



Horizons  
Innovation Summit

# Horizons Innovation Summit 2026

All demos

March 2026

© 2026 NTT DATA, Inc.

# Gamification for HR

## An interactive recruitment and onboarding experience

### Core Idea

We developed a videogame to enable HR teams to deliver an interactive and memorable onboarding experience by leveraging LLMs to provide interactive voices for the videogame characters.

In the game, candidates navigate a virtual Japanese town, collecting clues about the company's culture and values by answering questions. These questions assess their knowledge of the company and their ability to apply it to real-world scenarios. Some of the characters in the game are powered by generative AI, allowing them to generate new dialogue and behaviors based on player interactions, creating a dynamic, realistic and engaging experience.

### Challenge

The HR department sought an innovative recruitment method, as traditional approaches were less effective in engaging new hires and ensuring retention of critical company information. They needed a solution to attract top talent and provide an interactive, memorable onboarding experience, fostering a strong connection between new employees and the company from the start.

### Impact and Applications

→ Increased engagement: Improved information retention, and strengthened new hires' connection to the company's culture.



# Network in a box: your 5G bubble anywhere

Deploy secure, low-latency private 5G anywhere in 30 minutes

## Core Idea

We packaged a fully operational standalone private 5G network into a single portable unit to enable enterprises and engineering teams to instantly deploy secure, low-latency connectivity wherever needed, by leveraging an integrated 4G/5G core, edge compute, and built-in networking capabilities.

Rather than treating private 5G as a complex infrastructure program, we transformed it into a production-ready, plug-and-play business asset. The solution integrates Celona Edge (4G/5G converged core), 5G RAN (AP22), edge compute for microservices, and integrated core switching, routing, and firewall capabilities, requiring only power and Internet access to operate.

MicroSlicing ensures deterministic QoS per device and application, while IOWN transport enables geographic separation of Core and RAN with precise PTP timing, preserving performance at scale.

## Challenge

Private 5G deployments are often slow, integration-heavy, and treated as large infrastructure programs.

Organizations need a faster, lower-risk way to deploy secure mobile edge connectivity as business demand for real-time, low-latency use cases accelerates. The traditional model limits agility and delays value realization.

## Impact and Applications

- Reduced deployment time: From months to under 30 minutes, lowering integration cost and accelerating time-to-value.
- Portable architecture: That enables rapid field activation for manufacturing, logistics, defense, and remote operations. Zero-touch provisioning simplifies scaling.
- What's next: Distribute Core and RAN over the IOWN optical network while preserving precise PTP timing.



# Federated edge monetization for automotive OEMs

## Unified API-driven 5G edge platform for connected vehicle services

### Core Idea

We designed and validated a federated multi-operator edge architecture to enable automotive OEMs to deploy low-latency Digital Twin and connected vehicle services across, by leveraging a standardized API-driven federation layer and real-time usage monetization.

Instead of integrating separately with each operator, OEMs connect once through a unified northbound interface, while a Federation Broker model orchestrates network slicing, edge workload placement, and roaming continuity across CSP environments.

Built on GSMA Open Gateway principles and CAMARA®-based APIs, the platform standardizes lifecycle management converts edge consumption into usage-based billing. A Teltonika® 5G dual-mode modem router was integrated within the Catalyst solution to validate end-to-end operability.


### Challenge

Automotive OEMs require deterministic, low-latency edge services across multiple operators to support connected vehicle use cases.

However, CSP environments remain fragmented, with proprietary APIs and inconsistent slicing models. This fragmentation slows global deployment, increases integration cost, and limits scalable edge monetization.

### Impact and Applications

- Federated orchestration: Enabled federated edge orchestration with real-time, usage-based billing across operators.
- Performance metrics: 10–15% lower integration costs, ~20% faster rollout, and 3–5% ARPU uplift.
- Use case support: Supports SLA-based connected vehicle services, Digital Twins, and geo-fenced mobility use cases.



# A4S Sandbox

## Testing environment for trustworthy AI validation and EU AI Act readiness

### Core Idea

The A4S Sandbox is a structured testing environment that helps organizations evaluate and strengthen the technical trustworthiness of their AI models. It focuses on four key dimensions: Robustness, Fairness, Cybersecurity, and Explainability.

Unlike high-level risk management tools, A4S provides in-depth, hands-on testing by simulating real-world conditions to identify weaknesses, ensure compliance with AI Act requirements, and support conformity assessments. It also enables continuous post-deployment monitoring and offers practical recommendations and countermeasures to improve AI systems.

Whenever an AI model is updated through retraining, fine-tuning, or architectural change, the A4S Sandbox supports validation and compliance. Overall, it helps European organizations deploy trustworthy AI with confidence, reduce compliance risks, protect customer trust, and maintain the integrity of AI-driven operations.


### Challenge

Interpreting and complying with the EU AI Act is challenging, especially for high-risk AI applications. Organizations need to ensure their AI systems are robust, fair, secure, and explainable, all while advancing AI development.

Additionally, there is a demand to accelerate time-to-market without compromising trustworthiness.

### Impact and Applications

- **Accountability:** A4S strengthens accountability and transparency through evidence-based compliance with the EU AI Act.
- **Ethical Innovation:** Helps detect and reduce bias to support equitable outcomes, and enables ethical innovation within a regulation-aligned environment - advancing Europe's leadership in responsible AI.



# Zero Trust Data protection powered by ABE

Protect sensitive data at its core making stolen data unusable by design

## Core Idea

Traditional Zero Trust Architecture (ZTA) relies on Role-Based Access Control (RBAC), but modern threats require more granular and resilient protection. By augmenting Zero Trust with Attribute-Based Encryption (ABE), we introduce a data-centric defense layer that protects information both inside and outside the traditional perimeter.

Unlike conventional models, ABE embeds access policies directly into the data, governed by dynamic attributes rather than static roles. We developed the industry's first Zero Trust Data Security (ZTDS) platform powered by ABE to protect modern organizational data estates.

Using patented fine-grained, data-centric encryption, the solution secures sensitive information regardless of system, format, or state; whether at rest or in transit, turning stolen data from an asset into a liability for attackers.

## Challenge

The escalating sophistication of cyberattacks has rendered perimeter-based defenses insufficient. Traditional security models, designed to protect network boundaries, provide limited resilience once a breach occurs.

In an “assume breach” world, defending the perimeter is no longer enough if the data inside remains exposed and accessible to anyone who bypasses the outer layer.

## Impact and Applications

- Data protection: Even if attackers gain access to critical data - financial records, employee information, or proprietary designs - the data remains unusable without the required attributes and policies.
- Security measures: The platform mitigates AI-driven automated attacks, prevents unauthorized agentic data exposure, and is designed to withstand emerging Post-Quantum Cryptography (PQC) threats, ensuring long-term data resilience.

# AGADE: Human-Tailored Robotics

Reducing injury risk while improving industrial productivity

## Core Idea

AGADE developed a patented semi-active robotic technology that enables smart occupational exoskeletons to provide adaptive, real-time muscular support while minimizing power consumption and maximizing user comfort.

The innovation combines an elastic mechanical system with electrical actuation to create a semi-active architecture that delivers assistance only when needed. Unlike fully active systems that require continuous power, or passive systems that provide static support, AGADE's approach integrates robotics and AI to recognize manual tasks and dynamically adjust the level of assistance.

By leveraging intelligent task recognition and energy-efficient semi-active actuation, the system ensures targeted muscular relief, reduces fatigue, and improves ergonomics while maintaining lightweight design and extended operational autonomy.

## Challenge

Industrial workers face repetitive movements, heavy lifting, and prolonged static postures that increase fatigue and musculoskeletal risk.

Traditional exoskeletons struggle to balance effectiveness, comfort, energy efficiency, and adaptability. Delivering lightweight, real-time, task-aware assistance remains a key technological challenge.

## Impact and Applications

- Adaptive support: AGADE's smart exoskeletons provide adaptive muscular support that reduces physical strain and improves worker comfort and productivity.
- Workforce safety: By lowering fatigue and injury risk while maintaining energy efficiency, the solution enables safer operations and long-term workforce sustainability.



HUMAN-TAILORED  
ROBOTIC



# Agent-driven Contact Center

## Enhancing customer care with Agentic AI capabilities

### Core Idea

We developed a cloud-native, single-instance omnichannel contact center platform that embeds Agentic AI across routing, customer interaction, and post-contact processes. By leveraging AI-powered routing, AI agents, and live APIs, the platform transforms the contact center from a traditional service hub into an intelligent orchestration layer.

The solution combines automation, personalization, and real-time sentiment analysis to improve decision-making and service quality. Its model-agnostic architecture ensures flexibility, observability, and long-term sustainability of AI capabilities while enabling scalable operations across channels.

### Challenge

Modernizing legacy on-premise contact center systems while managing high interaction volumes across voice and digital channels remains complex. It requires integrating heterogeneous enterprise systems, enabling multi-outsourcer decoupling, embedding advanced AI capabilities, and ensuring real-time data synchronization, all without introducing vendor lock-in or operational disruption.

### Impact and Applications

- Improves customer satisfaction: Delivers consistent, context-aware interactions while reducing handling times and increasing automation.
- Resource optimization: AI-driven routing and real-time assistance improve first-call resolution and efficiency.
- Scalable architecture: Modular cloud design enables innovation, personalization, and long-term value generation.



# Agentic payments: the economic layer of AI

## Enabling AI agents to transact autonomously with the x402 protocol

### Core Idea

We created a set of services to enable AI agents to execute secure payments by leveraging the A2A-x402 protocol and programmable stablecoin settlement.

We established the foundation for the emerging Agent Economy:

- Reduced friction from multi-step manual shopping to autonomous execution
- Enabled secure delegation of spending authority via DIDVC
- Introduced machine-to-machine revenue streams via payable APIs
- Improved decision quality by enabling agents to purchase third-party scoring services
- Enabled asynchronous autonomous execution (e.g., products purchased when available)

A2A-x402 protocol enables interoperable agent-to-agent payments while Demetra provides APIs for a merchant scoring service.

### Challenge

As AI agents become more capable, they increasingly need to access external services and complete economic actions autonomously. However, today's digital ecosystems are fragmented: payments are user-driven, APIs are not economically interoperable and monetization of services between agents is complex.

The opportunity was to create a standardized economic interaction layer that allows agents to transact securely, autonomously, and at scale.

### Impact and Applications

- New revenues: Clients can acquire new revenue streams from the growing volumes of machine-to-machine payments.
- Process efficiency: Increased performances by unlocking the potential of agents that can perform complex tasks by purchasing the resources they need.
- Scalability: The application can scale from consumer experiences to autonomous B2B supply chains and financial services.



# AGV fleets optimization with NVIDIA

## Real-time AGV fleet optimization enabled by NVIDIA Isaac Sim

### Core Idea

We continuously optimize and dispatch AGV missions in real time to help warehouse/factory operators increase throughput and reliability, by leveraging a digital twin that stays synchronized with the shopfloor.

The system ingests live LIDAR/camera signals to detect obstacles, blocked aisles, and route availability; it then recomputes feasible paths, assigns missions automatically, and updates execution plans without manual scheduling. Operators get real-time route and mission visibility through a lightweight KPI dashboard.

By treating routing and dispatch as a closed-loop control problem: instead of static planning, the fleet adapts continuously to layout changes, congestion, and interruptions, preventing deadlocks and reducing recovery actions.

NVIDIA Isaac Sim for scalable digital-twin simulation/validation + LIDAR/cameras for real-time environment awareness feeding the optimizer.

### Challenge

Industrial environments change continuously: aisles can be blocked, obstacles appear, and priorities shift during the day.

AGV fleets must coordinate safely while avoiding congestion and deadlocks, keeping missions on time despite route disruptions. When visibility is limited and re-scheduling is manual, small changes quickly translate into delays and unstable plans.

### Impact and Applications

- Faster operations: Fewer delays from dynamic re-routing and automated dispatch.
- Higher robustness: Immediate reaction to obstacles and route disruptions.
- Scalable foundation: Same architecture extends to larger fleets/sites and other AMRs.
- Leaner footprint: Fewer unnecessary movements and lower energy usage.



# 3D scene reconstruction

## From 2D video to 3D scene with Gaussian Splatting

### Core Idea

We developed a real-time 3D scene reconstruction pipeline that enables developers and content creators to generate photorealistic, interactive 3D environments from multi-view images.

The key innovation lies in representing the scene not through neural networks or dense geometric structures, but as a set of optimized anisotropic 3D Gaussians. This approach replaces heavy neural radiance field training with a direct, continuous scene representation that encodes position, color, opacity, and spatial extent within each gaussian primitive.

By combining anisotropic Gaussian primitives with GPU-accelerated rasterization, the system achieves high visual fidelity and real-time rendering performance, while significantly reducing computational overhead. The result is a scalable and efficient solution, even for complex scenes.

### Challenge

Traditional 3D reconstruction methods, including photogrammetry, are computationally intensive, slow to process, and dependent on heavy neural network training. They often struggle with noise, transparency, and reflective surfaces.

These limitations constrain scalability for immersive, high-fidelity 3D content production.

### Impact and Applications

- Efficiency: Reduces computational overhead and minimizes reliance on intensive neural network training, improving scalability and cost efficiency.
- Real-Time rendering: Enables real-time rendering with near-photorealistic quality, even in complex scenes involving transparency and reflective materials.
- Time-to-create: Accelerates 3D content creation for media, Digital Twins, simulation, and immersive enterprise applications.



# AI Communications Gateway

Your gateway between human beings and Agentic AI ecosystems

## Core Idea

We built an enterprise-grade AI Communications Gateway to enable global enterprises to deliver secure, human-like, multichannel AI interactions at scale, by leveraging real-time language models and knowledge-augmented orchestration within a governed, enterprise-controlled framework.

The gateway unifies voice, chat, email and digital channels into a single orchestration layer, ensuring consistent persona, guardrails, compliance and seamless human-in-the-loop escalation. It harmonizes agentic AI ecosystems, backend systems and human agents, while maintaining operational visibility, cost control and resilience.

The real challenge is not smarter AI agents; it is governing and orchestrating AI communications at scale. A centralized gateway ensures consistency, compliance, brand control and measurable performance across all interactions.

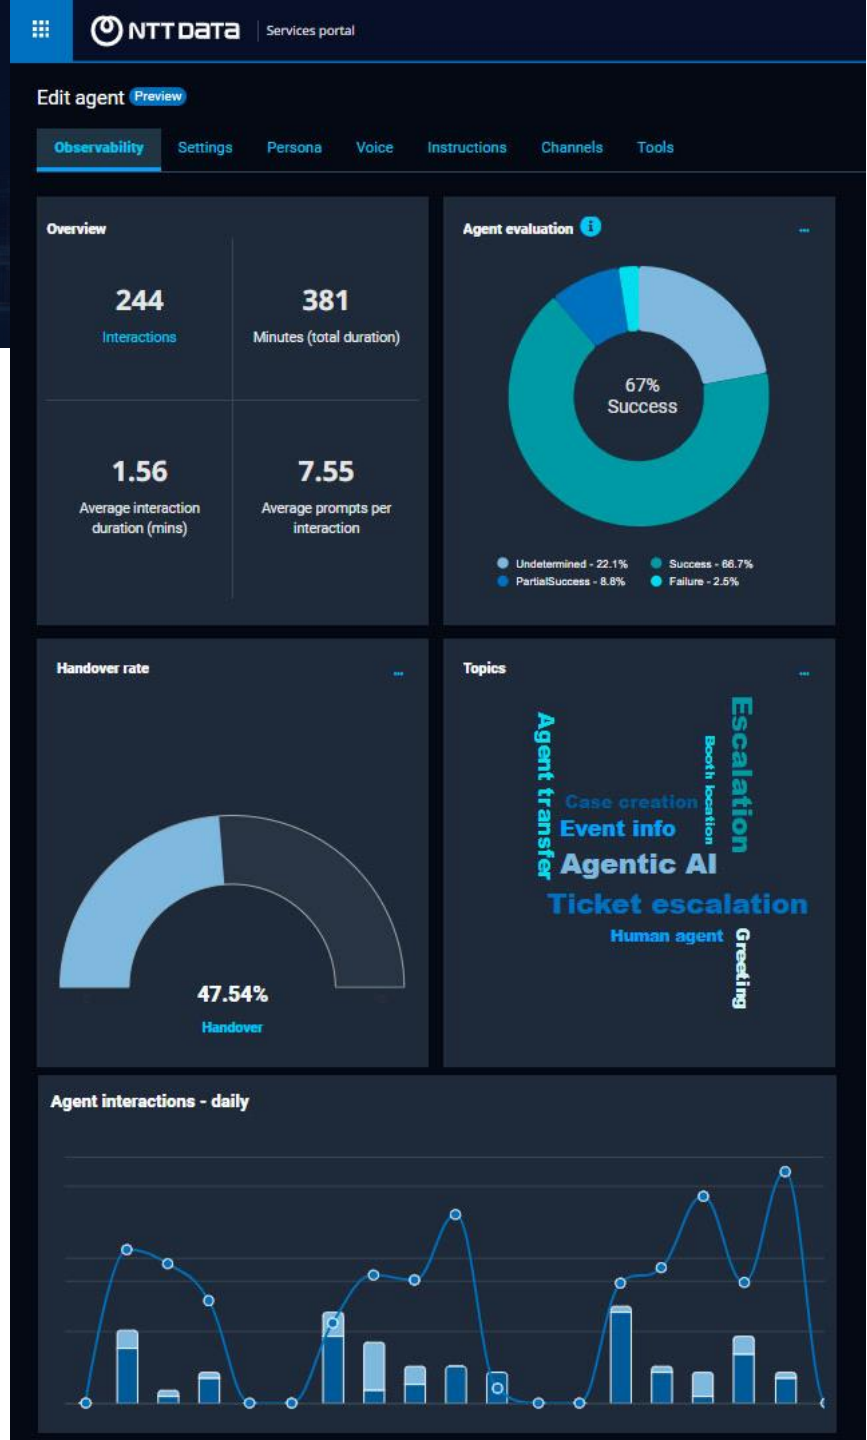
Advanced real-time LLMs combined with secure agentic orchestration protocols enable scalable, knowledge-augmented and controlled AI conversations.

## Challenge

As enterprises scaled AI agents across voice and digital channels, communications became fragmented, costly and difficult to govern. Disconnected systems, inconsistent personas and rising LLM consumption threatened customer experience, compliance and ROI. Organizations needed control, scalability and measurable performance, without slowing innovation.

## Impact and Applications

- Improves first-contact resolution and reduces handling time.
- Optimizes LLM consumption, lowering operational costs.
- Increases customer satisfaction (CSAT and NPS).
- Accelerates deployment and enables scalable operations.



# AI compliance reporting

## Ensuring trust, transparency, and regulatory readiness in AI

### Core Idea

We developed a centralized AI compliance reporting platform to enable organizations to manage regulatory obligations and governance at scale, by leveraging an AI registry, structured documentation workflows, and policy-aligned reporting controls.

The solution provides a comprehensive AI registry to document and track all AI initiatives, ensuring required technical, operational, and risk information is captured and maintained.

Designed for organizations running AI in production, the platform embeds governance into operational processes—transforming compliance from manual documentation into a continuous, auditable enterprise capability.

### Challenge

As AI adoption accelerates, organizations face growing pressure to demonstrate transparency, accountability, and regulatory alignment.

Evolving frameworks, such as the EU AI Act require structured documentation and oversight. Establishing enterprise-wide AI governance while scaling innovation has become a strategic and operational challenge.

### Impact and Applications

- Cost efficiency: Delivers enterprise-grade AI governance and reporting capabilities at a substantially lower total cost of ownership compared to traditional compliance platforms.
- Customizability: The platform can be flexibly tailored to align with regulatory requirements, internal policies, and the scale of AI adoption within each organization.



# Social simulation platform

## Enhancing AI personas for complex interactions and behaviour analysis

### Core Idea

We developed a social simulation platform with Digital Humans to enable businesses to achieve detailed insights into people's preferences and behaviours by leveraging real data and GenAI models to shape and enliven AI Personas in group interactions.

The social simulation platform includes a profiling system that uses advanced AI models to create dynamic Digital Humans representing various customer segments or employee prototypes. These Digital Humans engage in realistic debate within virtual chatrooms, providing valuable insights into preferences and behaviours through sophisticated social simulations and data analysis. The solution is available across different industry-specific scenarios.

### Challenge

Social Simulation enables the transformation of behaviour analysis from a months-long, expensive process into an immediate and affordable competitive advantage by creating virtual discussion groups featuring diverse, consistent digital personas that talk, sound, and behave like real people.

### Impact and Applications

- Reduced organizational complexity: Overcomes the limits of traditional social interactions such as focus groups, reducing organizational complexity and related costs.
- Future-ready research efficiency: Anticipates future people's needs, surpassing traditional research methods in efficiency and scalability.



# Life-saving AI

## Adaptive, data-driven crisis response at city scale

### Core Idea

We built an AI-powered crisis response system to enable cities to optimize large-scale evacuations and emergency coordination, by leveraging multimodal real-time data, collaborative AI, and human-in-the-loop decision support.

The platform generates adaptive evacuation scenarios based on continuously updated satellite imagery, video feeds, and audio signals. Rather than replacing human decision makers, the system augments coordination with dynamic, data-driven intelligence that evolves as conditions change.

The solution detects emerging risks, recommends response adjustments, and supports rapid decision-making at metropolitan city scale.

### Challenge

In major metropolitan areas, tens or even hundreds of thousands of people may be affected during a crisis.

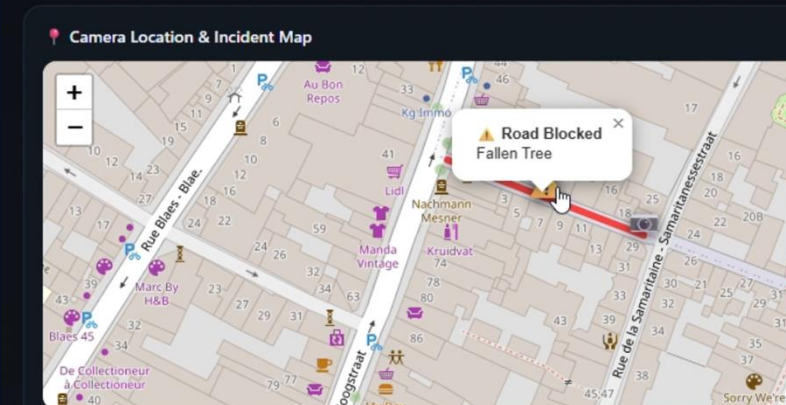
Static emergency plans cannot adapt fast enough to rapidly evolving conditions.

The real challenge is to integrate live, multimodal data and continuously adjust evacuation strategies in real time to protect lives.

### Impact and Applications

- Enhanced response: The system enables faster, data-driven emergency response through real-time analysis of satellite, audio, and video intelligence.
- Improved coordination: Decision-makers gain dynamic situational awareness, improving coordination and accelerating life-saving actions.
- Future expansion: The next step is scaling the solution to additional cities and advancing GenAI-powered crisis simulations.

### Video Event Detection



# Unlocking clinical data through conversation

Secure, instant access to complex population data for clinicians

## Core Idea

We developed an AI-powered conversational assistant to enable clinicians, researchers, and healthcare managers to access and analyze OMOP-CDM clinical datasets using natural language, by leveraging an agentic architecture with embedded privacy and governance controls.

The solution translates clinical questions into schema-aware SQL queries and returns compliant, anonymized results, eliminating reliance on specialized technical language teams. By bridging the gap between clinical reasoning and database query languages, it democratizes access to population-level insights while maintaining strict regulatory safeguards.

Built on LLM-driven intent detection and controlled function-calling workflows, the platform ensures secure, auditable data interactions that scale across healthcare institutions.

## Challenge

Healthcare professionals required access to complex OMOP datasets but depended on technical teams to generate SQL queries, causing delays and inefficiencies. At the same time, increasing regulatory scrutiny demanded strict privacy and governance standards. The opportunity was to democratize clinical data access without increasing compliance risk.

## Impact and Applications

- Efficiency: Reduced time to generate cohorts and extract clinical data.
- Privacy: More consistent query generation with built-in anonymization, reducing privacy and regulatory risk.
- Speed: Faster access to population-level insights improves research and operational decisions.
- Scalability: Modular cloud architecture enabled easy deployment across institutions.



# Real world user experience KPI's

## Virtualised test technologies for cloud environments

### Core Idea

We developed and deployed a virtualized testing platform to enable organizations to measure real-world user experience KPIs in live cloud environments, by leveraging instrumentation-grade methodologies embedded directly into cloud infrastructure.

The solution captures critical performance indicators, including latency, reliability, and service responsiveness across operational networks. By integrating testing into production environments, it provides objective, data-driven evidence of performance maturity rather than relying on benchmarks.

By validating real user data in operational environments, the platform provides objective evidence of performance maturity and demonstrates the tangible benefits of IOWN technologies and low-latency, high-reliability interconnect solutions.


### Challenge

Organizations need reliable ways to measure real user experience KPIs in live cloud environments.

The challenge is to objectively validate the performance and maturity of IOWN technologies, proving tangible benefits such as low latency and high reliability in real-world deployments.

### Impact and Applications

- Data validation: Enables data-driven decision-making for IOWN use cases by validating user experience KPIs on live networks and reducing risk for end users.
- Joint performance verification: Supports both partners and clients, providing measurable evidence of the advantages of IOWN technologies in real-world cloud deployments.



# Optical Network Digital Twin

Converting optical networks into flexible, robust systems for the AI age

## Core Idea

We developed a real-time optical network digital twin to enable telecom operators, hyperscalers, and data center providers to deliver reliable, scalable inter-data center connectivity for AI workloads, by combining physical-layer monitoring with software-driven feedback control.

The breakthrough lies in making deployed, multi-vendor optical networks fully visible and controllable in real time, without replacing existing infrastructure. By virtualizing optical devices and integrating advanced DSP-based physical-layer parameter monitoring, the solution transforms traditionally opaque networks into programmable, software-defined systems.

This architecture enables continuous optimization, improves operational control, and provides a scalable foundation for AI-driven infrastructure growth.


## Challenge

AI workloads are rapidly increasing traffic between geographically distributed data centers. Yet optical transport networks remain complex, opaque, and largely manually operated.

Limited real-time visibility and control constrain scalability, service assurance, and responsiveness at a critical moment for AI infrastructure expansion

## Impact and Applications

- Enhances service reliability through real-time visibility.
- Reduces operational costs in existing (brownfield) environments.
- Enables scalable, distributed connectivity for AI data centers.
- Next step: Expanding to multi-domain automation and AI-driven predictive optimization to support global standardization.



# Concepts of APN Controller technology

Accelerate nationwide APN delivery with automated optical network control

## Core Idea

We developed an APN Controller to enable telecommunications operators to accelerate nationwide APN service delivery while reducing operational complexity, primarily through optimized path route design.

The platform integrates advanced path provisioning applications, including route optimization, wavelength design, reachability checks, and parameter generation, into a unified architecture. By combining intelligent path design with configuration adapters, the solution simplifies complex optical network operations.

The architecture transforms intent-based service delivery orders into automated configurations and supports multi-vendor optical environments through modular adapters, enabling scalable management of thousands of devices across nationwide infrastructures.

## Challenge

APN service delivery involves complex tasks such as path route design, wavelength planning, reachability validation, and parameter generation.

In addition, operators must manage thousands of devices across nationwide networks while integrating multi-vendor optical transmission equipment, creating significant operational complexity.

## Impact and Applications

- Automation: Standardizes and automates optical network operations, accelerating APN service rollout.
- Efficiency: Reduces provisioning time while ensuring consistency across multi-vendor infrastructures.



# AR infrastructure intelligence

## Augmented Reality mobile app for underground infrastructures

### Core Idea

The Augmented Reality (AR) mobile app enables real-time augmented visualization of underground infrastructures directly on mobile devices.

By combining spatial tracking (ARCore), geolocation data, and structured geospatial datasets (Mapbox Vector Tiles (MVT)), the solution transforms infrastructure data into interactive 3D objects accurately positioned within the physical environment. Users can perform spatial calibration, dynamically toggle infrastructure layers, and inspect object metadata in a stable and aligned AR experience.

The system integrates a native Android AR module with a Flutter-based mobile application, ensuring seamless data exchange and intuitive user interaction.

### Challenge

Current solutions struggle to achieve precise spatial alignment between geospatial data and the physical environment despite GPS and sensor inaccuracies. They face limitations in managing performance constraints related to tile download, data parsing, geometry conversion, and real-time 3D rendering on mobile devices. In addition, they lack consistent and reliable ground height estimation under varying environmental conditions.

### Impact and Applications

- Enhanced field awareness: Overlays infrastructure data onto the real-world environment.
- Reduced interpretation errors: Improves safety through accurate 3D visualization.
- Interactive asset inspection: Enables real-time access to technical attributes.
- Targeted infrastructure analysis: Supports dynamic layer filtering for focused assessment.



# Automating industrial decisions with Physical AI

## Defect-free manufacturing with physical AI

### Core Idea

We developed a Physical AI foundation model running at the edge to access Standard Operating Procedures (SOPs) and verify that assembly tasks are performed correctly in real time. Ensuring that standard SOPs are followed results in higher product quality delivered to end customers.

This demonstration shows the Physical AI model monitoring a shopfloor worker completing an assembly task. Data collection and processing occur directly at the edge, enabling the worker to be notified in real time that the SOP has been successfully followed before the product moves to the next stage of assembly.

Physical AI foundation models can be customized with relatively small amounts of training data and deployed across manufacturing, energy, logistics, smart cities, robotics, and other sectors for use cases including process optimization, task verification, and safety.

### Challenge

Manufacturers need to reduce defects in production processes while lowering rework, repair, and RMA costs. Ensuring that Standard Operating Procedures are consistently followed during assembly tasks is essential to delivering high product quality and maintaining strong brand perception.

### Impact and Applications

- Minimal training deployment: A physical AI foundation model has been deployed with minimal training and prompt engineering.
- Real-Time SOP compliance detection: Identifies steps missed during manufacturing processes that do not comply with a predefined SOP.
- Defect and rework reduction: Catching errors and correcting them in real time results in fewer defects in finished products and reduced rework costs.



The interface displays a three-step assembly process for a yellow Duplo vehicle:

- 1 Position the Base**: Place the yellow Duplo vehicle frame with wheels flat on the table.
- 2 Add the Body**: Attach the light yellow rectangular yellow pieces on top of the base.
- 3 Secure the Top**: Align and snap the long, curved roof piece onto the vehicle.

The interface shows a timer at 00:30 and a 'REC' button. Below the steps, it displays 'Archetype® x NTT DATA' and 'CALCULATING YOUR ASSEMBLY SCORE...'. A yellow Duplo vehicle is shown in the center. At the bottom, it shows a '66% accuracy score' (2/3 STEPS COMPLETED) and a 'STEPS OVERVIEW' table:

Step	Description	Status
01	Place the yellow Duplo vehicle frame with wheels flat on the table, ensuring it is stable and properly oriented.	✓
02	Attach the rectangular yellow pieces to the base, snapping them into place to form the main body of the vehicle.	✗
03	Align and snap the roof piece onto the vehicle, adjusting as needed to ensure all parts fit snugly and securely.	✓

The interface also includes a 'NEWTON AI SUMMARY' section with a play button and a 'DONE' button at the bottom right.

# Autonomous Data&AI platform

Turning data operations into autonomous, value-driven performance

## Core Idea

The Autonomous Data & AI Platform draws inspiration from autonomous vehicles: autonomy does not mean rebuilding the car, but augmenting it with intelligence that can perceive, decide, and act. In the same way, we do not reinvent the data platform, we enhance it.

We introduce an AI-powered autonomous layer that surrounds the platform like an exoskeleton, continuously monitoring it, detecting anomalies, diagnosing root causes, and resolving incidents before they escalate. Operating on top of the existing architecture, it enables operations teams to manage complexity more efficiently, reduce manual intervention, and lower OPEX, without disruptive dismantling or replatforming.

The Autonomous Data & AI Platform is powered by Agentic AI enriched with contextual understanding of the platform itself, including logs, metadata, code bases, and documentation, enabling closed-loop automation across heterogeneous legacy and cloud-native environments.

## Challenge

Data is mission-critical, and operations must guarantee trust and availability. Yet scaling value often increases complexity and cost.

The challenge is to break this trade-off by shifting the operating model from human-intensive to human-empowered, with progressively automated detection, diagnosis, and actuation.

## Impact and Applications

- Significant reduction in incident resolution time: Many issues resolved automatically
- Decrease in recurring operational errors
- Lower total cost of ownership (TCO) for Data & AI platforms
- Increased operational continuity and resilience
- Operations teams freed to focus on high-value, strategic initiatives



# Critical infrastructures monitoring via drone fleet

## Orchestration and intelligence for collaborative drone fleet missions

### Core Idea

Critical infrastructures such as dams, energy facilities, industrial plants, and large campuses require continuous monitoring to ensure safety, resilience, and operational continuity.

We developed an autonomous drone fleet platform designed to monitor and protect critical infrastructure at scale. The solution combines AI-driven perception, cooperative UAV coordination, and real-time data processing to enable continuous surveillance, threat detection, and adaptive coverage across large and complex environments.

By integrating robotics, computer vision, intelligent control systems, AI-based anomaly detection, and situational awareness, the platform delivers scalable, resilient, and automated infrastructure monitoring. It enables coordinated multi-drone fleet operations, dynamic coverage adaptation based on risk or events, and rapid deployment in emergency or temporary scenarios.


### Challenge

Traditional surveillance approaches face several limitations as limited coverage and scalability of fixed sensors, delayed response to dynamic or remote threats, high operational costs for manual inspection and inflexibility in adapting to changing risk scenarios

Additionally, large and geographically distributed infrastructures demand mobile, autonomous, and intelligent systems capable of real-time threat detection and coordinated response.

### Impact and Applications

- Enhanced infrastructure security and resilience: Supports continuous, adaptive monitoring, reducing human exposure to risk while improving response time to potential threats.
- Provides a scalable foundation for AI-driven autonomous protection systems.



# RFQ Companion

## The tool for bid management automation

### Core Idea

We developed a GenAI-powered tool to enable the client's employees to optimize tender analysis procedures in order to automatically generate an RFQ draft by leveraging multimodal GenAI models.

The solution includes an advanced search engine for retrieving previous tender documentation matching the required technical expertise, an AI-based agent for answering questions and summarizing tender documents with image-processing capabilities, and an automated system for generating response drafts to Requests for Quotation (RFQs).

### Challenge

Tender response processes are often manual, document-heavy, and time-consuming. Organizations must extract and interpret information from multiple heterogeneous sources, align technical expertise, and prepare high-quality RFQ drafts under tight deadlines. This creates effort-intensive workflows and slows response times.

### Impact and Applications

- Faster tender response: Shortens response time by automating information extraction and draft generation.
- Higher documentation quality: Delivers more consistent and optimized RFQ documentation.
- Simplified information retrieval: Improves comprehension and reuse of relevant tender knowledge.



# Vehicle under attack

## Making hidden cyber risks visible for secure-by-design automotive systems

### Core Idea

We designed and executed a live cyberattack simulation to enable automotive OEMs and engineering teams to make cyber risks tangible, by demonstrating how unsecured vehicle communications can directly impact real-world system behavior.

Using an Automotive Test Rig integrating cockpit, infotainment, and selected ECUs, we simulated controlled CAN and Ethernet message manipulation from an external device. Rather than “hacking” a vehicle, the objective was to translate abstract vulnerabilities into observable cyber-physical effects.

This hands-on approach bridges the gap between theoretical compliance and operational resilience, helping decision-makers understand how communication weaknesses can escalate into safety, regulatory, and reputational risks.


### Challenge

As vehicles become software-defined and increasingly connected, the attack surface expands significantly. Security features such as authentication and secure communication are effective only when rigorously validated and tested.

Automotive leaders must move beyond documentation-based compliance (UNECE R155/R156 and ISO/SAE 21434) toward demonstrable, real-world cybersecurity resilience.

### Impact and Applications

- Risk visibility: Cybersecurity testing services and a tool that accelerates TARA and SBOM analysis.
- Faster TARA: Precise security requirement selection for each ECU, ensuring appropriate protection levels and saving money and time.
- Compliance confidence: Support ISO/SAE 21434 and UNECE R155 compliance and supplier coordination, reducing risk, time, and cost while strengthening cyber resilience.



# Cisco routed optical networking

## Converging IP and optical, communication infrastructure for the AI Era

### Core Idea

We address the rapid growth of data center traffic by evolving communication infrastructure to support higher capacity while reducing power and space consumption, by converging IP routing and optical transport through Cisco Routed Optical Networking while automating provisioning and troubleshooting.

The solution integrates high-speed 400Gbps and 800Gbps QSFP-DD Digital Coherent Optics (DCO) interfaces directly into routers and switches, eliminating the need for dedicated, power-hungry WDM transponders.

Routers aggregate, protect, and restore traffic, simplifying the WDM Line System, reducing ROADM node complexity, and enabling highly efficient point-to-point links.

Standard and open network management solutions and APIs (NETCONF, OpenConfig, TAPI) simplify operations and enable intent-based network automation.


### Challenge

AI workloads place unprecedented demands on network infrastructure, requiring higher bandwidth, ultra-low latency, and scalable power efficiency.

At the same time, networks must evolve toward secure, observable, and highly scalable fabrics capable of supporting massive, synchronized GPU workloads. Solutions such as Cisco RON and IOWN provide a practical and sustainable path to address these evolving requirements.

### Impact and Applications

- Simplified architecture: Converged IP and optical transport architectures enable content providers, service providers, and enterprises to meet bandwidth and latency requirements through a simplified and energy-efficient network design.
- Scalable growth: The introduction of 800Gbps DCO interfaces and the roadmap toward 1.6Tbps in the near future- enables continued scaling while leveraging the same Open Line System infrastructure.



# Automatic analysis of citizens' complaints

## Enhancing problem-solving efficiency in city administration

### Core Idea

We developed a GenAI-powered system to enable city officials to achieve faster and more effective troubleshooting by leveraging text classification and sentiment analysis.

The goal is to process citizen complaints by extracting key insights such as urgency, sentiment, and topic, which are then visualized through an Augmented Reality (AR) application that overlays data onto an interactive city map, showing real-time issues at various points of interest.


### Challenge

Modern cities rely on citizen feedback to improve public services and ensure effective resource allocation. Cities receive a large volume of complaints daily and, therefore, manually processing them is time-consuming, making it difficult for administrators to prioritize urgent issues effectively.

This system helps city officials prioritize urgent issues and make data-driven decisions to improve public services and resource allocation.

### Impact and Applications

→ Faster, citizen-centric city operations: Enhances decision-making, improves citizen engagement, and accelerates issue resolution by providing city officials with real-time, intuitive insights into pressing concerns, ultimately improving overall quality of life.



# Security on the microfactory

## Strengthening OT cybersecurity without disrupting production

### Core Idea

We implemented an Operational Technology (OT) cybersecurity solution within the Microfactory to enable manufacturers to gain full visibility and threat detection across their XIoT environments, by leveraging passive, real-time network monitoring that does not impact production operations.

Network traffic exchanged among factory components, such as PLCs, sensors, and controllers, is mirrored to a dedicated security sensor for real-time packet analysis. This enables comprehensive asset discovery, identification of device characteristics, detection of vulnerabilities, and monitoring of network anomalies.

The solution is powered by the Claroty platform, a market leader in XIoT cybersecurity, with NTT DATA serving as a Global Elite Partner to deliver secure-by-design industrial environments.

### Challenge

Many industrial organizations lack comprehensive visibility into the assets deployed across their OT environments. Without clear insight, vulnerabilities remain hidden and risk exposure increases.

Cyber protection must be continuous and non-intrusive, ensuring security improvements without disrupting production continuity.

### Impact and Applications

- Risk reduction: Identifies vulnerabilities and anomalies early, reducing exposure to operational and safety incidents.
- Compliance readiness: Supports NIS2 alignment.
- Incident response: Accelerates detection and containment of cyber threats across industrial environments.



# Patient clinical care summary generator

## Generation of clinical care summaries for healthcare professionals

### Core Idea

We developed a GenAI-powered clinical summarization system to enable clinicians quickly understand a patient's overall condition.

The solution integrates structured EHR data and unstructured clinical notes, automatically generating concise, context-aware summaries that synthesize fragmented medical information into a coherent patient overview.

Recognizing that critical clinical information is often split between structured fields and free-text notes we combined them within a reasoning framework that allows GenAI to interpret, cross-reference, and contextualize heterogeneous inputs.

This shift from data aggregation to AI-driven clinical reasoning transformed the approach from reporting to intelligent synthesis.

### Challenge

Healthcare professionals operate in high-pressure environments where patient data is fragmented across multiple systems and formats. Clinicians need to manually review EHR fields, lab results, imaging reports, and text notes to form a complete picture of a patient's status.

The problem was clear: clinical time is scarce, and fragmented data increases cognitive load, risk of oversight, and inefficiencies in care.

### Impact and Applications

- Efficiency: Significant reduction in time spent reviewing and synthesizing patient information.
- Quality & Safety: Lower risk of overlooking critical insights hidden in free-text notes.
- Decision-making: Faster access to comprehensive patient overviews improves clinical confidence and coordination.
- Scalability: Cloud-based modular architecture enabled deployment across institutions.



# From static planning to living network intelligence

## Sharper coverage insight, stronger capital decisions

### Core Idea

We integrate high-resolution satellite intelligence into RAN planning environments to enable telecom operators make more accurate decisions relying on a continuously updated, data-driven foundation for network expansion and optimization.

The platform delivers environmental inputs at resolutions down to 40 cm, significantly improving radio planning accuracy, particularly for higher frequency bands where terrain and urban detail directly affect coverage modeling.

Through seamless integration with existing RAN planning tools and real-world device performance data, the platform enables operators to evaluate infrastructure investments, validate model outputs, and quantify quantify population impact, and prioritize site deployments with greater confidence and capital efficiency.


### Challenge

Operators must expand coverage and optimize capacity while justifying significant CAPEX in rapidly evolving urban environments. Yet planning models often rely on static terrain data and infrequent updates, limiting accuracy.

Without continuously refreshed environmental intelligence, investment decisions carry uncertainty and hidden coverage gaps persist.

### Impact and Applications

- Capital efficiency: Targeting investments where coverage and population impact are highest.
- Planning precision: Improve modelling accuracy through 40cm resolution and continuous updates.
- Validation & automation: Align planning outputs with real device performance and advance towards autonomous networks.



# AI for CV matching

## Automation in candidate evaluation

### Core Idea

We developed an AI-based system to support HR departments in managing high volumes of applications by leveraging Generative AI models to automatically extract, analyze, and match information from CVs with job descriptions.

The system enables automated evaluation of candidate profiles by identifying relevant skills, experiences, and qualifications, even when CVs and job descriptions are written in different languages. Matching parameters can be customized according to specific HR needs and role requirements.

### Challenge

Managing large volumes of applications manually is time-consuming and prone to inconsistencies. HR teams must evaluate CVs written in multiple languages, ensure fair and accurate assessments, and effectively match candidate skills to job requirements. This creates inefficiencies and slows the hiring process.

### Impact and Applications

- Accelerates the hiring process: Reduces manual screening time, increases consistency in candidate evaluation.
- Enables HR teams to focus on higher-value recruitment activities: Facilitates tedious multilingual CV assessment and performs more accurate skill-to-role matching.



# Cyber resilience

A unified interface for civil cyber operations

## Core Idea

We developed a human-centered cyber operational interface built around a territorial Digital Twin of the entire infrastructure network assets, interconnections, dependencies, and cyber risk exposure.

This cyber-focused Digital Twin provides a unified Common Operational Picture, enabling operators to visualize how threats dynamically impact systems, services, and mission continuity. By integrating threat intelligence, asset monitoring, and risk assessment into a single environment, the platform transforms fragmented cyber data into structured operational insight.

AI-assisted decision support orchestrates the full lifecycle; Plan, Simulate, Execute, reducing cognitive load and accelerating response in mission-critical civil environments.

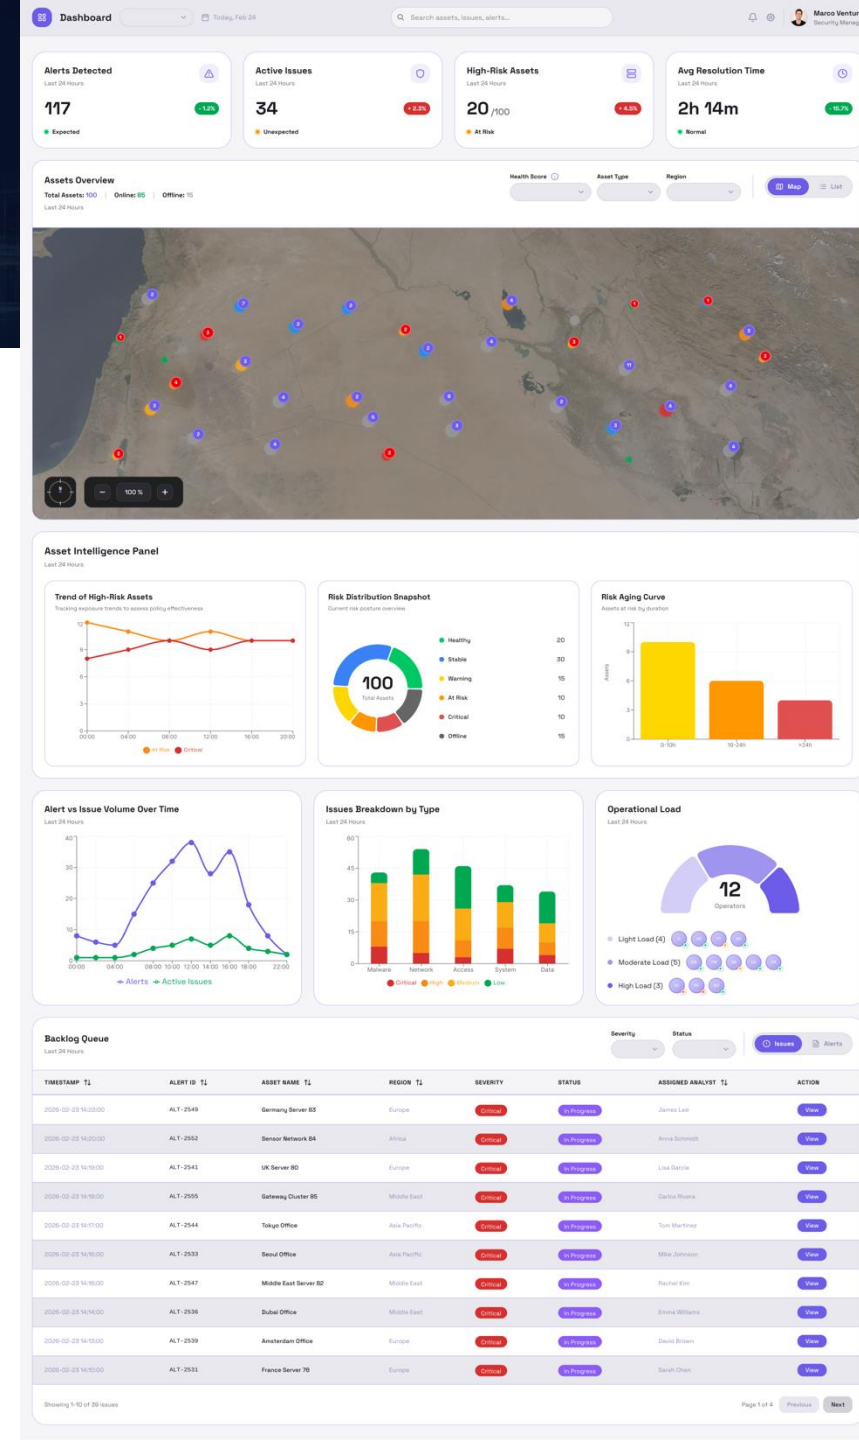
## Challenge

Civil critical infrastructures face increasingly complex and fast-evolving cyber threats, while tools and data remain fragmented across systems.

This fragmentation limits situational awareness, slows prioritization, and increases operational risk. Cyber resilience requires unified, secure, and operator-centered control environments.

## Impact and Applications

- Enhances operational efficiency and response readiness: Accelerates risk evaluation and decision-making across the cyber lifecycle.
- Designed with security-by-design principles: With Interoperability, and data sovereignty at its core, it strengthens resilience of civil critical infrastructures and multi-agency operations.



# Stay ahead of exposure

Accurate risk intelligence that turns portfolio visibility into strategic control

## Core Idea

We unified satellite intelligence, trusted third-party risk data, and insurer-owned information into a dynamic 3D Risk Intelligence Platform, enabling insurers to assess, validate, and actively manage exposure across underwriting and portfolio management.

We transform fragmented hazard and exposure data into verified, asset-level intelligence that directly informs underwriting and risk decisions.

By reconciling multiple data sources, validating geolocation accuracy, and linking intelligence directly to insured assets, we create a continuous 3D portfolio view that supports disciplined growth and active risk protection.

High-resolution Earth observation combined with a Digital Twin foundation provides spatial clarity, while an agentic AI layer continuously interprets multi-source signals and generates forward-looking, explainable exposure insights within insurer systems.


## Challenge

Climate volatility and capital pressure demand smarter portfolio expansion and tighter exposure control.

However, static maps and fragmented tools limit visibility into how catastrophe risk evolves across existing and prospective assets. Without a unified view, aggregation risk increases, underwriting opportunities are missed, and loss prevention becomes reactive.

## Impact and Applications

- Confident growth: Evaluate new and existing assets within a validated 3D risk context to improve portfolio quality and premium performance.
- Proactive exposure management: Monitor portfolio-wide risk in real time, detect emerging concentrations, and enable preventative action before events translate into losses.



# Trusted identity for AI Agents

## Secure delegation and verifiable autonomy in AI-driven transactions

### Core Idea

We establish a trusted identity and delegation mechanism to enable users and AI agents to perform autonomous transactions securely and with full data control, by leveraging decentralized identifiers (DID) and verifiable credentials (VC).

The solution cryptographically verifies delegation, allowing service providers to confirm user intent and authorization without relying on centralized identity providers. This introduces a trust layer for autonomous AI actions.

Powered by DID & VC, which ensure tamper-resistant identity, selective disclosure, and decentralized validation - critical for establishing trust between users, AI agents, and service providers.

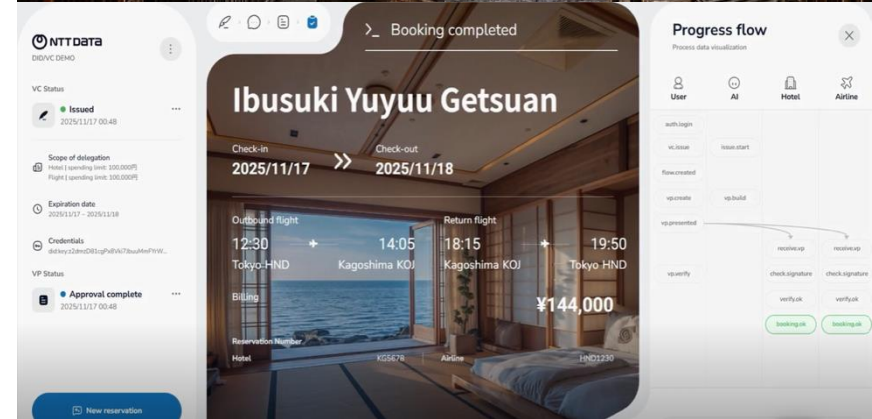
### Challenge

The challenge arises from the inability to verify whether AI-generated requests truly reflect a user's intent and authorized delegation, limiting trust in autonomous transactions. As AI agents increasingly handle tasks such as bookings and payments, service providers lack a reliable way to confirm that actions are genuinely approved by the user.

A verifiable trust layer is therefore required to reduce risk and enable secure, scalable AI-driven services.

### Impact and Applications

- Efficiency gains: The solution adds value by enabling trusted, automated transactions through verifiable delegation, reducing manual checks and lowering privacy and authorization risks. This increase efficiency and reliability across service interactions.
- Scalability: The same trust framework can scale to other domains - such as public services or multi-agent environments - expanding reuse and ecosystem reach.



# Digital Twin of a robotic arm

## Enabling brand-agnostic robotics through real-time 3D Digital Twins

### Core Idea

We developed a Digital Twin of industrial robotic arms that integrates 3D simulation, real-time connectivity and AI-based control into a unified and interoperable platform.

By leveraging NVIDIA Omniverse for photorealistic simulation, bi-direction communication, 5G connectivity and AI algorithms for object recognition and motion control, the solution enables seamless integration of robots from different manufacturers within a common digital environment.

The platform allows manufacturers to simulate, test and optimize robotic operations before deployment, reducing dependency on proprietary ecosystems and improving operational flexibility.

### Challenge

Manufacturers face increasing complexity in integrating robotic systems from different brands while ensuring real-time communication, interoperability and precise motion control.

Traditional robotic environments are often vendor-specific, limiting flexibility and increasing integration costs.

### Impact and Applications

- Test before deploy: Enables realistic simulation and validation of robotic operations in a connected 3D environment aligned with physical processes.
- Seamless multi-Brand integration: Accelerates integration across heterogeneous robotic systems and improves efficiency through AI-driven control.
- Lower risk, greater flexibility: Reduces production risk and integration costs while providing a scalable foundation for industrial digitalization.



# The optical foundation for Ultra-Low-Latency AI

Proven deterministic performance with open, scalable control

## Core Idea

We developed an open, disaggregated optical architecture that delivers ultra-low latency and jitter-free performance essential for next-gen innovation, from AI and data replication to any mission-critical service demanding absolute precision.

The key is pure light transport. By bypassing traditional electrical conversions, this architecture eliminates the latency and jitter that have long constrained standard network performance.


By integrating high-performance optical hardware with software-defined management, the solution transforms rigid networks into agile, programmable ecosystems, allowing infrastructure to scale at the speed of innovation.

## Challenge

Traditional networks have become a bottleneck for modern innovation. They struggle to keep pace with the fluid, high-speed reality of the digital era, acting as a silent brake on the progress that next-generation services demand.

## Impact and Applications

- Network as a growth catalyst: Transforms the network into a catalyst for growth, delivering superior performance, cost efficiency, and seamless operational simplicity.
- Enabling next-generation workloads: Establishes a new industry benchmark for optical data center interconnection. Supports distributed GPU training for AI, high-velocity data replication, mission-critical disaster recovery, telemedicine, and remote device control.



# Bring the power of Cloud everywhere!

## How to adopt a platform approach to make real Edge Cloud services

### Core Idea

We designed a sovereign, open-standard Edge Cloud platform to centrally manage distributed infrastructure through automated host, fault, software, configuration, and service management. Built to support mission-critical scenarios such as O-RAN, Smart Grid, Automotive, and Cloud for Defense, the platform addresses the need for low latency, data sovereignty, operational resilience, and real-time processing at the edge.

By combining cloud-native automation, open ecosystems, and strong governance models, the Distributed Edge Cloud platform accelerates innovation while ensuring security, compliance, and operational excellence across complex distributed environments. The architecture leverages a containerized version of OpenStack on Kubernetes, enhanced with KubeVirt and best-of-breed CNCF components, delivering scalable virtualization capabilities and an open, future-ready foundation for edge services.

### Challenge

To deliver next-generation services, computing capabilities must be deployed as close as possible to end users.

At the same time, centralized control is required to reduce OPEX while ensuring high performance, low latency, security at the edge, geographic distribution, data and source-code sovereignty, and edge isolation for critical environments.

### Impact and Applications

- Scalable Edge deployment: Enables distributed edge infrastructure across critical sectors including O-RAN, Smart Grid, Smart Manufacturing, Autonomous Driving, and Defense—supporting real-time, mission-critical workloads at scale.
- Open architecture: Avoids vendor lock-in, improves economic efficiency, and supports faster ROI through scalable, sovereign edge infrastructure.



# Accelerating drug discovery

## AI-driven computational foundation for the next era of drug discovery

### Core Idea

We optimize molecular structures to enable pharmaceutical R&D teams to identify promising compounds faster and improve research efficiency through multi-objective AI-driven search and advanced mathematical optimization techniques.

Our approach shifts molecular design from an experience-driven, trial-and-error paradigm to a structured, data-driven optimization framework in which trade-offs between efficacy, stability, and pharmacological properties are explicit and measurable.

By combining multi-objective optimization methods, including Bayesian techniques, with AI-based molecular structure exploration, we enable the simultaneous evaluation of multiple performance criteria and a more intelligent navigation of the chemical search space.

### Challenge

Drug discovery is a long and high-cost process, and identifying promising compounds remains a critical bottleneck. Even minor structural modifications can significantly impact efficacy, stability, and pharmacological performance, making optimization complex and uncertain. Traditionally, these adjustments have relied on expert intuition and iterative experimentation, limiting scalability and slowing innovation.

### Impact and Applications

- Faster candidate identification: Accelerates molecular exploration and shortens the path toward identifying high-potential drug candidates.
- R&D efficiency: Reduces reliance on manual trial-and-error processes, supporting more informed decision-making.
- Next step: The integration of Quantum Machine Learning will further enhance predictive capabilities and enable broader adoption across drug discovery domains.



# e-Longevity: where health meets emotions

## Turn everyday wearable data into clear and objective health guidance

### Core Idea

We analyze time-aligned vital signals from multiple wearables to enable individuals and clinicians to detect patterns and correlations that anticipate health changes. The solution leverages multivariate time-series analytics and an LLM (Gemini) to translate findings into clear, objective guidance.

The real value lies not in isolated metrics, but in relationships over time, how signals interact and evolve. We run multivariate analyses to surface cross-signal correlations and recurring patterns (e.g., sleep degradation preceding glucose instability, or stress markers correlating with reduced recovery), transforming raw data streams into interpretable signals that support proactive action.

An LLM (Gemini), combined with medical history summaries, converts detected patterns into transparent explanations and actionable next steps.

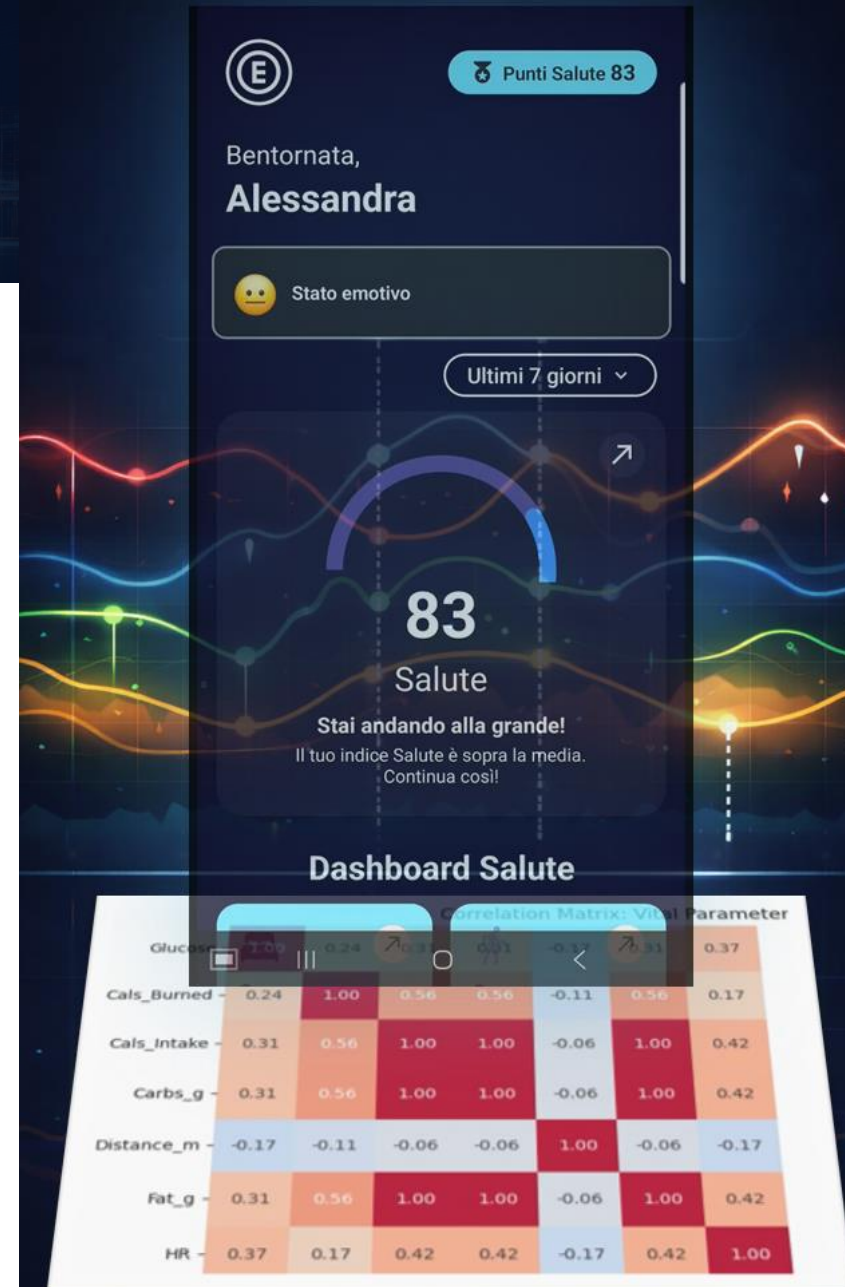
### Challenge

Current wearable applications focus on isolated dashboards, making it difficult to understand why metrics change and which signals move together. Correlations across devices and time are rarely extracted in a robust, clinically meaningful way.

As wearables become mainstream, the bottleneck is no longer data collection, but data interpretation—particularly for early detection and prevention.

### Impact and Applications

- Improved decision-making: Correlation-driven insights and stronger trust enabled by evidence-based explanations (“what changed and which signals co-moved”).
- Faster identification of multi-signal patterns through automation.
- Next steps: Expand the correlation engine to incorporate lagged correlations, seasonality modeling, and anomaly clustering.



# Automatic ESG report analysis

## Improvement of sustainability performance analysis process

### Core Idea

We developed an Agentic AI system to enable sustainability teams to achieve automatic analysis of ESG reports by leveraging Intelligent Document Processing techniques to gather relevant insights from data sources.

The goal is to efficiently analyze both quantitative metrics and qualitative sustainability narratives at scale to support strategic decisions and improve ESG positioning.

The solution developed includes an AI-based system that extracts numerical KPIs and qualitative ESG information from reports, along with a dashboard for viewing the extracted data and making immediate multi-player comparisons. The solution is the prototype of a broader asset that is being developed to support our clients.

### Challenge

Monitoring and comparing sustainability performance against competitors can be difficult and time-consuming due to the heterogeneous and unstructured nature of ESG reporting in the industry. With investors increasingly focused on ESG issues and reporting regulations becoming more stringent, there is a pressing need for an automated solution.

### Impact and Applications

- Comprehensive sustainability benchmarking: Enables a comprehensive view of the company's sustainability performance compared to competitors and best-in-class companies, providing insights into customer positioning.
- Simplified analysis process: Reduces manual work and simplifies sustainability performance analysis.

# Scaling human performance

## Turning brain data into competitive advantage for teams & organizations

### Core Idea

We developed a modular textile-based EEG wearable platform designed to enable individuals, teams, and high-performance professionals to objectively measure and train cognitive states such as focus, mental resilience, and energy regulation.

The system integrates high-fidelity frontal EEG, HBR, temperature, and IMU sensing with Bluetooth connectivity, edge processing, and AI-driven analytics to deliver real-time insights in motion and in professional environments.

The key shift was moving from rigid, single-user headband devices to a scalable, multi-user neuro-performance ecosystem. By embedding a removable core module into flexible textile form factors and enabling dual streaming of raw and processed data, the platform bridges consumer usability with professional-grade performance.

Its architecture combines advanced biosignal acquisition, real-time processing, and secure multi-device connectivity, creating a foundation for scalable B2B deployment and continuous expansion into complementary performance data domains.

### Challenge

Despite rapid growth in the wearable market, there is still no reliable and comfortable solution capable of operating with multiple users simultaneously, in motion, within professional training environments and real-world high-performance settings.

The challenge is therefore strategic as well as technological: transforming consumer neurofeedback devices from personal gadgets into deployable professional infrastructure.

### Impact and Applications

- Objective KPIs: Introduces measurable performance metrics in domains traditionally guided by subjective assessment—enhancing decisions, training, and outcome tracking.
- Critical sectors: Applies to sports, defense, aviation, safety-critical roles (PPE), and leadership performance where cognitive readiness is essential.
- Market opportunity: Addresses ~€12B TAM, ~€4B SAM, and ~€120M initial SOM.



# 未来 Futures

## From insight to action: foresight in the era of mass intelligence

### Core Idea

We transformed NTT DATA Technology Foresight into an interactive GenAI-powered Foresight Assistant to enable executives and clients to translate macrorends into concrete, industry-specific actions, by leveraging agentic AI grounded in curated strategic research.

Foresight often remains static, published, admired, and archived. Our key insight was to shift from “report as artifact” to “report as intelligence layer.” By embedding the six macrorends of Technology Foresight 2026 into a governed, explainable AI assistant, we enable dynamic exploration across industries, functions, and strategic questions in real time. The assistant does not replace human judgment; it augments it, combining structured foresight architecture with conversational AI to create contextual, forward-looking dialogue.

LLM is merged with a formal foresight ontology, enabling structured strategic reasoning rather than simple content retrieval. The LLM generates possibilities. The ontology enforces structure.

### Challenge

Executives face accelerating complexity: mass intelligence, sovereign infrastructure, AI trust, and sufficiency pressures converge simultaneously.

While Technology Foresight 2026 provides strategic clarity, leaders need to operationalize insights in real time, across industries, geographies, and board-level decisions. Static reports alone cannot support continuous strategic dialogue.

### Impact and Applications

- Faster translation of trends into use cases and investment priorities.
- Improved cross-industry ideation in foresight workshops.
- Reduced time-to-insight for executive discussions.
- Grounded in method, enhanced by AI.



# Gregorian Tower digital reconstruction

## Transforming real-world environments into interactive experiences

### Core Idea

We digitally reconstructed the Gregorian Tower (also known as “Torre dei Venti”) to enable Expo Dubai visitors to experience Vatican cultural heritage within an immersive environment, leveraging high-resolution photogrammetric scanning and interactive VR technologies (e.g., Leap Motion).

The key insight was that immersive digital reconstruction can transform a physical environment into a scalable and reusable experience, enabling interactive access beyond physical and operational constraints.

### Challenge

At Expo Dubai 2020, the Vatican Pavilion needed to deliver a distinctive and meaningful visitor experience within physical space limitations and COVID-related constraints.

The challenge was to provide immersive access to a cultural masterpiece while ensuring scalability, accessibility, and compliance with operational requirements.

### Impact and Applications

- Global recognition: Showcased at Expo Dubai 2020, where the pavilion was nominated among the best at the Expo.
- Permanent deployment: Now permanently installed in Dubai and extended into a standalone application for Meta Quest 3.
- Scalable reuse: Adaptable to other physical spaces, enabling reusable immersive environments and advanced digital representations of real-world assets.



# Gyōkō

## Autonomous multi-agent knowledge discovery at enterprise scale

9 Orbit Levels • Distance = Technology Readiness Level • Size = Disruption Potential

### Core Idea

We developed Gyōkō an internal AI-powered platform for continuous research and insight generation across companies, technologies, products, and emerging trends.

Unlike traditional automation platforms, Gyōkō operates without predefined flows or rigid pipelines. Instead, it relies on a goal-based planner agent that continuously plans, executes, evaluates, and replans after every action and reasoning step, adapting dynamically to what it discovers along the way.

This architecture is grounded in two foundational research paradigms: ReAct, which synergizes reasoning and acting in language models, and the Puppeteer agent architecture, which enables flexible orchestration of specialized agents without hardcoded execution paths. Each agent has a specific role and access to specific tools, and the planner coordinates them autonomously toward the research objective, with no human intervention required during execution.

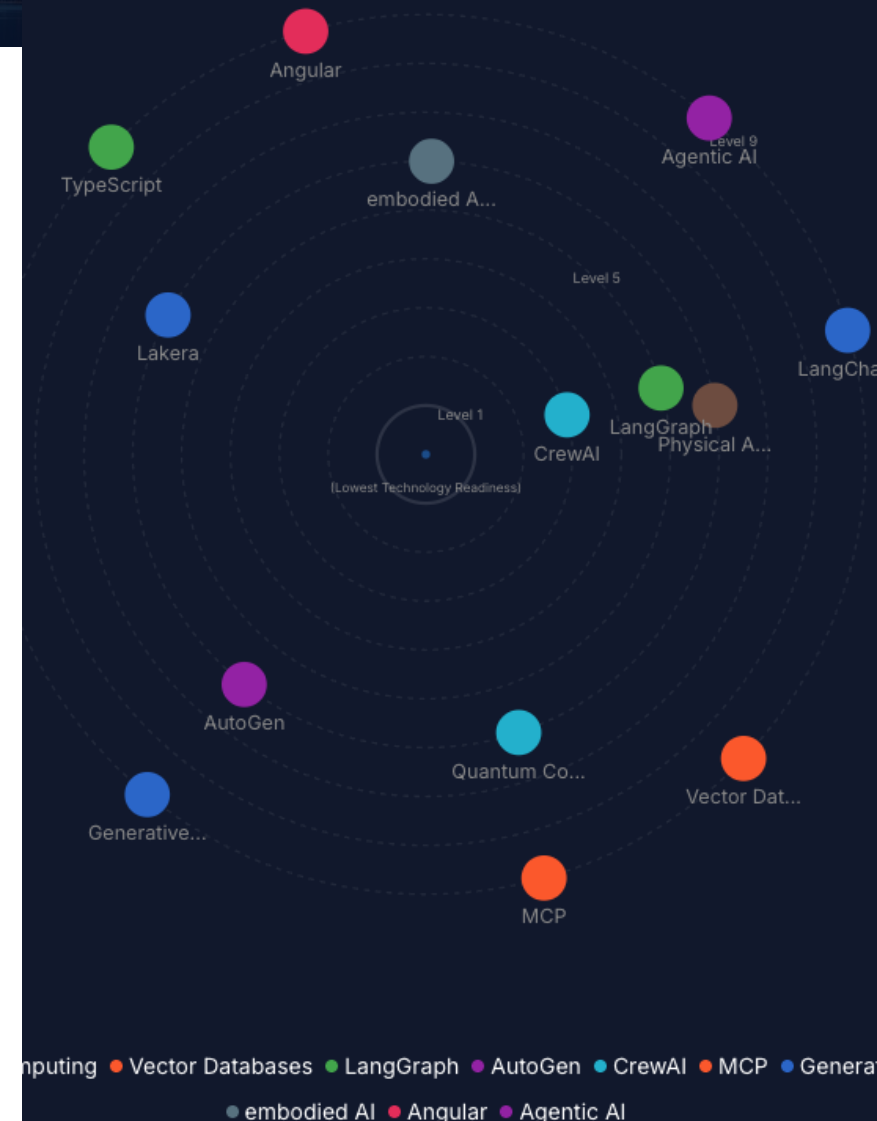
### Challenge

Strategic research is fragmented, slow, and difficult to scale. Teams spend considerable time gathering, validating, and structuring information across multiple sources, leaving limited capacity for analysis and strategic thinking.

Existing automation tools rely on rigid workflows that break when conditions change, making them unsuitable for the open-ended nature of real research tasks.

### Impact and Applications

- Continuous intelligence: Gyōkō transforms scattered information into continuously updated, structured knowledge that accelerates strategic and innovation-focused decision-making.
- Enterprise scale: By tracking companies, technologies, trends, and products from a single workspace, it enables innovation and sales teams to identify opportunities, monitor competitors, and generate client-ready insights faster than any manual process.



# Control the corridor

## High-resolution, tree-level intelligence for proactive risk control

### Core Idea

We transformed satellite intelligence into a dynamic 3D Digital Twin of transmission corridors to strengthen grid resilience by continuously monitoring vegetation encroachment, structural clearance, and wildfire exposure risk.

Powered by Marble Visions, our high-resolution LEO constellation, the platform delivers frequent, up-to-date visibility across transmission lines and the wider buffer zones on both sides of the right-of-way.

Multi-view photogrammetry generates high-fidelity 3D models of terrain, lines, and individual trees, including precise height and clearance analysis. Advanced analytics assess vegetation health and dryness to identify high-risk trees and elevated wildfire exposure before they threaten grid reliability.

The result is a continuously updated, asset-level risk view that supports targeted maintenance, optimized field deployment, and proactive resilience planning.


### Challenge

Vegetation-related outages remain one of the leading causes of transmission disruption. Traditional inspection cycles - manual patrols, helicopters, periodic LiDAR - are costly, infrequent, and reactive.

Without continuous, high-resolution insight, risk builds between inspection cycles, increasing outage probability and operational uncertainty.

### Impact and Applications

- Proactive risk reduction - Identifies encroachment and structurally weak trees early, reducing outage events and emergency response costs.
- Targeted maintenance optimization - Shifts from blanket trimming cycles to precision-based dispatch, focusing crews only where intervention is needed.
- Scalable monitoring: Provides continuous corridor visibility at regional scale without recurring LiDAR dependency.



# Smart textile biosensing platform

## From wearable ECG to cloud-ready physiological intelligence

### Core Idea

The Hitoe® kit is an Android based wearable application designed to collect, visualize, store, and export biomedical data acquired from the hitoe transmitter wearable sensor. The system enables real-time monitoring of multiple physiological and environmental parameters, including ECG, heart rate, three-axis acceleration, temperature, humidity, step count, and battery status.

By combining textile-based biosensing technology (hitoe® conductive fiber) with mobile connectivity, the app transforms a garment-integrated sensor into a portable health and activity monitoring platform.

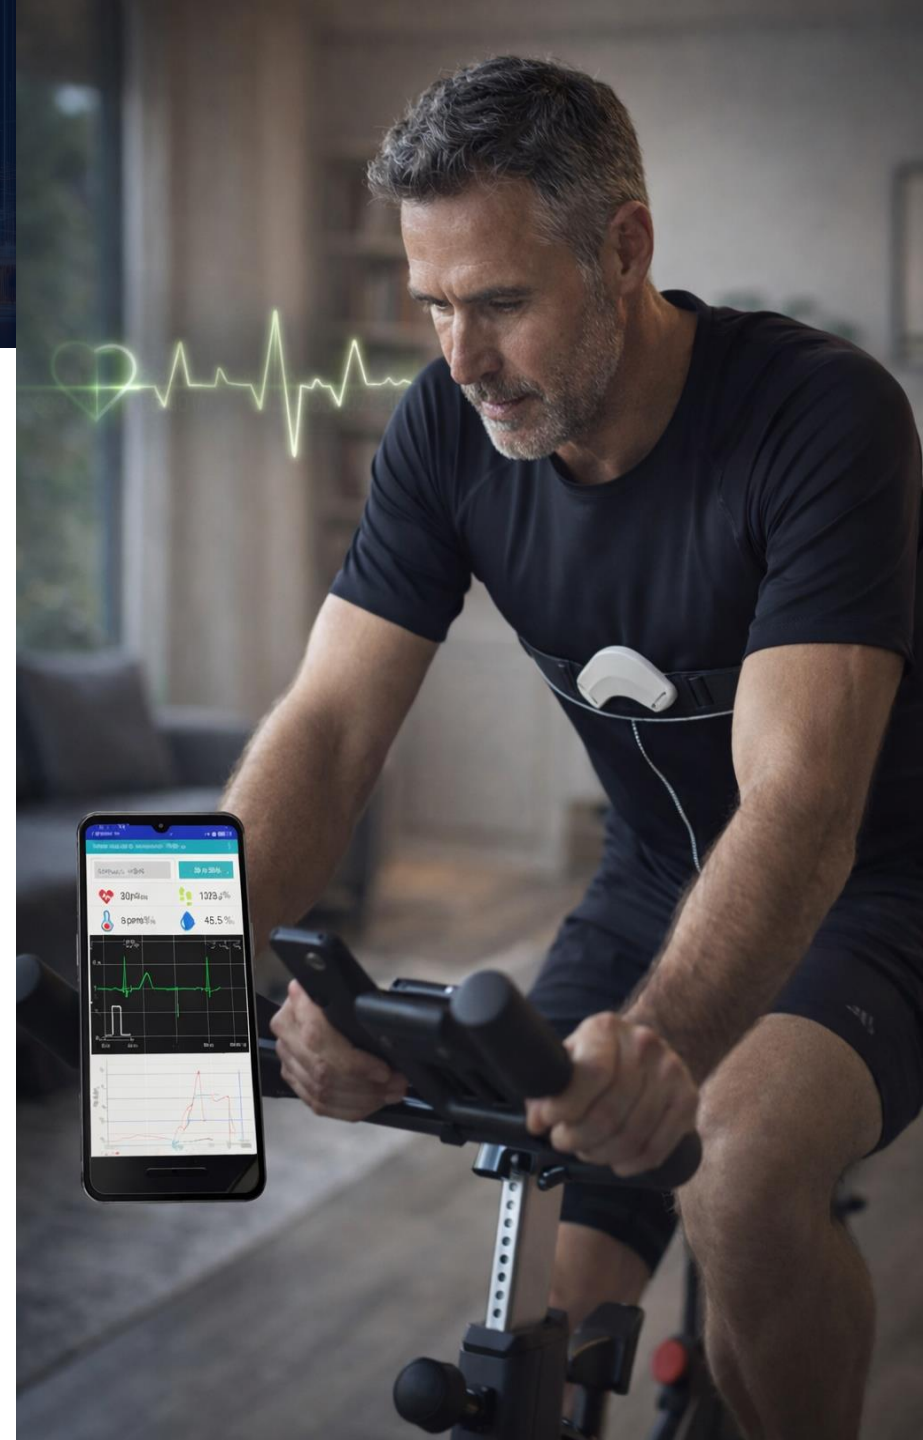
Hitoe ecosystem supports data-driven decision-making across healthcare, industry, and research domains.

### Challenge

Continuous and reliable acquisition of physiological signals in real-world conditions presents technical and operational challenges. Maintaining stable ECG signal quality requires proper sensor contact and effective motion artifact mitigation. In addition, uninterrupted background data collection depends on careful management of OS permissions (Bluetooth, location, battery settings) to ensure continuous transmission and recording.

### Impact and Applications

- Occupational health and safety: Monitors workers' physical strain, fatigue, heat stress, and environmental conditions in real-time.
- Sports and performance optimization: Continuously tracks heart rate, ECG dynamics, & movement data to enhance training programs.
- Remote health monitoring and research: Collects longitudinal ECG and activity datasets for telemedicine, clinical studies, and population health analysis.



# Humanoid robots

## The future of industrial efficiency and safety

### Core Idea

We developed an intelligent humanoid robotics framework that enables human-like machines to perceive, reason, and interact safely within real-world environments.

The innovation lies in integrating advanced AI models, multimodal perception (vision, audio, force, and spatial sensing), and real-time motion control into a unified cognitive and physical architecture.

Unlike traditional industrial robots designed for repetitive and structured tasks, humanoid robots are built to operate in environments designed for humans; adapting to dynamic conditions, understanding context, and collaborating naturally with people.

By combining Physical AI, intelligent vision, and edge-based decision-making with whole-body control and dexterous manipulation, the system enables autonomous navigation, object handling, and task execution with high precision and situational awareness.

### Challenge

Industries today face increasing challenges such as labor shortages, rising costs, and the need for higher efficiency. Many industries struggle to find skilled workers, particularly for repetitive, dangerous, or physically demanding jobs.

Humanoid robots offer a practical solution by seamlessly integrating into human-centric work environments, improving productivity, and reducing risks.

### Impact and Applications

- Workforce augmentation: Can support labor-constrained environments by performing repetitive, hazardous, or physically demanding tasks.
- Flexibility: Can operate in environments designed for humans without the need for significant modifications to factories or assembly lines.
- Safety: Their capability to interact with the surroundings and perform complex tasks ensures reliability and precision while using human tools and machinery.



# Redefining responsibility of information technologies

## Turning digital carbon visibility into measurable performance advantage



### Core Idea

Our ICT carbon calculator baselines, forecasts, and optimizes IT-related emissions, enabling enterprises to reduce their digital carbon footprint and improve decision-making through an intelligent SaaS platform with advanced analytics and automated reporting.

IT emissions are growing rapidly but remain fragmented across applications, cloud, devices, networks, and data centers. Traditional reporting is backward-looking. The breakthrough is a shift from static measurement to predictive, scenario-driven carbon management, embedding sustainability into IT investment and architecture decisions.

Advanced analytics and AI-driven anomaly detection power forecasting, what-if simulations, and automated standards-based reporting across hardware, software, applications, cloud, networks, and data centers all within a vendor-agnostic architecture.

### Challenge

As digital estates expand, carbon visibility becomes increasingly complex. CIOs lack unified, forward-looking tools to quantify emissions and guide optimization decisions. Boards demand measurable reduction pathways, not static reports.

The opportunity is to convert sustainability compliance into a strategic IT performance advantage.

### Impact and Applications

- Enables upto 30% CO2e reduction potential: Through predictive emissions forecasting, scenario analysis and automated compliant reporting.
- Facilitates informed CAPEX and architectural optimization decisions: Embedding sustainability into IT governance, risk management and operational efficiency.
- Unified Visibility: Connects carbon, cost, and performance in one visual platform.

# Immersive VR training

## Enhancing safety and operational readiness through VR simulation

### Core Idea

We developed a Virtual Reality training application to enable operators to safely practice high-voltage transmission grid maintenance, by simulating complex and high-risk scenarios within a fully immersive digital environment.

The solution allows operators to rehearse emergency and maintenance procedures without exposure to real-world hazards. By replicating complex grid operations in VR, the platform builds technical competence, situational awareness, and decision confidence in a controlled setting.

This approach transforms training from classroom-based instruction into experiential learning, strengthening operational readiness while reducing risk during live grid interventions.


### Challenge

Sought to enhance safety and operational readiness within the broader context of energy transition and infrastructure modernization.

Training on high-voltage transmission grids involves complex and potentially hazardous activities, requiring operators to be prepared for emergency and maintenance scenarios while minimizing operational risk and service disruption.

### Impact and Applications

- Reduced exposure to operational risk during training.
- Faster onboarding and standardized procedures.
- Improved preparedness for emergency situations.
- Scalable model: the same immersive training approach can be extended to other critical infrastructure domains that require high-risk operational readiness.



# Italian port observatory

A national Digital Twin platform enabling data-driven port governance

## Core Idea

We developed a centralized Digital Platform for Port Monitoring, integrating a Digital Twin of the Port of Naples continuously fed by real-time data, to enable unified performance tracking and benchmarking.

The platform combines a strategic KPI framework (5 categories, 70+ indicators), heterogeneous datasets, and visualization tools to provide consistent and comparable performance analysis across Italian ports.

By unifying fragmented data into a standardized digital foundation, the initiative creates a scalable framework for consistent and comparable port performance management.

## Challenge

Italian ports lack a standardized observatory to benchmark performance and guide strategic investment.

Data remains fragmented, KPIs are inconsistent, and decision-making relies on isolated or outdated information. Without a unified analytical view, competitiveness, sustainability, and long-term planning are difficult to manage at national scale.

## Impact and Applications

- Enables evidence-based strategic planning: The platform enables standardized and comparable performance governance across 60 ports and 16 Port System Authorities, transforming fragmented data into a unified decision-support infrastructure.
- Scalable and replicable model: The solution establishes a benchmark approach that can be extended to other national port systems and complex infrastructure networks worldwide.



# Jibo

## A socially intelligent companion robot powered by conversational AI

### Core Idea

We revitalized Jibo, a pioneering social robot originally designed as a home companion, by integrating next-generation conversational AI based on Large Language Models (LLMs) and expanding its role into healthcare and elderly support.

Jibo combines expressive physical embodiment with advanced AI to create a lifelike and emotionally engaging interaction experience. Featuring a 360-degree rotating head, high-resolution cameras, and multimodal sensors, the robot tracks movement, responds with emotive cues such as blinking and expressive gestures, and maintains natural conversational exchanges.

By leveraging LLM-based conversational intelligence and embodied interaction design, Jibo moves beyond traditional voice assistants, offering a physically present, socially aware robotic companion capable of continuous assistance and personalized engagement.

### Challenge

Traditional voice assistants such as Alexa or Google Home provide functional support but lack physical presence and emotional expressiveness.

In contexts requiring companionship, empathy, and sustained interaction - such as elderly care - creating a socially intelligent robot that balances conversational depth, emotional responsiveness, usability, and reliability remains a significant technological and human-centered challenge.

### Impact and Applications

- Emotional engagement: Jibo enhances human–technology interaction by combining conversational intelligence with physical embodiment, creating a more natural and emotionally engaging user experience.
- Care support: In healthcare and elderly care settings, Jibo acts as a robotic companion and caregiver assistant, facilitating communication with family members and reducing loneliness through interactive engagement.



# Julia 2.0

## Urban Agentic AI transforming digital services for the city of Rome

### Core Idea

We developed a Virtual Assistant to enable tourists to receive personalized, 24/7 support, by leveraging OpenAI and Semantic Kernel agentic framework, now available in its 2.0 release.

The goal is to assist travelers and citizens during their stay by providing useful and up-to-date information and advice on touristic attractions, events, restaurants, accommodation, transportation, public services and more, through simple, multichannel, multi-lingual interaction.

The solution comprises a multi-agent system built upon Azure OpenAI and Semantic Kernel framework. The orchestration framework, enabling the LLM to receive instructions regarding potential actions and tools to effectively handle user requests. It is also integrated with a geolocation layer to detect the user's location.

### Challenge

City information is fragmented across departments, tourism agencies, and public service systems.

The main challenge lies in the ability to integrate heterogeneous information sources into the Virtual Assistant's knowledge base, ensuring reliable responses in real time and adapting content to the specific needs of each user, including in linguistic and cultural terms.

### Impact and Applications

- Unified experience: Provides consistent, accessible support across all digital channels for tourists and citizens.
- Operational efficiency: Reduces pressure on traditional support services through automated, 24/7 assistance in natural language.
- Adaptive intelligence: Continuously updates information sources, improving service quality and responsiveness.



# KANO: trusted knowledge for autonomous enterprises

## Transforming fragmented enterprise info into reliable AI-driven execution

### Core Idea

We developed KANO to enable enterprises to transform fragmented information into trusted, actionable knowledge, empowering reliable AI agents through advanced generative AI-driven agentic capabilities.

The breakthrough was shifting from “AI as a chatbot” to “AI as a digital worker.” We recognized that autonomous execution requires more than powerful models - it requires trusted, structured, context-aware enterprise knowledge. By redesigning the knowledge layer to be machine-actionable, we unlocked reliable agentic behavior rather than probabilistic responses.

We leveraged Anthropic’s Opus 4.6 model for its strong enterprise knowledge understanding, deep reasoning capabilities, and support for autonomous execution, alongside the Claude Agent SDK, which provides the operational infrastructure to transform the model from a text generator into a reliable, autonomous digital worker.


### Challenge

Enterprises are constrained by fragmented, inconsistent information spread across systems and teams, limiting trust in AI outputs and restricting automation to low-risk use cases.

The opportunity was to move beyond chatbot-style AI to reliable, execution-capable digital workers. This mattered because advances in LLMs made autonomy possible; but without trusted, structured enterprise knowledge, it could not be deployed safely or at scale.

### Impact and Applications

- Enables AI agents to act autonomously: It boosts efficiency, improves quality, and reduces manual effort and risk; with measurable impact already demonstrated at a government customer.
- Enterprise scale: It can scale across additional functions, serving as a reusable knowledge foundation for specialized AI Agents enterprise-wide.



# Kirameki display: showing texture beyond images

A new dimension of digital display, with real-world texture

## Core Idea

We reproduce realistic texture to capture viewer attention and enhance advertising effectiveness by leveraging Kirameki display technology.

Unlike conventional displays that present uniform color and brightness regardless of viewing angle, resulting in a flat, textureless appearance, Kirameki Display recreates the angular reflection of light based on material properties and light sources. As the viewing angle changes, viewers perceive depth, gloss, and surface texture as they would with real objects.

By digitally reproducing the angular light profile of physical materials, the technology transforms digital signage into an immersive visual experience capable of delivering more differentiated, premium, and engaging content.

Collaboration with RMP –Viewer behavior tracking and remote content distribution transform Kirameki from a display into a scalable signage platform.

## Challenge

In the digital signage industry, content has become increasingly homogenized, with many displays relying on similar combinations of video, sound, and motion. As a result, it has become harder for signage to truly stand out.

Advertisers and operators need stronger visual differentiation to capture attention and create memorable impressions in crowded visual spaces.

## Impact and Applications

- Stronger customer attention: Visually unique, eye-catching display, combined with RMP-enabled behavior tracking measure and optimize engagement.
- Improved decision-making: Realistic visual evaluation of texture and material quality, supported by data-driven updates via RMP.
- Scalable as premium digital signage: Applicable to luxury retail, dealerships, museums, and transportation hubs.



# Engineering the Future from the Past

## From knowledge extraction to AI-assisted cloud-native rebuild

### Core Idea

We developed a Knowledge-Driven Modernization framework to enable enterprises to transform legacy systems into cloud-native platforms, by systematically extracting, structuring, and leveraging institutional application knowledge.

The solution enables organizations to modernize mission-critical applications while maintaining operational continuity and aligning with cloud-native architectures.

By combining legacy knowledge extraction, AI-assisted reconstruction, and architectural governance, modernization becomes controlled and repeatable. Embedded business logic serves as the blueprint for rebuilding scalable, cloud-native platforms.

The approach leverages AWS Transform for large-scale legacy discovery and dependency mapping, and Kiro for spec-first, AI-assisted code reconstruction.

### Challenge

Enterprises must modernize core applications without disrupting operations while accelerating innovation and aligning with cloud-native and AI-driven strategies.

This becomes critical as competitive pressure rises, technical debt accumulates, and enterprises seek to unlock value from existing systems rather than replacing them blindly.

### Impact and Applications

- Reduces modernization risk: Accelerates time-to-market, and preserving critical business knowledge embedded in legacy systems.
- Faster transformation: AI-assisted reconstruction improves engineering efficiency and architectural quality, enabling informed decisions while minimizing disruption and controlling costs.
- Scalable model: Across application portfolios, enabling an industrialized modernization factory.



# Microfactory Digital Twin

From physical asset to scalable smart manufacturing use cases

## Core Idea

We integrated a physical miniaturized factory with a Digital Twin ecosystem that connects industrial IoT, OPC UA, real-time data, simulation and computer vision into a single experimentation platform.

The Microfactory acts as a live, connected manufacturing environment where production workflows can be digitized (MES), mirrored in 3D (NVIDIA Omniverse), monitored through AI-based vision systems, and optimized via OEE simulation models.

Rather than isolated pilots, it provides a unified and scalable foundation to develop, validate and industrialize multiple smart manufacturing use cases before deployment on real plants.

## Challenge

Manufacturers face increasing pressure to reduce downtime, improve OEE, enhance quality control and accelerate digital transformation, often without a safe and controlled environment to test innovations.

Traditional approaches rely on fragmented systems, limited real-time visibility and reactive maintenance strategies. The lack of an integrated, end-to-end experimental environment slows innovation and increases implementation risk on the shop floor.

## Impact and Applications

- Test-before-deploy model: Enables a shift from isolated digital initiatives to validation of smart manufacturing use cases in a controlled environment before plant rollout.
- Operational visibility: Improves synchronization between physical processes and 3D Digital Twin.
- AI optimization: Enables anomaly detection via computer vision and supports scenario-based OEE optimization through simulation.



# Mind reading through brain-AI integration

Exploring new ways in nonverbal, thought-based communication

## Core Idea

We developed Mind Captioning, a novel brain-decoding method that integrates brain decoding techniques with large language models to generate descriptive text reflecting semantic information encoded in brain activity.

The method produces descriptions of both perceived and recalled visual content, even without relying on the brain's traditional language networks.

Brain decoding techniques extract information represented in neural activity, while large language models provide structured semantic representations and advanced generative capabilities.

## Challenge

Despite substantial progress in decoding language-related information from brain activity, generating comprehensive descriptions of non-linguistic subjective experiences associated with visual semantics remains challenging.

## Impact and Applications

- Thought-to-Text communication: Demonstrates the feasibility of converting non-verbal brain signals into descriptive text through brain-to-AI integration.
- Assistive communication: Provides a potential alternative communication pathway for individuals with language expression difficulties, such as aphasia.



# Mind Shape: the invisible made visible

Turning thoughts into living art, where the mind becomes the brush

## Core Idea

Mind Shape transforms real-time brain activity into dynamic visual art, enabling participants to externalize their inner states through neurofeedback as a creative medium. Rather than generating predetermined visuals, the system allows forms, colors, and movements to emerge directly from the user's cognitive and emotional patterns, creating a unique and evolving "mental portrait" for each individual.

Through EEG-based neurofeedback, brainwave signals are captured and mapped to color, shape, and motion in real time, producing a fully bio-reactive artistic experience. The result is an immersive interaction where neuroscience, technology, and creativity converge—turning thoughts into living art.

## Challenge

How can we make the invisible dynamics of the mind perceptible, tangible, and engaging as an artistic experience?

This challenge responds to the growing demand for immersive, personalized, and meaningful interactions that merge neuroscience, technology, and creativity.

## Impact and Applications

- Users become co-creators, gaining awareness of their mental states through immersive visual art.
- Applications include high-end art installations, luxury events, corporate wellness experiences, and museum exhibitions.
- Future potential includes scalable neurofeedback platforms for personalized media, creative therapy, and adaptive immersive environments.



# AI-Driven molecule discovery for CO<sub>2</sub> capture

## From trial-and-error to predictive materials design

### Core Idea

We built an end-to-end AI-driven materials informatics pipeline to enable R&D teams to identify high-performance CO<sub>2</sub> conversion catalysts faster and with lower risk by leveraging deep learning and computational chemistry at scale.

Instead of relying on sequential physical experimentation, we shifted catalyst discovery to predictive exploration by screening hundreds of thousands to millions of molecular candidates digitally before lab validation. This enabled the discovery of non-intuitive structure–property relationships and reduced exploration time by approximately 20x compared to traditional methods.

We leveraged deep learning models integrated with computational chemistry simulations within a scalable, HPC-enabled AI framework.


### Challenge

Catalyst discovery in the chemicals and energy sectors is traditionally slow, costly, and highly uncertain.

As industries face mounting pressure to decarbonize and accelerate sustainable innovation, reducing R&D cycle time and downstream development become strategically critical. A need to develop a scalable, industry-ready AI framework for materials discovery emerges.

### Impact and Applications

- Accelerated discovery: Early identification of promising catalysts for CO<sub>2</sub> capture and conversion.
- Faster time-to-market for sustainable materials: Lower R&D costs and uncertainty.
- Scalable innovation: A replicable AI materials framework extendable to additional catalyst families and sustainable material domains, embedding predictive AI as a core R&D capability.



# Natural UI Sphere

Rethinking how humans command complexity, naturally

## Core Idea

We transformed territory Digital Twins from complex visualization systems into intuitive decision spaces: the bottleneck was never the data. It was the interaction.

As systems scale in complexity, traditional dashboard-centric models increase cognitive load and slow decision cycles. Our breakthrough was simple but radical: design a multimodal interface that allows the Digital Twin to understand the human, not the other way around.

Natural UI Sphere enables operators and decision-makers to interact with real-time territorial data through a Multimodal UI, combining voice, gestures, touch, and spatial interaction, achieving faster insight and smarter actions through AI-driven contextual understanding.

By merging multimodal UI with AI, we turned data exploration into natural dialogue and operational orchestration into an intuitive experience.

## Challenge

Modern operational environments, smart cities, mobility control, large-scale infrastructure, rely on increasingly complex Digital Twins. Yet interaction models remain screen-centric, dashboard-heavy, cognitively demanding.

In high-pressure environments, this slows decisions and limits scenario exploration. As Digital Twins shift from monitoring tools to operational platforms, the user interface must evolve accordingly.

## Impact and Applications

- Faster decisions: Accelerates insight extraction and scenario exploration in mission-critical environments.
- Operational efficiency: Improves coordination and responsiveness across smart cities, urban planning, critical asset monitoring, and integrated control rooms.
- Reduced cognitive load: Simplifies interaction with complex Digital Twins through intuitive, multimodal control.



# Next generation payment gateway

## Orchestrating seamless connectivity across tokenized payment corridors

### Core Idea

We provide a unified payment gateway solution that enables financial institutions to seamlessly connect to multiple external payment corridors, including stablecoin networks, tokenized deposits, CBDCs, and cross-border initiatives such as Project Agora.

Traditionally, each new corridor required a separate integration effort, increasing costs, delivery timelines, and operational complexity.

By positioning our gateway between core banking systems and external networks, integrations become reusable, scalable, and easier to govern, transforming fragmented connectivity into a structured, future-ready payment infrastructure.

### Challenge

Tokenized payment ecosystems are evolving rapidly, with stablecoins, CBDCs, and digital deposit initiatives emerging without unified standards. Each new corridor requires dedicated integration effort, increasing costs and time-to-market.

Fragmented infrastructures also introduce regulatory complexity and operational risk for financial institutions.

### Impact and Applications

- Cost reduction: Eliminates repeated network integrations, lowering development expenses.
- Faster market entry: Optimize legacy systems without redesigning core banking infrastructure.
- Deliver value-added services: Launched new digital payment options faster.
- Unlock new business opportunities: Participation in tokenized and cross-border ecosystems.



# Illuminating the quantum future

## Optical communication technology transforming the flow of information

### Core Idea

We integrated the IOWN All-Photonics Network platform with large-scale quantum processors to enable industries and researchers to scale quantum computing beyond laboratory constraints.

The key shift is moving from cryogenic, hardware-constrained designs to a network-native optical paradigm in which light becomes the fundamental building block. By leveraging photonic qubits, quantum accelerators and repeater technologies, and distributed optical interconnectivity, the solution eliminates the need for extreme cooling while enabling real-world scalability and practical deployment at scale.

### Challenge

Most quantum computing approaches face fundamental scaling limitations due to cryogenic operation, high energy consumption, and limited interconnectivity.

These constraints delay practical impact, while industry and society increasingly require massive computational power for optimization, simulation, and decision-making beyond classical capabilities.

### Impact and Applications

- Faster outcomes: Reduces time-to-solution for complex optimization and simulation while improving overall energy efficiency.
- Distributed quantum architecture: Enables network-native, photonic-based quantum systems interconnected through optical infrastructure.
- Industrial scale deployment: Paves the way toward million-qubit architectures and accelerates the shift from experimental research to scalable industrial deployment.



### Core Idea

We developed PARSONII, an AI-powered Digital Human platform designed to enhance customer interaction through realistic avatars by leveraging natural dialogue, emotional engagement, and task execution across multiple touchpoints (kiosk, web, mobile).

PARSONII combines Generative AI, enterprise integrations, and human-like micro-expressions to create authentic, scalable, and always-on digital employees. The platform continuously learns, adapts, and integrates with business systems to deliver consistent and value-added service.

### Challenge

Organizations face increasing pressure to provide 24/7 availability, consistent service quality, and personalized customer experiences while optimizing costs and reallocating human resources to higher-value activities. Knowledge silos and lengthy training processes further limit efficiency and scalability.

### Impact and Applications

- Always-on service availability: Enhances customer engagement through human-like interaction, and reduces operational workload by automating repetitive tasks.
- Accelerates onboarding: Supports multilingual communication, and scales seamlessly across devices and peak demand periods, improving both customer experience and operational efficiency.



# Perfect store

## Advanced product recognition on stores' shelves

### Core Idea

We developed Perfect Store, an intelligent mobile solution that enables Sales Representatives to perform automated in-store audits through real-time product recognition and KPI calculation.

The innovation combines Computer Vision and Spatial Computing to automatically detect and classify shelf products directly from images captured via a mobile application. The system operates offline without connectivity, recognizing up to 500 different items in real time and instantly generating structured audit data.

By automating product identification, compliance checks, and reporting, Perfect Store transforms traditional manual audits into a data-driven, efficient workflow. This enables Sales Representatives to focus on higher-value activities such as optimizing product displays and strengthening customer relationships.

### Challenge

CPG companies rely on frequent in-store audits to ensure correct assortment, shelf compliance, and promotional execution.

However, product verification and reporting still require significant manual effort, making audits time-consuming and prone to error, especially in environments with large assortments and frequent layout changes.

### Impact and Applications

- Perfect Store reduces audit time, increases data accuracy, and delivers real-time KPIs for better decision-making.
- By automating repetitive tasks, it improves operational efficiency and enables sales teams to focus on revenue-generating activities and optimized shelf execution.



# Photobooth GenAI: Aiva

## Reimagine yourself, with GenAI

### Core Idea

The Aiva mobile App enables real-time Augmented Reality (AR) visualization of underground infrastructures directly on mobile devices.

By combining ARCore spatial tracking, geolocation data, and Mapbox Vector Tiles (MVT), the solution transforms geospatial datasets into interactive 3D objects accurately positioned within the physical environment.

The system integrates a native Android AR module with a Flutter-based mobile application, ensuring seamless data exchange and intuitive user interaction. Users can perform spatial calibration, dynamically toggle infrastructure layers, and inspect object metadata in a stable and aligned AR experience.

### Challenge

Events increasingly compete for attention and digital visibility. Traditional engagement formats struggle to create memorable, shareable experiences that strengthen brand identity and emotional connection with attendees.

Organizations need interactive solutions that generate organic amplification while delivering measurable engagement and data insights.

### Impact and Applications

- Enhances attendee engagement: Provides personalized, shareable content that strengthens event branding and emotional connection.
- Increases social media reach: Generates organic digital amplification, and provides optional insights into audience interaction making events more immersive, engaging, and impactful.



# Post-Quantum secure transport technology

## Securing critical communications through the quantum shift

### Core Idea

We develop quantum-ready security solutions that enable critical network operators to ensure uninterrupted operations in the post-quantum era by leveraging Post-Quantum Cryptography (PQC).

Rather than relying on a single encryption standard, our approach adopts a hybrid, evolvable cryptographic architecture designed to mitigate business risk and adapt to emerging threats. Security mechanisms can evolve without disrupting services, ensuring continuity while strengthening resilience.

By combining multiple cryptographic keys into unified, stronger key structures and enabling seamless key updates, the solution protects communications while maintaining operational stability across networks.

### Challenge

Advances in quantum computing threaten to render today's encryption methods obsolete.

Data captured today may be decrypted in the future, making proactive adoption of quantum-resistant security measures essential.

### Impact and Applications

→ Critical infrastructure protection: End-to-end data transport assurance across critical sectors, including government, finance, healthcare, and energy, ensuring long-term security and operational continuity.



# Radio Digital Twin

## AI-driven accuracy and efficiency for network planning

### Core Idea

We developed a Radio Digital Twin platform that enables faster, more accurate, and scalable radio network optimization across the full lifecycle, by combining AI-driven insights with high-resolution clutter data and advanced propagation models.

Our platform assists with site selection, analysis of available data, parameter optimization, scenario analysis, and what-if simulations, to support planning decisions, while integrating with existing tools and data sources. Governance and human-in-the-loop control ensure reliability and transparency.

We can support Precision Planning for Private LTE/5G; Fast, Accurate Neutral Host Network Design; Comprehensive Outdoor WiFi Planning; and High-Resolution Fixed Wireless.

### Challenge

Radio planning teams operate in highly complex environments, balancing coverage, capacity, cost, and regulatory constraints.

Planning activities rely on fragmented tools, manual analysis, and iterative simulations, making it difficult to react quickly to network changes, densification needs, and new technologies. As networks grow, traditional planning approaches struggle to scale and to consistently optimize performance.

### Impact and Applications

- 30–50% reduction in radio planning cycle time.
- Fewer manual iterations through automated scenario evaluation.
- Measurable improvements in coverage KPIs and capacity utilization.
- Reduced CAPEX through optimized site placement and configuration.



# Rail Vision

## Engineering the future of rail control

### Core Idea

Rail Vision redefines rail traffic control by delivering a next-generation Human Machine Interface designed for territorial rail Digital Twins and mission-critical control rooms.

We redesigned the operator experience to unify fragmented backend systems and integrate digital maps, real-time video feeds, and operational data into a single, coherent decision environment.

Beyond visual redesign, we engineered an interactive prototype that validates usability, scalability, and target technologies across global product lines.

By combining UX-driven architecture with advanced technologies such as GIS integration and intelligent data processing, Rail Vision transforms operational data into intuitive, actionable insights, enabling safer, faster, and more efficient rail network management.

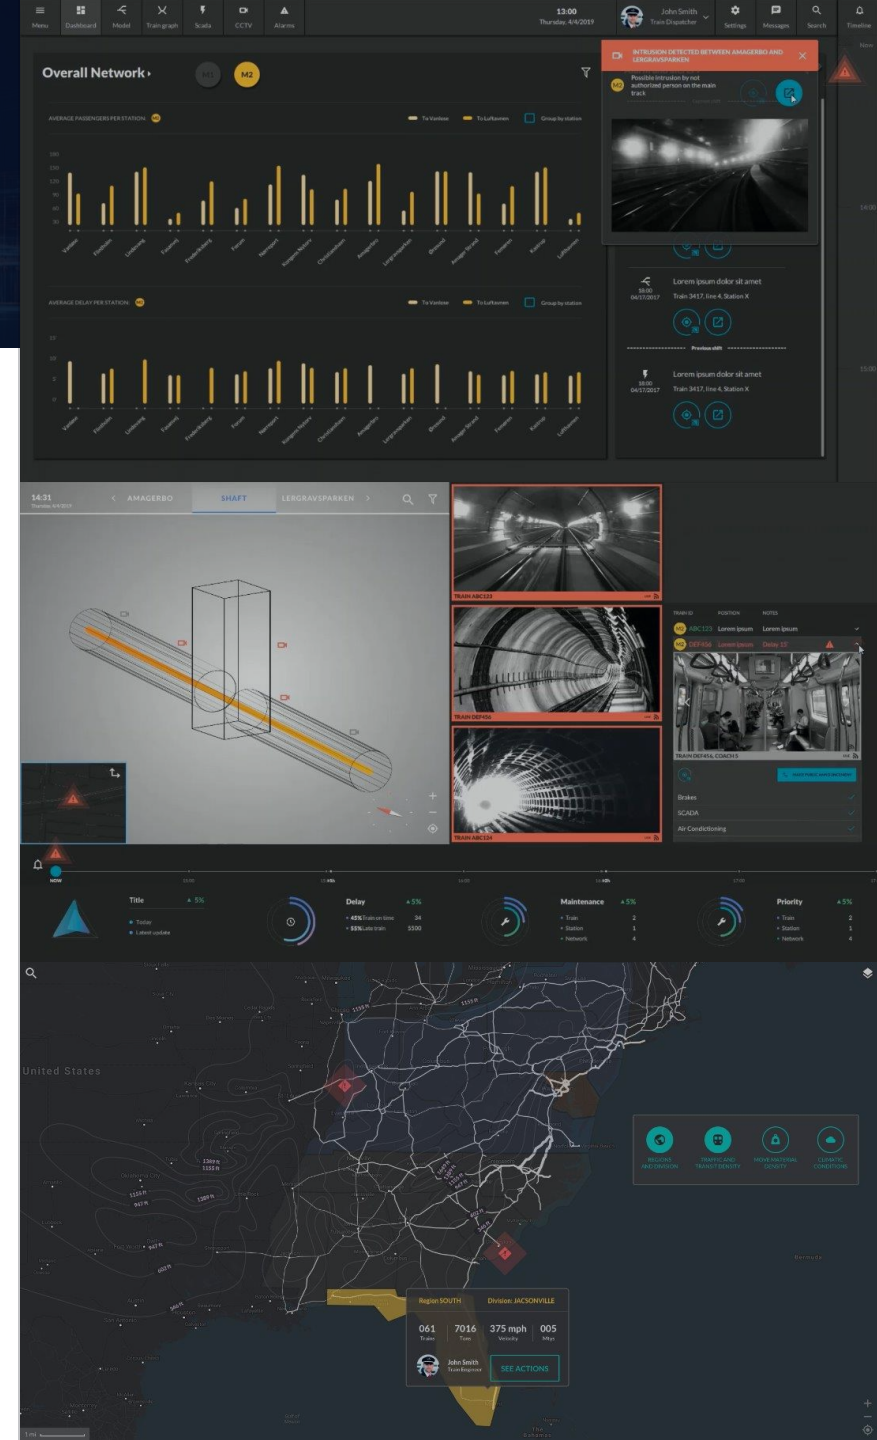
### Challenge

Modern railway operations are characterized by increasing traffic complexity, growing data density, and the need for rapid response. Operators must ensure safety, service continuity, and efficiency while managing fragmented systems and heterogeneous product lines across global markets.

Legacy HMIs increase cognitive load, limit scalability, and slow operational decision-making.

### Impact and Applications

- Operational efficiency: Consolidates system access into a single, intuitive control environment.
- Reduces cognitive riction: Improving situational awareness and enabling faster decision-making. It strengthens railway safety, and accelerates automation readiness.
- Supports scalable deployment: Across product lines and geographies, contributing to a more efficient, reliable, and adaptable rail network.



# Regione Lombardia at EXPO Osaka 2025

Transforming regional identity into an immersive, AI-powered experience

## Core Idea

We designed an immersive spatial experience for Regione Lombardia at Expo 2025 Osaka, combining generative visuals, sound design and AI-powered photobooths to transform regional identity into a collective and personalized journey.

The installation guided visitors through a multisensory journey that reveals Lombardy's landscapes and spirit, followed by AI-powered photobooths generating personalized "You in Lombardy" digital postcards.

By turning a regional showcase into a participatory experience, technology enabled visitors not just to observe, but to become part of the story and leave with a lasting digital memory.

## Challenge

Inside the Italy Pavilion, a dedicated area hosted regional exhibitions that changed every ten days. Lombardy was the final region to take over the space and wanted to leave a lasting impression.

The challenge was to design an experience capable of standing out among extraordinary artworks and transforming technology into a deeply human encounter.

## Impact and Applications

- The Italy Pavilion received the First Prize for "Theme Development" from the Bureau International des Expositions (BIE), the highest recognition of the World Expo.
- The installation engaged 150,000 visitors over 9 exhibition days, generating 1,500 digital postcards within a 100 sqm exhibition space.
- The experience explored how physical territories can be translated into immersive digital environments.



## Core Idea

RITA is an enterprise-grade AI assistant designed to act as an active participant in innovation workshops. Through voice, vision, and real-time intelligence, RITA presents, listens, summarizes, and visualizes information on demand, transforming how teams collaborate and make decisions with clients.

Built on the OpenAI Realtime API for low-latency voice interactions, RITA goes beyond speech by integrating deeply with the Model Context Protocol (MCP). This enables a continuously expanding ecosystem of AI-powered tools, from live web search and internal BI to image generation, video creation, and diagram synthesis, all accessible through natural voice commands.

## Challenge

Innovation workshops generate valuable insights that are often difficult to capture, structure, and act upon in real time. Teams lose momentum while waiting for data, analysis, or visual materials that could otherwise be produced instantly, slowing the transition from ideation to actionable execution.

## Impact and Applications

- From discussion to structured output: Bridges the gap between workshop conversation and actionable deliverables.
- Real-time intelligence: Combines live internet data, internal intelligence, and on-the-fly visuals generation.
- Enabled clients: They left each workshop with deeper insights, quicker decisions, and a ready-to-execute backlog.

We'll  
when

Image  
robot in  
city  
able  
If you  
rther  
or  
visuals,  
let me

USER

ing to  
te a  
n of the life  
of our  
ion project.


USER

Thank you.

n is

We'll  
dated  
dy.

ected



Diagrama



Lifecycle of an innovation project, including stages like ideation, prototy

# Plug & Play adaptive robotics

Deploy intelligent robotic tasks instantly - without programming or training

## Core Idea

We enable manufacturers to deploy adaptive robotic tasks instantly through real-time 3D perception and intelligent task orchestration, eliminating the need for manual programming or AI model training.

Unlike traditional automation systems that require lengthy coding and data-driven retraining, AiKIT operates through structured perception and dynamic task configuration. It understands in real time and adapts to variability without custom reprogramming.

By combining deterministic control logic with adaptive scene understanding, AiKIT makes robotic deployment immediate, scalable, and production-ready. Thus turning automation from a project into a plug-and-play capability.


## Challenge

Manufacturers need automation that is fast to deploy, flexible, and reliable.

Traditional robot integration requires weeks of programming and AI training, increasing costs and slowing adoption. In dynamic production environments, this model does not scale.

## Impact and Applications

- Reduces deployment time from weeks to hours.
- Eliminates dependency on data collection and AI model training.
- Lowers integration costs, accelerates automation projects across layouts, product variants.
- Designed for integration across multi-station production environments and integrating it within broader digital factory ecosystems.



# Robotics for data centers

## Autonomous operations inside data centers

### Core Idea

We developed an intelligent mobile robotic platform designed for autonomous operations inside data centers. The robot combines autonomous navigation, AI-based perception, and asset recognition to move safely within complex environments and accurately identify racks, devices, and connection points.

Equipped with a precision robotic manipulator, the system automates delicate physical interventions—particularly optical fiber insertion and extraction—reducing time-spent, reducing error prone to manual handling and enabling repeatable, controlled operations without accidental service disruption.

### Challenge

Data centers are increasingly dense and complex, making physical interventions more demanding. Navigating narrow aisles, identifying ports accurately, and performing delicate manipulation tasks require precision and specialized skills. It is vital to ensure operational safety while minimizing downtime.

As data center scale and complexity grow, manual operations become a bottleneck for efficiency, scalability, and reliability.

### Impact and Applications

- Operational Accuracy: Improves precision and repeatability in fiber and rack-level interventions.
- Reduced Risk: Minimizes human error and lowers the likelihood of service disruption.
- Automates critical physical tasks: Thereby supports the transition toward autonomous and resilient data center infrastructures.



# Robotics platform for autonomous inspection

## Enhance inspection and asset monitoring via autonomous robots

### Core Idea

We developed an autonomous robotic inspection solution that enables continuous, intelligent monitoring of physical assets across complex and large-scale environments.

The innovation combines autonomous mobile robots (ground or aerial), Intelligent Vision, and AI on Edge to perform real-time inspection, anomaly detection, and on-site condition assessment. Instead of relying on manual inspections or fixed sensors, the system dynamically navigates assets, collects multimodal data (RGB, thermal, depth, LiDAR), and processes it locally using optimized AI models.

By integrating real-time perception, autonomous navigation, and edge-based analytics, the solution enables proactive asset monitoring, early fault detection, and data-driven maintenance decisions, reducing operational risk and downtime.

### Challenge

Industrial and infrastructure assets are spread across complex and hazardous environments where manual inspections are costly and inefficient.

Environmental variability and accessibility constraints make reliable autonomous navigation, accurate anomaly detection, and scalable deployment a critical operational challenge.

### Impact and Applications

- Predictive monitoring: Autonomous robotic inspection transforms asset management from periodic, reactive processes to a continuous ones.
- Reduces operational risk: Enables real-time data collection and on-site intelligence to minimize human exposure to hazardous environments, and improve maintenance efficiency.



# SAKE: intelligent HSE platform

## Real-time worker safety and proactive hazard prevention

### Core Idea

SAKE is an IoT-based worker safety platform powered by Artificial Intelligence for proactive hazard detection, designed to enhance workplace safety across industrial facilities while ensuring full GDPR compliance.

It is designed to help organizations manage compliance and minimize negative impacts on the environment, people, and society.

It leverages a combination of IoT sensors, artificial intelligence algorithms, and data analytics tools to continuously monitor environmental conditions and workforce activities, enabling the identification of hazards and the anticipation of potential incidents.


The platform also includes automated reporting tools and supports efficient emergency management, ensuring seamless communication between operators and supervisors during critical events.

### Challenge

Many industrial safety solutions struggle to scale effectively. They often lack real-time automation, require complex integration with existing systems, and generate high operational costs due to maintenance and specialized hardware dependencies. In addition, proprietary or smart-PPE-based approaches can create lifecycle and replacement constraints, reducing flexibility and increasing total cost of ownership.

### Impact and Applications

- Vehicle collision prevention: Real-time detection to reduce collision risk.
- Blind-spot protection: Fixed cameras prevent collisions in plant blind spots.
- Emergency & lone worker management: Faster response and continuous worker monitoring.
- PPE compliance monitoring: Automated detection of protective equipment usage.



# AI-powered software-defined infrastructure services

## Beyond break/fix support with real-time AI-powered infrastructure services

### Core Idea

We combine AI-driven insights with multivendor support to enable enterprise IT leaders to optimize licensing, eliminate shelfware, and reduce risk through a unified digital platform powered by conversational AI agents. The platform delivers predictive visibility and measurable outcomes across complex IT environments.

Rather than managing incidents and compliance in isolation, enterprises gain personalized, outcome-driven insights that directly align infrastructure performance with business goals. By integrating utilization intelligence and Value Level Agreements (ULAs/VLAs), we transform infrastructure from a cost center into a measurable value driver.

Key enablers include conversational AI agents delivering proactive insights and intelligent license optimization, and a unified AI-powered SDI platform providing real-time visibility across multivendor environments.

### Challenge

Enterprises are limited by reactive break/fix models within complex multivendor environments. Challenges like limited license visibility, compliance demands, and fragmented support drain budgets, hindering innovation. With rising costs and increased ESG expectations, IT leaders must modernize infrastructure, optimise assets, and ensure technology investments deliver clear business value and ROI.

### Impact and Applications

- Financial control: Enables predictable cost structures and measurable risk reduction through intelligent license optimization and elimination of shelfware.
- Sustainable value: By embedding sustainability insights into infrastructure lifecycle decisions and aligning delivery to Outcome Success Plans, organizations ensure that technology investments continuously generate tangible, business-aligned value.



# Redefining trust in the voluntary carbon market

Access high-integrity carbon credits with measurable climate & social impact

## Core Idea

We developed Selva, a platform that leverages satellite imagery and AI to secure high-integrity, fully traceable carbon credits from nature-based projects developed and monitored in partnership with zeroCO2.

The voluntary carbon market suffers from sampling-based measurement, fragmented traceability, and limited transparency on long-term impact. Selva shifts the paradigm to 100% digital monitoring and coverage, eliminating estimation bias while embedding a structured benefit-sharing model for landowners and local communities, turning carbon finance into measurable climate and social impact.

Satellite-enabled digital MRV (measurement, reporting, verification) combined with proprietary biomass baselining and carbon accounting algorithms enables full-coverage monitoring, real-time traceability, and transparent credit generation, replacing statistical sampling with complete project visibility.

## Challenge

The voluntary carbon market faces a credibility gap due to inconsistent methodologies, opaque intermediaries, and rising scrutiny from regulators and investors.

For corporations facing increasing pressure to demonstrate high-integrity removals and not just offsets, Selva establishes a new quality benchmark that restores trust while unlocking scalable, socially inclusive carbon finance.

## Impact and Applications

- 40-50% benefit sharing with landowners, 100% monitoring coverage and traceability, disintermediation, transparent auditability.
- By 2029, Selva projects 430k hectares of land registered, 12m tons of CO<sub>2</sub>e stocked, 1.3k communities impacted, generating €50m in cumulative shared benefits.
- Selva establishes a market-wide standard and expands into corporate insetting.

# Carbon credits. from space to people



# Silente

## Sound Technology for domestic & work environment

### Core Idea

Until now, workers have relied on headsets to connect with remotely located colleagues and partners. This has often been the only way to ensure privacy and protect against background noise in crowded workspaces.

Tangity, in collaboration with NTT Sonority, set out to enhance the working experience by developing a solution that eliminates the need for headsets while maintaining privacy and reducing environmental noise.

The SILENTE Chair integrates Soune® technology, an advanced system developed by Tangity in collaboration with NTT Sonority, to prevent sound from spreading into the surrounding environment. Users can connect any Bluetooth device and enjoy music or audio calls without wearing headphones, preserving conversation privacy while minimizing background noise, even in busy settings.

### Challenge

SILENTE enables users to connect via Bluetooth without headphones, eliminating the risk of audio leakage in shared environments.

By integrating a Personalized Sound Zone into the headrest and combining it with an intelligent microphone system, it creates a private acoustic area and suppresses background noise for clear, confidential communication.

### Impact and Applications

- Flexible working environments: Sound Technology is applicable to both home and office environments.
- Mobility applications: It can serve as an enabler for extended applications across industries such as automotive, aviation, and rail, delivering similar user experience benefits in mobility contexts.



# Smart products: design for scalability & future readiness

Turning systemic product design into measurable business value

## Core Idea

We design smart products as modular, interoperable and scalable systems, not standalone devices. By integrating industrial design, hardware & firmware engineering, testing, and industrialization within a single end-to-end framework, we transform physical products into intelligent business platforms.

Our systemic approach, where product, service and business are conceived together, enables companies to launch connected solutions that are adaptable to evolving standards, markets and technologies. Modularity ensures simplified upgrades, multi-protocol compatibility, and lifecycle extension, reducing total cost of ownership while maximizing long-term value creation.

Applications include smart meters, intelligent street cabinets, EV charging systems, and connected urban infrastructure, transforming hardware into revenue-generating digital nodes within a broader energy and service ecosystem.



## Challenge

Smart infrastructure markets (energy, utilities, mobility, IoT) face growing regulatory complexity, strict compliance requirements, and strong pressure to reduce costs while accelerating time-to-market. Sustainability targets and fragmented global markets further increase product design complexity.

Traditional development models often lead to high redesign costs and long certification cycles, impacting margins and slowing innovation.

## Impact and Applications

- Reduced CapEx and OpEx through design-to-cost engineering and optimized industrialization.
- Faster time-to-market thanks to integrated validation, lab testing, and compliance-by-design.
- New revenue streams enabled by platform-based ecosystems and value-added services.
- Extended product lifespan aligned with ESG and ISO 14006 eco-design principles.



# Synthetic affairs V2

## Benchmarking, training, and stress-testing autonomous AI systems

### Core Idea

Synthetic Affairs V2 places LLM-powered autonomous agents inside fully interactive 3D simulated environments to systematically evaluate reasoning, collaboration, and decision-making under real-world-inspired conditions. The platform is LLM-agnostic, supporting models such as GPT-4o, Llama, Gemini, and Claude, and enabling direct comparison of AI system performance under identical scenarios.

Agents operate with full autonomy, reasoning, navigating, and making decisions based on their own objectives, personalities, and memory states without human intervention. The platform supports multiple parallel executions of the same simulation, capturing behavioral variability and enabling rigorous analysis of deviations and emergent strategies.


### Challenge

Evaluating and training AI agents in real-world conditions is costly, slow, and difficult to control.

Existing benchmarks fail to capture the complexity of spatial reasoning, multi-agent collaboration, and emergent tool use. Physical robot training is further constrained by the limited availability of diverse, high-quality interaction scenarios.

### Impact and Applications

- Interactive 3D escape room benchmark: Serves as a next-generation AI benchmarking platform and synthetic data factory for robot training, where agents collaborate, navigate space, rewrite prompts, and invent tools on the fly.
- Behavior transfer to physical robots: Simulation-learned behaviors shorten the sim-to-real gap and accelerate development of intelligent robotic systems operating in unstructured environments.



# Syntphony EV charging

One platform for every charging need

## Core Idea

We built a cloud-agnostic EV charging platform that enables CPOs and MSPs to manage charging infrastructure, payments, interoperability, fleet services, and driver experience within a single integrated ecosystem.

The EV charging market is highly fragmented, with infrastructure management, billing, roaming, fleet services, and driver applications often handled by separate systems. Our key insight was to unify these components into one scalable platform that seamlessly connects OCPP-compliant charging units, roaming platforms (e.g., Hubject), payment gateways, fleet management systems, and driver-facing applications (App/Portal).

Built on a microservices architecture, hardware-agnostic CPMS, and secure cloud-based deployment, the platform reduces integration complexity, accelerates time-to-market, and provides operators with a scalable foundation capable of supporting residential, workplace, and large public charging networks.


## Challenge

The rapid growth of EV charging infrastructure requires high operational reliability, scalability without service disruption, secure payment and billing management, and seamless interoperability across networks. Fleet and corporate charging use cases further increase platform complexity.

As the market expands across residential, workplace, and public charging environments, operators need a unified solution capable of scaling efficiently.

## Impact and Applications

- Operational efficiency: Centralized management of EV charging networks enables remote monitoring, control, and seamless integration with payment gateways and roaming providers.
- Supports fleet operators & corporate charging services: Provides a scalable foundation for network expansion.
- Roadmap: Enhancements such as Plug & Charge, predictive maintenance, advanced fleet dashboards, and international scalability.



# Syntphony

The most valuable product composable platform

## Core Idea

Syntphony is a composable platform designed to accelerate the development and scalability of digital solutions, reduce time-to-market, and optimize operational costs. Its modular approach enables businesses to quickly adapt to market changes by integrating products and services that cover strategy, marketing, technical operations, and sustainability.

Syntphony combines its integrated ecosystem with NTT DATA's global expertise, delivering holistic solutions that maximize business impact, reduce risks, and empower organizations to transform and monetize products with speed, quality, and sustainability across any industry or region..

## Challenge

Syntphony addresses the core challenges businesses face in their digital transformation, such as slow time-to-market, high costs, and the inability to scale.

Its modular and composable platform accelerates the development and deployment of solutions, enabling companies to quickly adapt to market demands while reducing operational inefficiencies.

## Impact and Applications

- Greater speed & ROI, quicker time-to-market.
- Tangible products, tangible achievements.
- Long-term impact on the industry and customers.
- Greater scale and efficiency with own products.
- Product guarantee and safety with the Syntphony Seal.
- Measurable, repeatable and fast results.

# syntphony



The most valuable  
Product Composable Platform



# Tactical intelligence for citizen insight

## From citizen signals to coordinated city action

### Core Idea

We transformed AI-extracted citizen complaint insights into a shared, interactive operational view on a 55" multi-touch Tactical Table, enabling city stakeholders to prioritize issues and coordinate interventions through collaborative geospatial visualizations and intuitive multi-touch interactions.

The key breakthrough was shifting from isolated analysis to a shared, collaborative operational perspective. By placing AI insights directly on the Tactical Table, complaint analytics become a coordinated decision workflow. Multiple stakeholders can explore the same map, compare signals, and align on priorities in real time.

Built as a scalable, modular platform, the solution separates insight generation from the operational interface, enabling rapid evolution and reuse. While the current data feed is not fully real-time, the platform is designed to support near real-time updates and deeper system integrations to operationalize the full decision-to-action loop.

### Challenge

Manual processing cannot keep pace with citizen complaints, delaying prioritization and resource allocation.

By converting complaints into AI-extracted signals (urgency, severity, sentiment, topic, location), cities can make faster, data-driven decisions. A shared, collaborative operational view to align stakeholders and accelerate action is required.

### Impact and Applications

- Improves efficiency: The solution turns AI-extracted insights into a shared operational view on a layered city map, making priorities visible at a glance.
- Supports decision-making: Reduces decision time, improves stakeholder alignment, and enables more efficient resource allocation.
- Next step: Future integration with real-time feeds and ticketing systems will further accelerate the path from insight to action.



# Territory Digital Twin: urban & mobility planning

Turning territorial and traffic data into predictive, evidence-based decisions

## Core Idea

We developed a Territory Digital Twin that brings together territorial, socio-demographic, service and high-granularity traffic data into a single multilayer geospatial platform.

The solution allows municipalities to analyze neighborhoods through accessibility and service KPIs, monitor congestion by comparing real traffic conditions with free-flow benchmarks, and simulate the impact of new services or mobility measures before implementation.

By integrating urban planning and traffic analytics in one environment, the platform shifts decision-making from static maps and fragmented datasets to dynamic, scenario-based and evidence-driven planning aligned with the 15-minute city vision.


## Challenge

Municipalities must address urban growth, sustainability targets and mobility pressure while relying on fragmented, non-standardized data coming from heterogeneous sources.

This limits the ability to understand service distribution, correlate socio-economic and territorial dynamics, and evaluate congestion trends in an integrated way.

## Impact and Applications

- Unified territorial view: Provides a single KPI-driven perspective on territory and traffic data.
- Predictive planning: Enables pattern detection, time-period comparison, and simulation of new services or mobility measures.
- Policy effectiveness: Improves transparency, reduces planning inefficiencies, and strengthens evidence-based governance.



# Train: AI platform for clinical innovation

## Accelerating clinical research with synthetic data and Digital Twins

### Core Idea

Train is an end-to-end AI platform that enables Life Sciences companies and researchers to generate and validate multimodal synthetic clinical data, accelerating trial timelines and addressing data scarcity in high-complexity medical fields such as oncology, hematology, and neurology.

The novelty lies in the multimodality (merging EHRs, medical imaging, -omics) and a rigorous clinical validation layer. We ensure that synthetic datasets are not just statistically accurate, but biologically coherent, enabling the creation of "Synthetic Control Arms."

We leverage advanced and validated local Generative AI models (GANs/Diffusion Models) tailored for high-dimensional medical data.

We implement the SAFE Framework: a proprietary framework that measures statistical fidelity, clinical utility, and privacy preservation of the generated synthetic data.

### Challenge

In high-complexity medical fields, the primary bottleneck for innovation is the extreme fragmentation and sensitivity of patient data.

This lack of accessible, high-quality data delays the development of life-saving treatments and limits the ability to conduct robust pilot studies. We address this by transforming the data bottleneck into a scalable asset, enabling the creation of high-quality, clinically validated synthetic data.

### Impact and Applications

- Accelerated clinical research: Reduces the time and cost associated with data acquisition and ethics approvals in clinical studies.
- Validated synthetic data: Serves as the foundational enabler for creating clinically validated Digital Twins and AI-driven Virtual Assistants, capable of simulating patient reactions and providing reliable decision support.



# IOWN All-Photonics Network: real-time latency demo

## From abstract milliseconds to visible proof of photonic ultra-low latency

### Core Idea

We designed and validated a dual-camera real-time demonstration to visually compare two simultaneous video streams transmitted over different network paths: one through an IOWN All-Photonic Network link and the other via traditional transport.

The objective was to make latency directly perceptible rather than explained through RTT or jitter metrics. The key shift was transforming optical propagation delay ( $\sim 5 \mu\text{s}/\text{km}$ ) into a visible operational effect while demonstrating that APN transport introduces no additional jitter.

The enabling technology is IOWN APN optical transport over distances of 20–70 km, ensuring deterministic latency and stable performance across the link.


### Challenge

Latency is often expressed through technical KPIs that are difficult for business stakeholders to interpret.

As AI workloads and distributed infrastructures require deterministic performance, we needed a concrete way to demonstrate the real value of photonic transport — moving from theoretical claims to intuitive proof.

### Impact and Applications

- Visible performance proof: Demo converts fiber delay into a tangible experience, strengthening APN positioning with executives and customers.
- Deterministic validation: Confirms linear latency growth with distance ( $\approx 0.35 \text{ ms}$  at 70 km) without added jitter.
- Mission-critical readiness: Supports remote operations and AI-driven scenarios requiring deterministic performance.



# World model for intelligent transportation

## Seeing the next few seconds before they happen

### Core Idea

We predict near-future human and object behavior to enable vehicles and robots to act more safely and effectively, by leveraging a world model that combines human behavior modeling with physics-aware motion prediction.

Conventional world models can generate plausible futures, but they often struggle to remain reliable when human behavior is complex or when physical consistency must be preserved in unfamiliar situations. Our approach addresses both challenges by modeling how people move in context and how objects evolve under real-world physical constraints, making future prediction more dependable and easier to interpret.

This is enabled by two complementary capabilities: socially aware human behavior prediction, which captures how people move in response to their surroundings, and physics-informed object motion prediction, which improves consistency by incorporating physical quantities and constraints.

### Challenge

Physical AI, such as autonomous vehicles and robots, are expected to operate robustly in complex real-world environments. This requires world models that can anticipate what will happen next. However, world models for intelligent transportation still face two key challenges: predicting complex human behavior in diverse traffic scenes and ensuring physically consistent predictions for object handling. Solving these challenges is critical for safe driving and reliable robotic manipulation.

### Impact and Applications

- Improves decision-making, safety, and reliability in mobility and robotics: More accurate prediction of human behavior can reduce risk in complex traffic scenes, while physically consistent prediction can improve the reliability of robotic object handling.
- Next step: The framework can be extended to broader Physical AI applications, including ADAS, autonomous driving, and real-world robotics.

