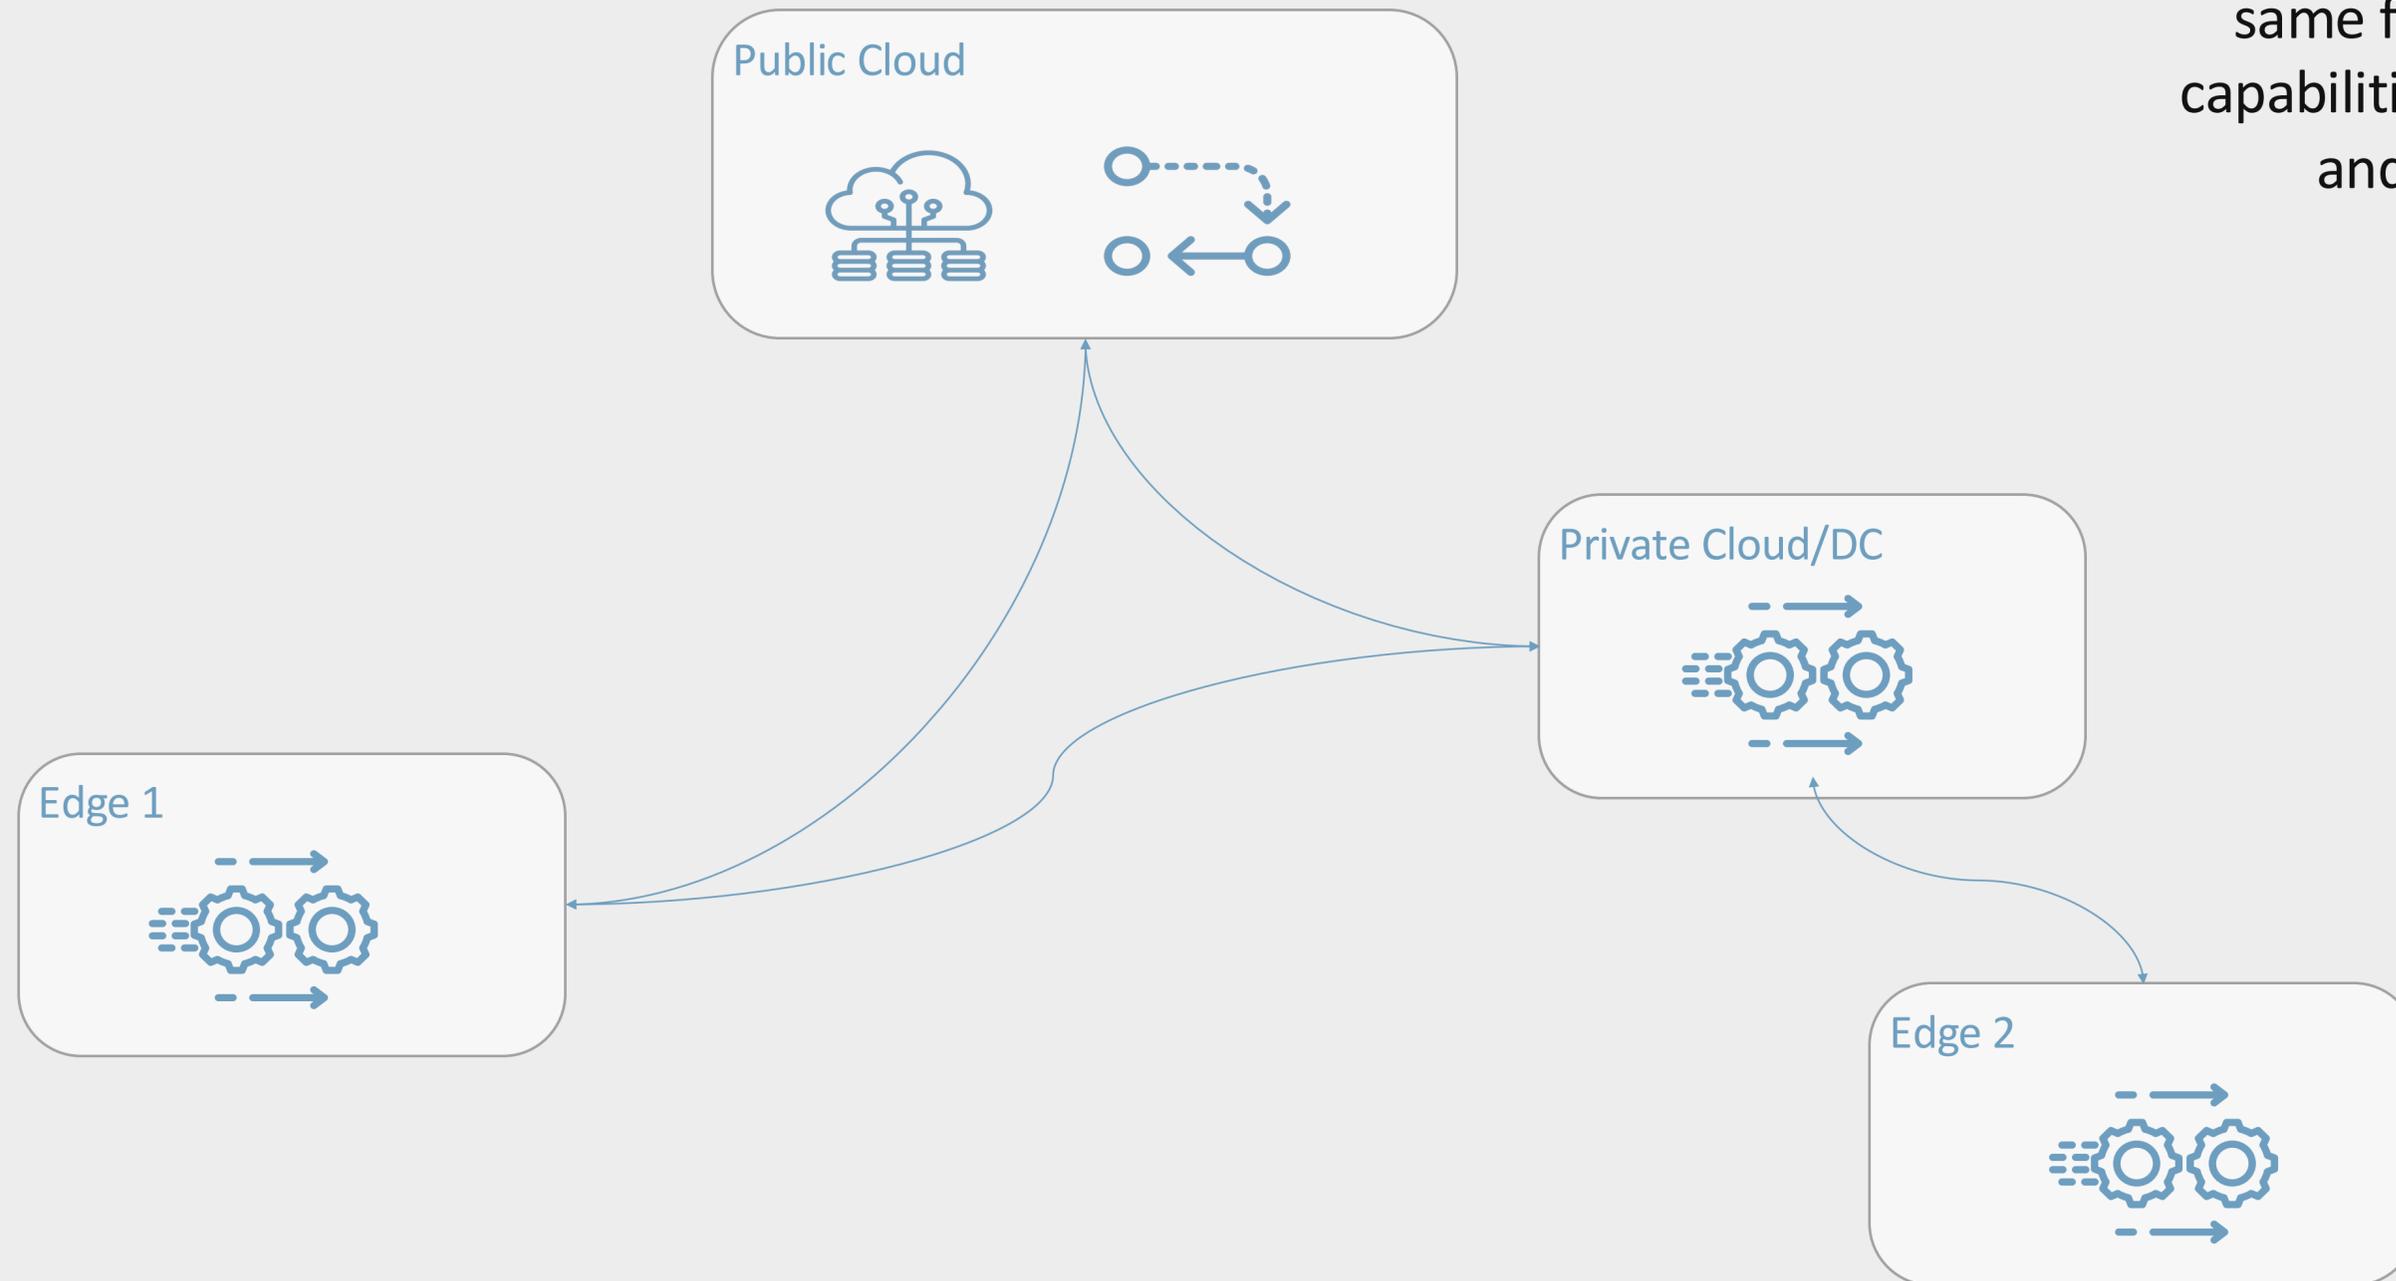
↑ **Big Data**  
*Apps/ESB/...*

↑ **Enterprise**  
*Hadoop/NoSQL/...*

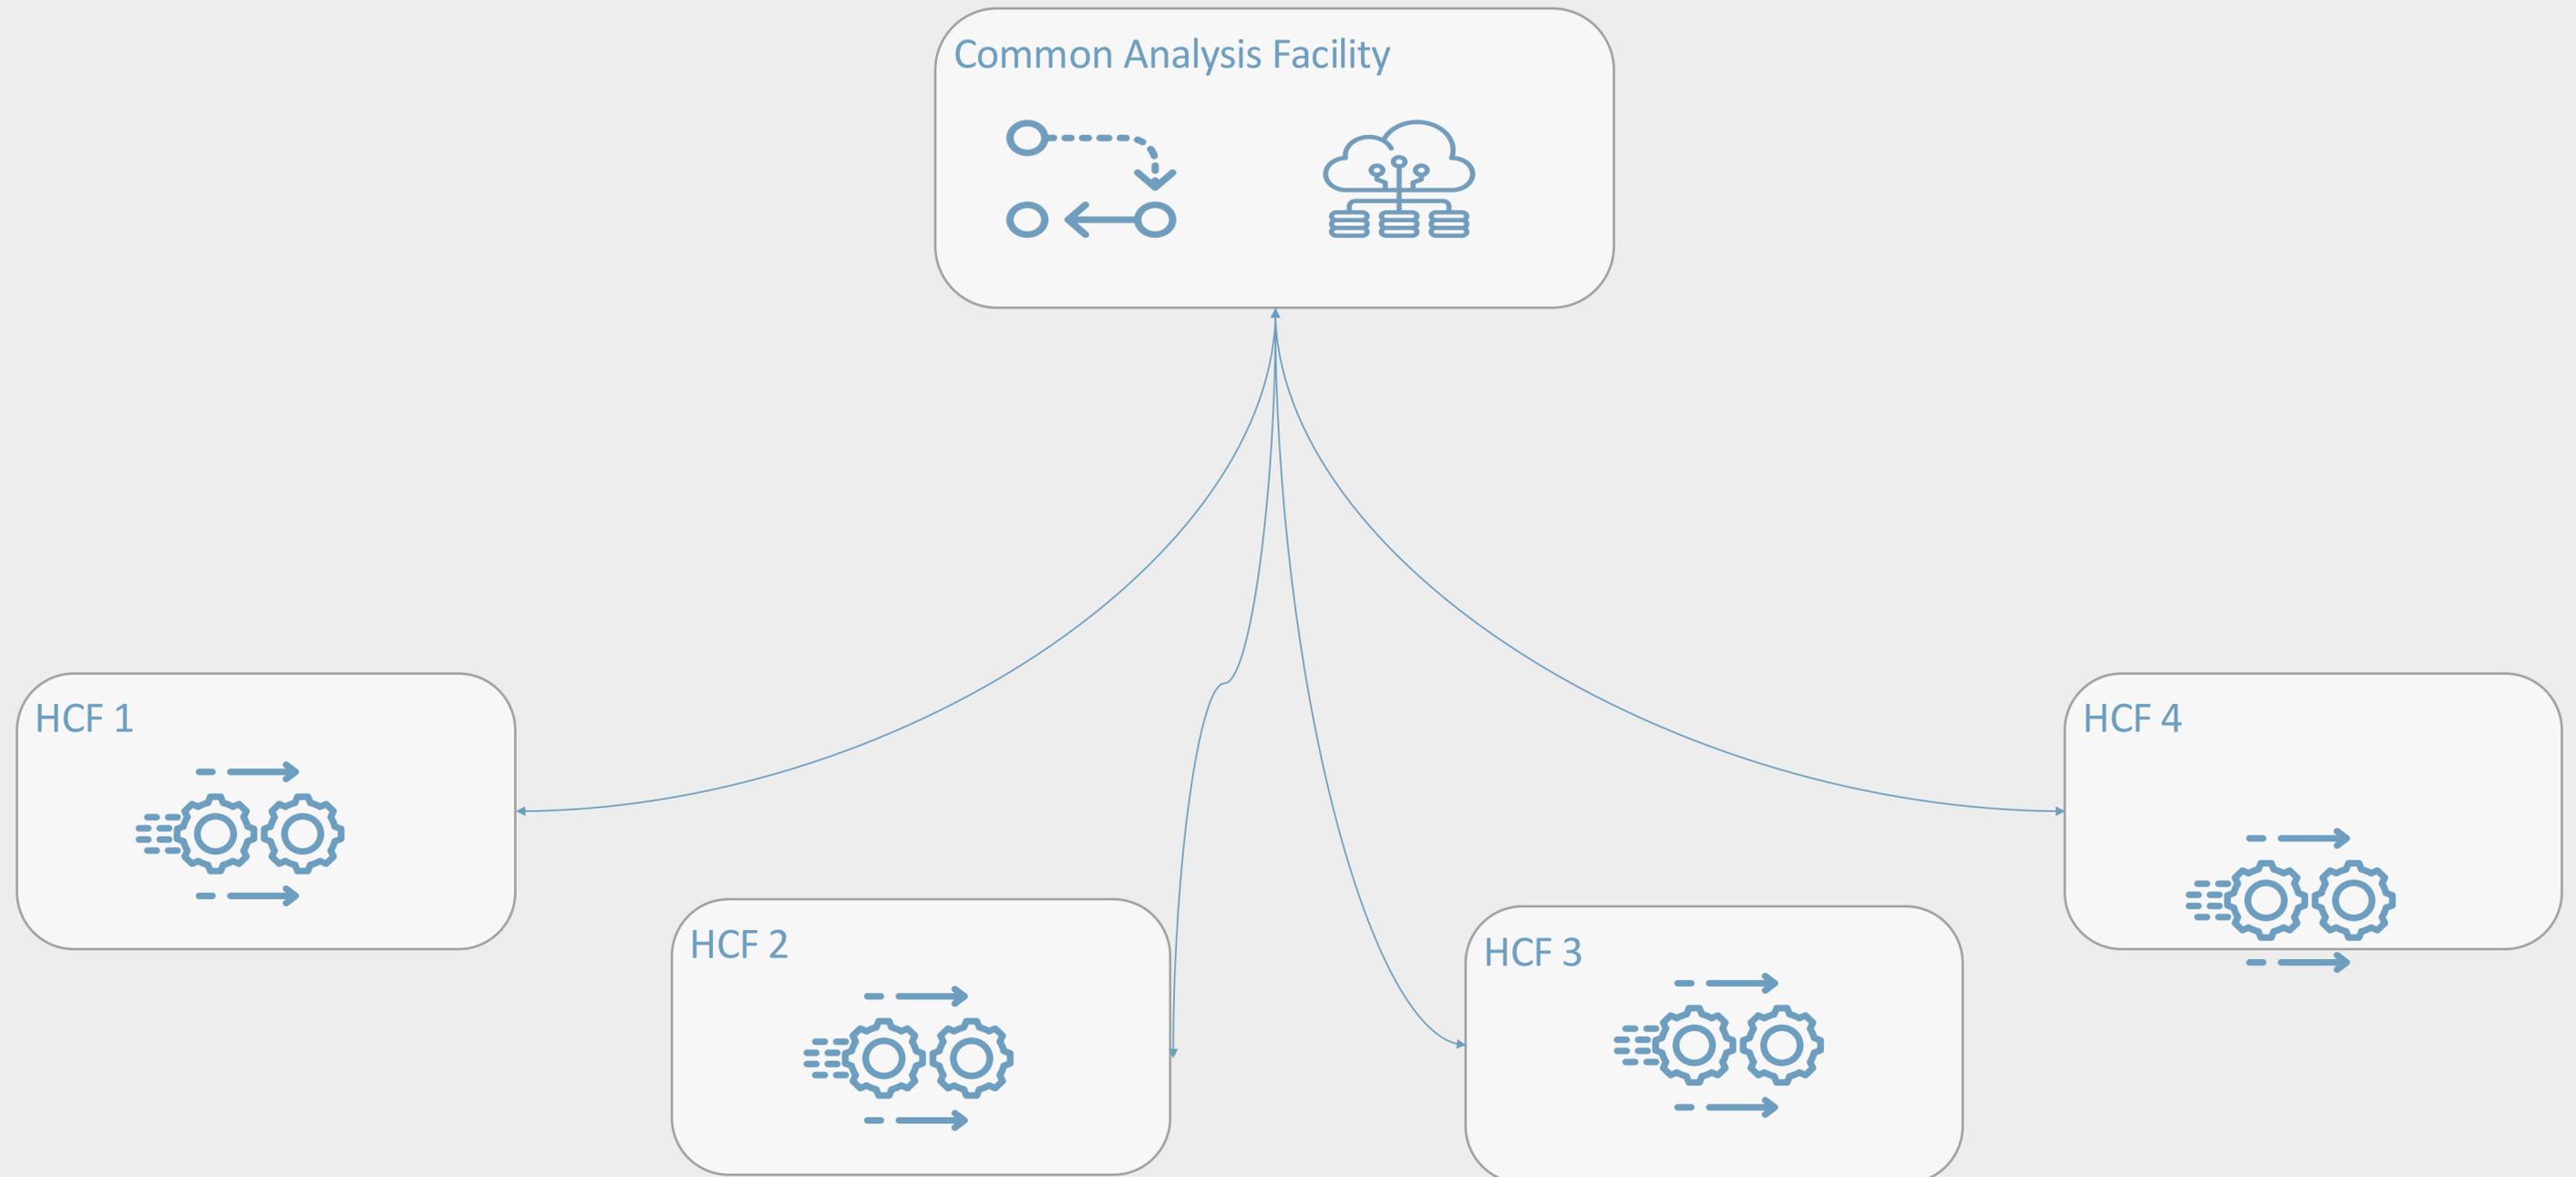
→  
**Visualization**

# ARBITRARY distributed topologies

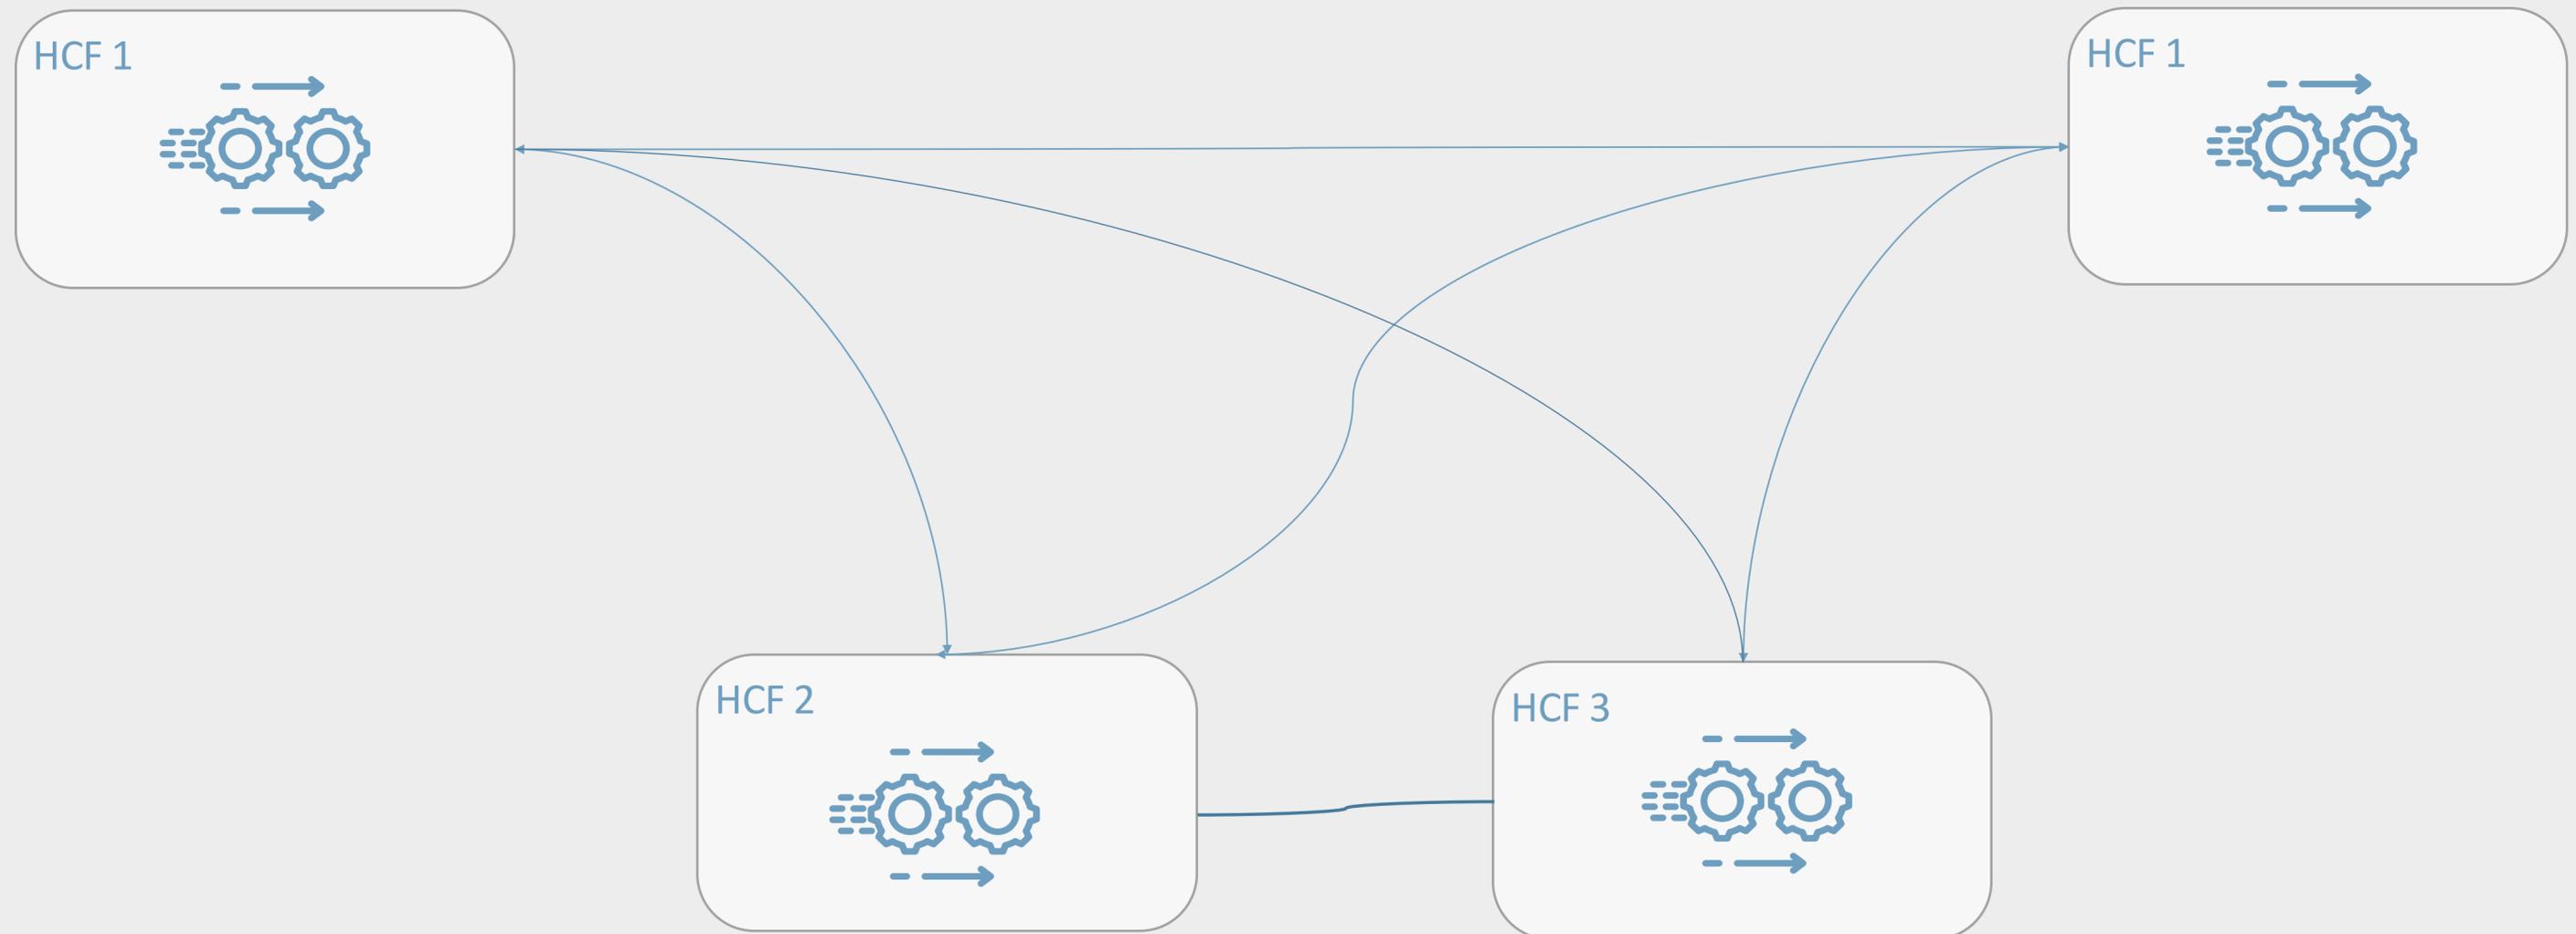
System supports same functional capabilities on cloud and edge



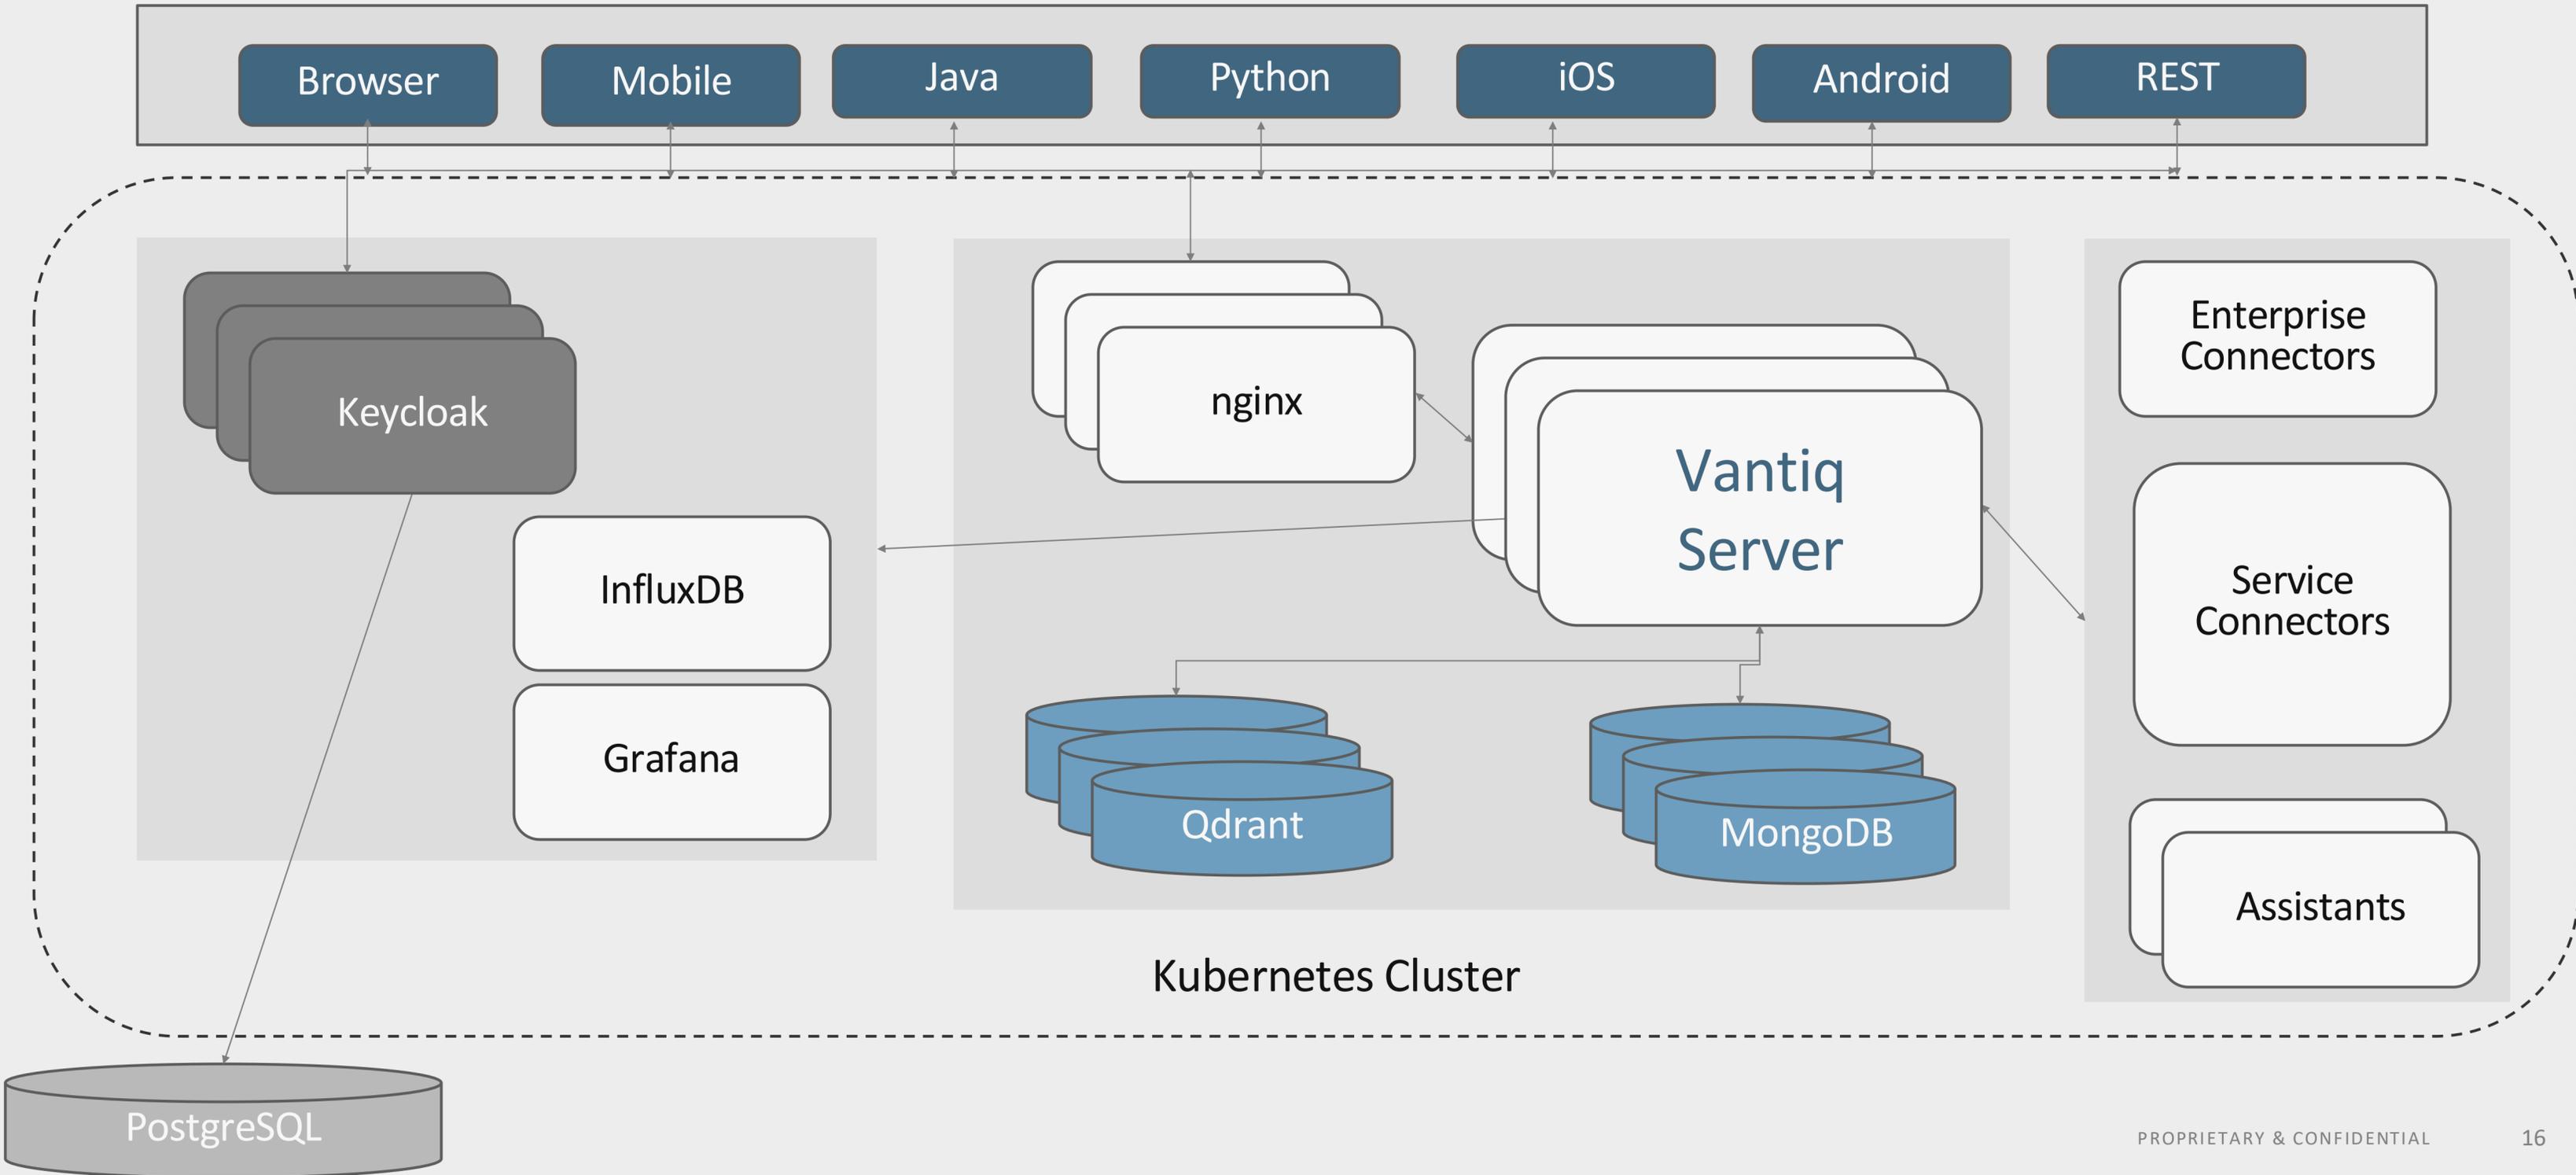
# PHI protecting distributed topology



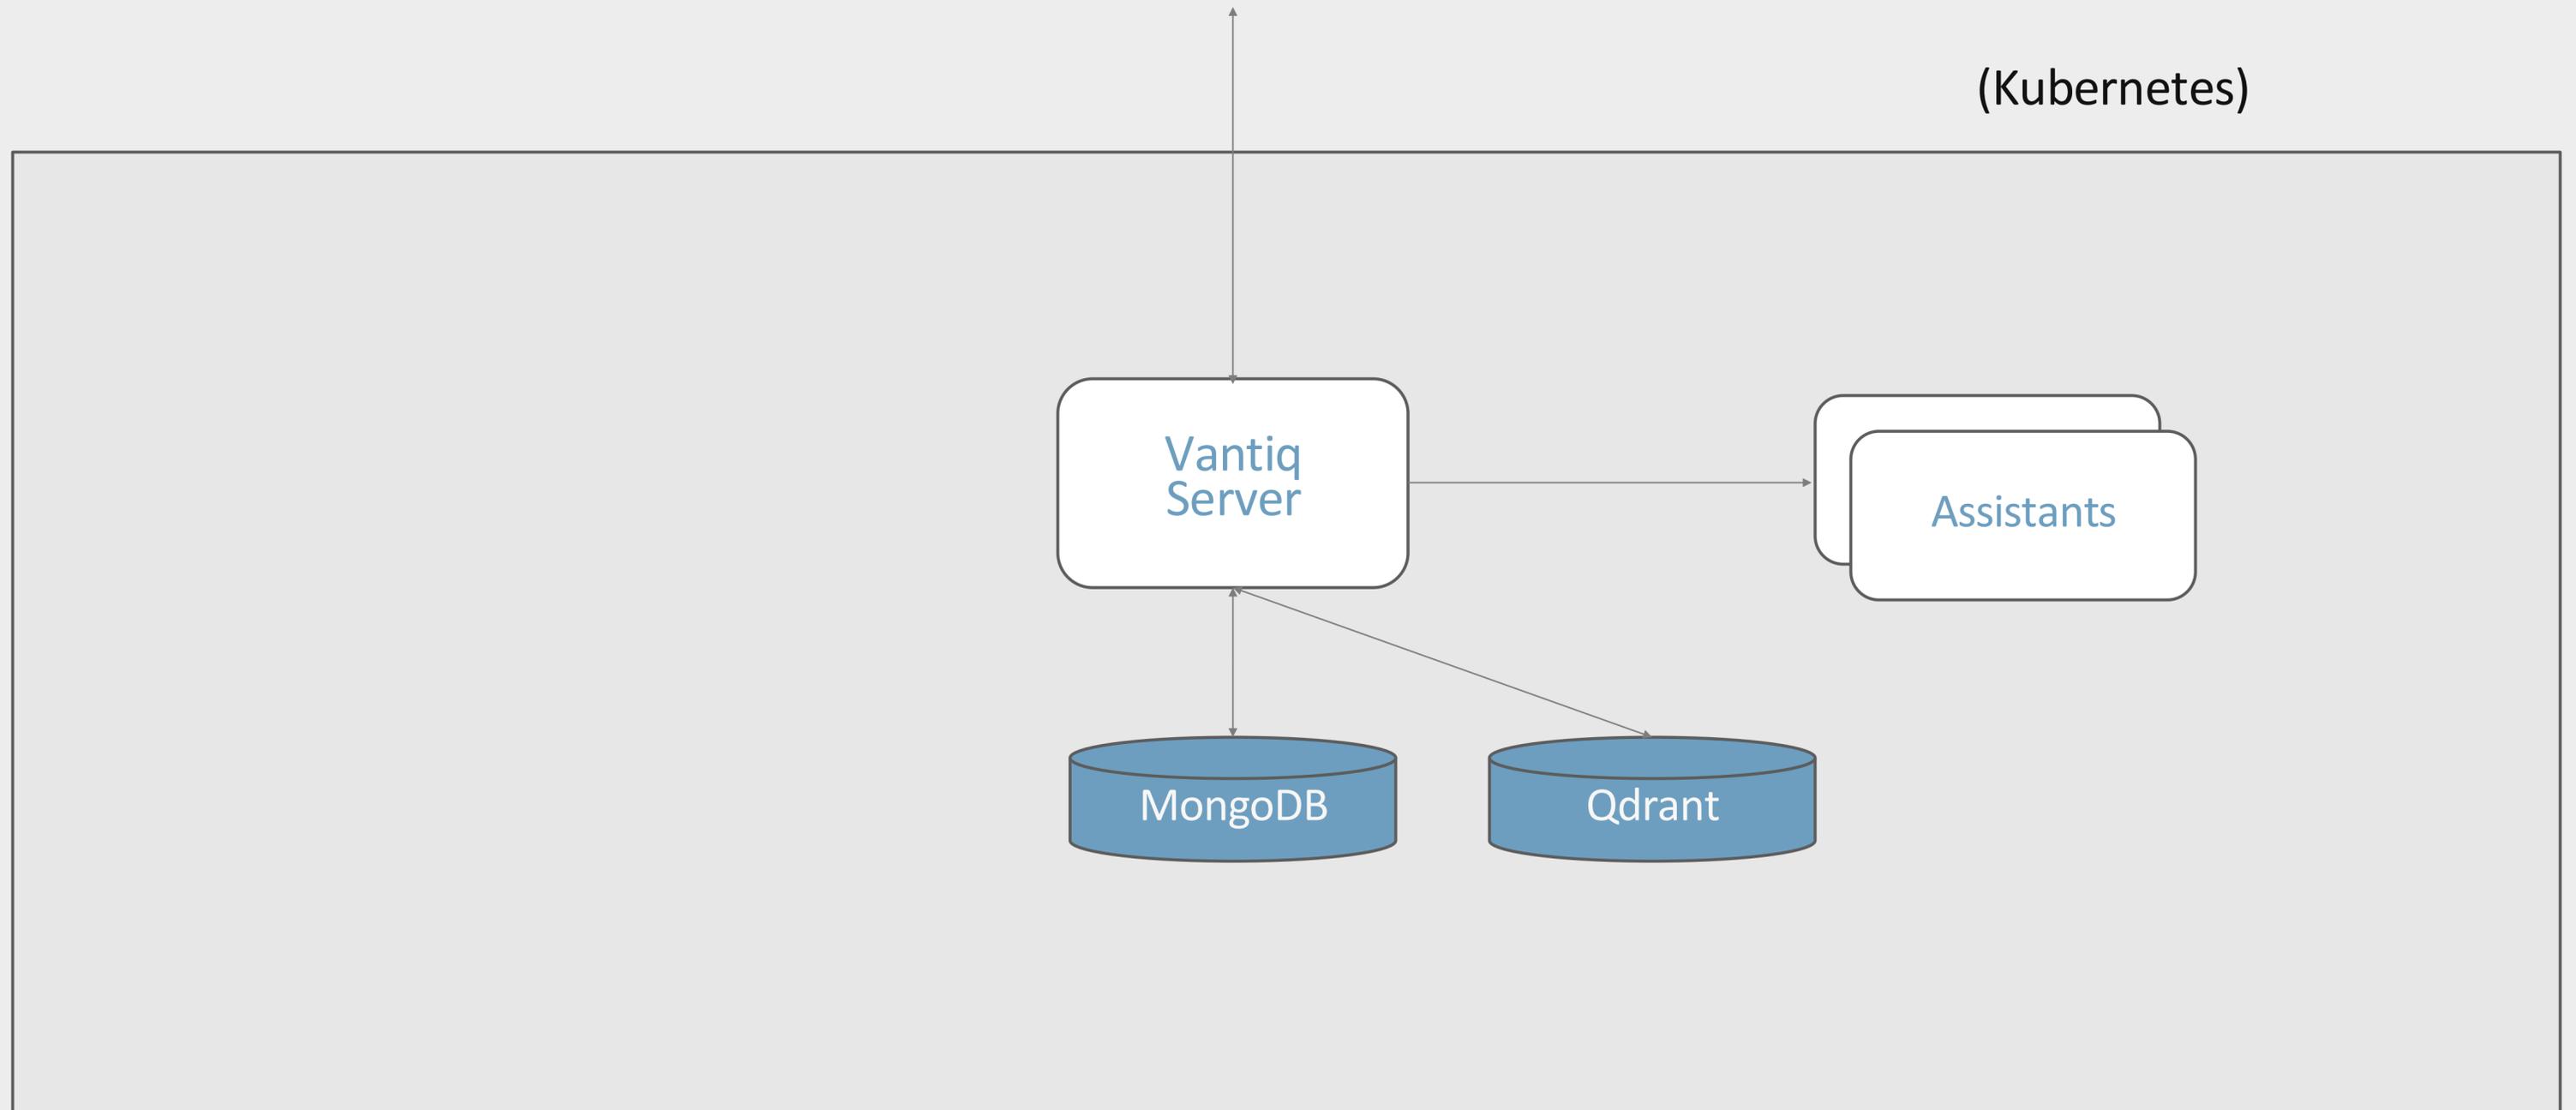
# Node recovery distributed topology



# Cluster deployment configuration



# Edge deployment configuration



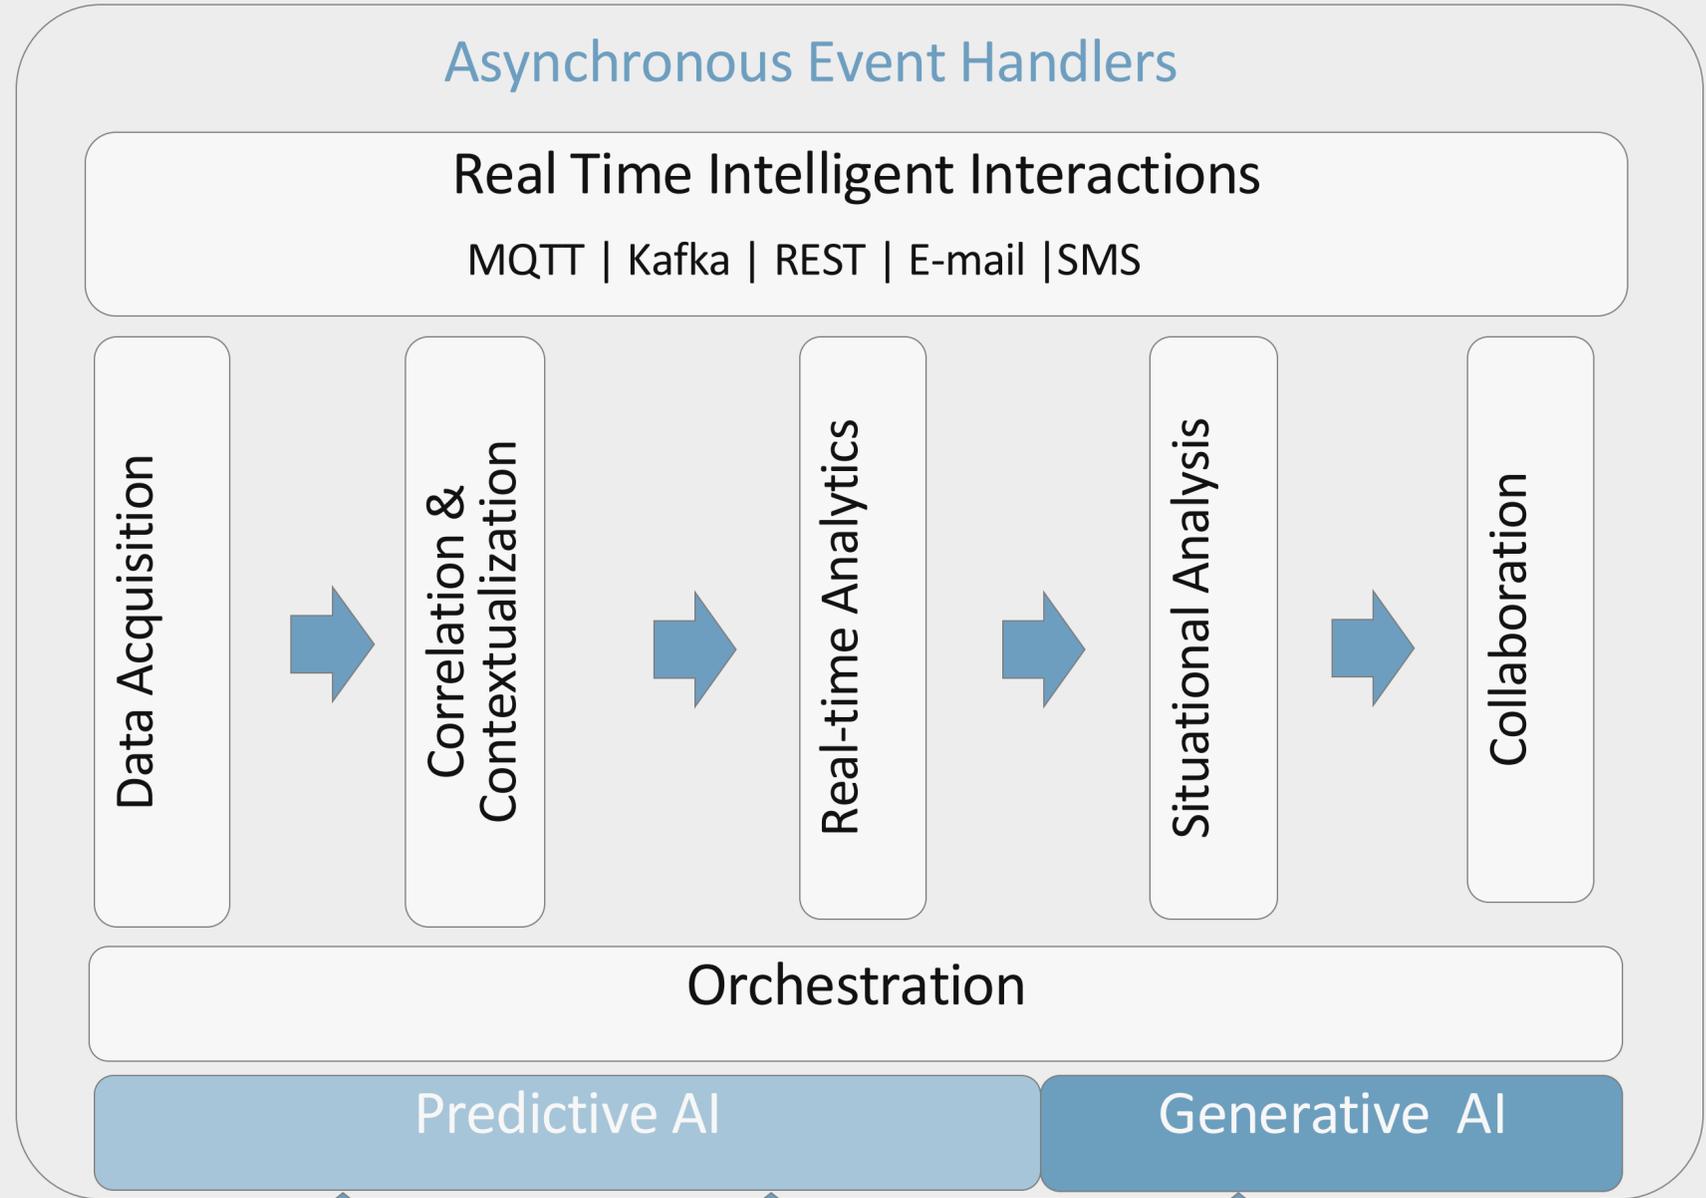
# AI Capabilities

# Applying AI to real-time systems

- UI Tools**
- Requirements Model
- Design Model
- Service Builder
- Client Builder
- Collaboration Builder
- System Modeler
- Test Builder
- Assembly Manager
- Partitioning
- Deployment
- Performance Monitor

- Devices  
*Beacons/IoT/...*
- Streams  
*Machine/Logs/...*
- Social  
*Twitter/Yammer/...*

Real Time Data Sources



Notification

Actions request

Visualization

AI Models  
*OpenAI/AWS/watsonx/...*

Big Data  
*Apps/ESB/...*

Enterprise  
*Hadoop/NoSQL/...*

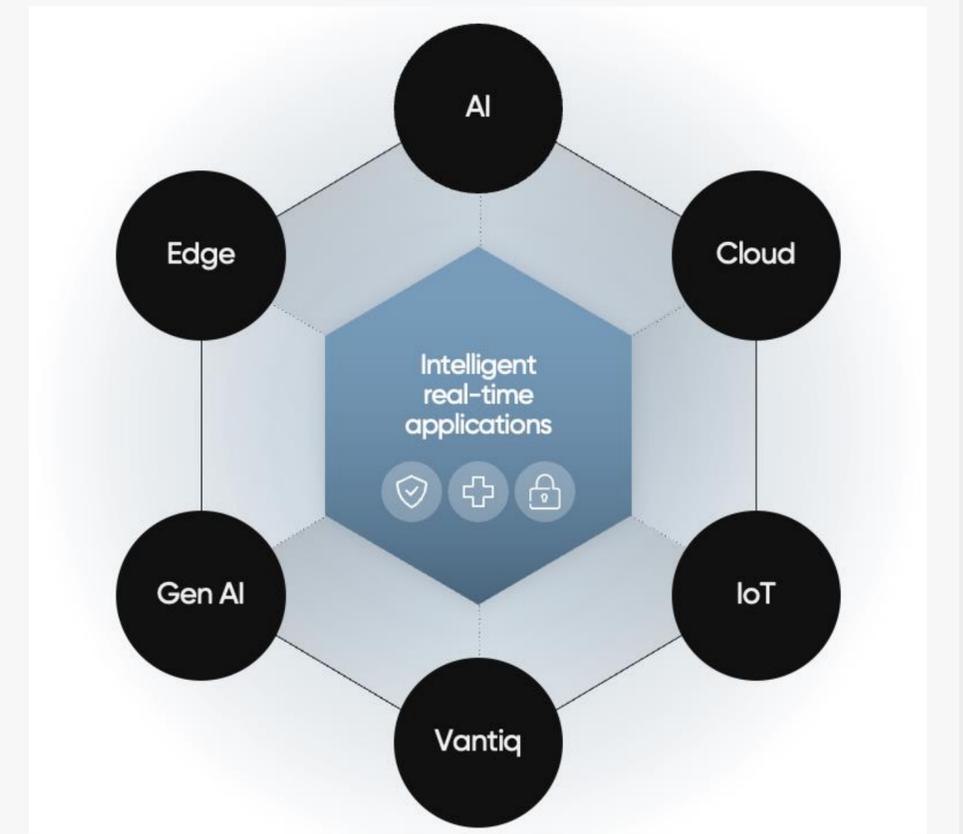
# Generative AI capabilities at Vantiq

Incorporate Generative AI into Vantiq real-time systems

- Real-time Prompt Templates
- Semantic Search
- GenAI flows
- Function Invocation, MCP
- Agents & Multi-agent Systems (Agentic AI), A2A
- Conversation & Discussion support

Improve productivity during the development of a Vantiq real-time system

- Client Builder – interactively build clients
- System Development – interactively architect systems
- Documentation Search



# Advanced Generative AI capabilities in Vantiq

1. Improve LLM response accuracy by consulting multiple LLMs

- Additional LLMs used to determine consensus response

3. Real-time augmentation of RAG & GRAPH RAG assets

4. Agentic AI and Multi-agent Systems

- Fully asynchronous for autonomy
- Distributed coordination protocols
- Persistent, local and replicated shared memory for context
- Planning and dispatching
- Componentization and discovery

2. Policy-based, real-time selection of optimal LLM to respond to a particular class of prompts such as questions vs commands

# Predictive AI integration

Integration of arbitrary predictive engines and models

- Image recognition systems
- Facial recognition systems
- Image to geo coordinates
- Object tracking
- Waveform analysis
- Audio recognition

Examples:

SAFETY AND SECURITY



HEALTHCARE—WOUND CLASSIFICATION, TELE-CRITICAL



# Rules engine

The heart of the system is an **event-driven rules Engine**



Each event that is ingested triggers one or more rules:

- Orchestrate the response(s) to the event
- Correlate events
- Encapsulate algorithmic analysis
- Incorporate predictive AI and generative AI activities into the event handling
- Trigger response collaborations

# Platform Capabilities

# Mission critical

## Reliability

- Guarantee processing of important events
- Shed unimportant events when under load via policies
- Increased LLM response accuracy via redundancy and policy-based selection

## Availability

- Fault tolerant clusters
- Automatic event handler restart from each activity pattern
- Replicated shared memory
- Automatic request redirection for external service failures
- Replicated databases

## Security

- Isolation levels – org & namespace
- Zero trust model – all requests are authenticated using Oauth2 or passwords
- Storage encryption
- Message encryption

# Distributed

Systems can be partitioned to run across any or all of:

- Cloud
- Data center
- Edge
- MEC

Partitioning allows the developer to dynamically allocate each component of the system to different distributed node sets

Greatly simplifies the development of distributed systems

- Construction
- Tuning & reconfiguration

# Productivity & Agility



1. Full life cycle support

2. Develop at four levels of abstraction

- Natural Language specifications
- Visual diagramming
- Declarative specifications
- Abstract scripting

3. AI integration

4. Automatic & dynamic deployment

5. Operational support tools

# Componentization

Key to agility - rich reusability models



VantIQ offers highly developed componentization models

- Services – all server-side logic is encapsulated as services. Services can be automatically published supporting REST and websocket interface by default
- Visual components
- Agents
- Multi-agent systems
- External integration components
- Database components
- Assemblies allow multiple resources, including services, agents, multi-agent systems to be aggregated into a single, configurable unit (Packaged Business Components)

# Orchestration

Event-driven orchestration – optimized for ingest, analysis, response

- Completely asynchronous execution model increases flexibility
- Integrated with platform decisioning capabilities
- Fault tolerant
- Orchestration as components

GenAI orchestration – optimized GenAI activities

- Sequence activities related to GenAI activities
- Access to full set of platform capabilities and GenAI resources

Collaborations – optimized M2M and User activities

- Flexible asynchronous execution model

# Real-time



Process events as they occur independent of arrival rate:

- 1 event every 15 minutes
- 100,000's events every second



Determine the situation in real-time



Respond to the identified situation in real-time



Response time from as low as a few milliseconds to minutes/hours



Real-time event streams can be unpredictable workloads

# Performance

Measure response time in milliseconds

Measure throughput in 100,000s events/second

Use memory-based algorithms for maximum throughput and minimum response times

Benchmarked on single machine at:

- 200,000 events/sec
- 8 msec average response time



# Current and Future Directions

# Event-driven, Mission Critical Systems

Always improving enterprise readiness: reliability, availability, performance, security.

Increasing fault tolerance across more topologies and cases.

# VIA – Creating Intelligent Systems

- Agentic assistant for rapidly creating intelligent systems
- Compresses the Intelligent System Life Cycle: ideation to operations.
- Conversational user experience – text and voice
- Collaborative
  - Users directs, approves and adds detail where needed
  - VIA does the low-level work
- Proactive: VIA offers suggestions across the systems life cycle
- Learns organizational knowledge and personal preferences
- Enterprise requirements satisfied by default:
  - Reliability, Availability, Scalability, Security
  - Documentation
  - Regression testing

# Adaptive Automation

Automation works for relatively simple business processes

Does not work for processes that have large numbers of variations or large numbers of distinct errors that must be recovered

Example: hospital processes such as admittance and discharge have thousands of variants

Adaptive automation customizes a basic process instance by instance by exploiting GenAI to adjust the process using available context and learning systems.

Each process variant is logged for approval by the governing authority and for use in increasing the learning capabilities of the system.

# Federated Natural Language Query

GenAI models can be reasonably adept at issuing natural language queries against a single database assuming the database exhibits reasonably simple semantics.

Many systems require queries to be run against a set of federated databases each with its own semantics at varying degrees of complexity.

FNLQ consists of a collection of agents that specialize in NL query, analytics, ontologies, schemas and the specific semantics of each participating database.

The agents work together to determine the query semantics, the target databases that contains the desired data, the queries that must be applied to each database and the queries that combine the individual results.

The system is driven by domain-specific context and a knowledge of the schemas and ontologies associated with each database and the federated system as a whole.

# Next Generation User Interfaces

Use natural input models – eliminate most typing

voice

vision

Capture context and input from environment including IoT devices

Eliminate navigation via gestures – pointing, clicking, swiping by understanding context

More flexible information display

Exploit A/R devices – voice, vision

# Thank you



Paul  
Butterworth CTO

---

CTO AND FOUNDER | FORTE

---

CTO | AMBERPOINT

---

CTO | EMOTIVE

---

CHIEF ARCHITECT | INGRES

---

VANTI 

Thank you