




INTERNET OF THINGS FOR AGENTS

IoT4Agent

IoT built for agents

Lindos OS spatial intelligence
Reference architecture and application

Lindos Technology · October 2026 · English edition



See ahead. Act ahead.

Trusted spatial facts, then decisions and actions.

CONTENTS

Inside this paper

Abstract 3

1. Agents enter physical sites: a new infrastructure gap 4

2. From displays for people to infrastructure for agents 5

3. Reference architecture: the five layers of Lindos OS 7

4. Key technologies 9

5. End-to-end example: an overheating incident in a business park 12

6. Typical use cases 17

7. Safety, trust and governance 19

8. Implementation path and open ecosystem 20

9. Conclusion: spatial intelligence, from insight to action 21

Appendix A. Glossary 22

Appendix B. Data and test conditions 23

OVERVIEW

Abstract

Agents are moving from conversation to action. In business parks, construction sites, transfer stations and laboratories, their first challenge is infrastructure: the physical site is not sufficiently visible, trustworthy or controllable. Earlier IoT systems were built for dashboards viewed by people. Data reached operators as charts, its meaning remained in human experience, commands travelled by phone and work order, and people confirmed the result.

This white paper proposes **IoT4Agent - IoT built for agents**. It asks physical sites to provide three capabilities: **readability**, with people, equipment, places and assets continuously connected through comprehensive, low-power sensing; **trustworthiness**, with a source, time, location and version attached to spatial facts; and **actionability**, with commands executed, acknowledged and verified through two complementary engines.

Lindos OS is a reference implementation of IoT4Agent. Its five layers - CELL, GRID, PORT, CORE and AXIS - support three propositions: Sensing as a Service, Source of Truth and Sense to Act. A 3D simulated overheating incident in a business park illustrates the full process, from sensing and assessment to dispatch and review.

We advocate starting with one use case: establish trusted spatial facts first, then involve agents in on-site decisions and actions. Open protocols, interfaces and data contracts can help partners build physical-world infrastructure that agents can use.

Keywords: IoT4Agent · agents · spatial intelligence · spatial facts · dual-engine architecture · action transactions · Lindos OS

CHAPTER 01

1. Agents enter physical sites: a new infrastructure gap

1.1 From conversation to action

Large language models enable software to understand goals, plan steps and use tools, allowing agents to undertake sustained work. In the digital world, they work with documents, databases and APIs. State can be queried at any time, and most operations can be reversed.

Physical sites work differently. Their objects are people, equipment, vehicles, materials and spaces. Conditions change continuously and enter a system only when sensed. Actions have real consequences, such as switching off power, restricting access or sending someone to a location. Before acting, an agent must reliably answer three questions: what is happening, where is it happening, and who or what is involved? It can then determine what to do, who should do it and how completion will be verified.

1.2 Three gaps between agents and the site

Table 1-1. Three gaps between agents and the site

Gap	What happens on site	Impact on agents
Not visible	People and mobile assets lack continuous location records. Equipment and environmental readings depend on manual transcription or occasional inspections. Short device battery life makes sustained coverage difficult.	Objects that cannot be seen cannot be considered in a decision.
Not trustworthy	Data is scattered across systems, with inconsistent object IDs and timestamps. A notification may lack a clear source, capture time, location or device version.	Assessments lack evidence, conclusions cannot be checked and responsibility is difficult to establish.
Not controllable	Instructions travel by phone, group chat or paper work order, with completion relying on manual feedback. Permissions and safety prerequisites exist only in policy documents.	Agents cannot confirm delivery, execution or effect.

In dashboard-centred systems, people bridge these gaps. They use experience to judge data quality, call the site to check conditions and compile records afterwards. Agents lack that buffer: their interfaces must already expose complete, traceable and actionable information.

1.3 The same problems recur across industries

Waste collection, equipment rental, warehouse handling and hazardous-chemical operations involve different objects but reveal very similar problems.

Table 1-2. Five recurring problems across industries

Problem	Typical situation	Missing spatial facts
Waiting	Transfer vehicles queue at stations; lorries wait for loading; production lines wait for materials or maintenance staff.	Vehicles, goods, loading bays and work plans lack a common timeline.
Utilisation	Construction machinery and mobile medical equipment can be idle in one place and scarce in another. Productive staff hours are unclear.	Whether people and equipment are present, in use or idle.
Authorisation	Who may enter a controlled area or carry out high-voltage or hot work?	A live link between identity, qualifications and location.
Responsibility	Who was present or away from their post during an incident, and which party caused a delay?	Continuous, replayable movement and event records.
Settlement	Outsourced services are paid by hours, trips or coverage; equipment is billed by productive operating hours.	A measurement basis and evidence chain accepted by both parties.

These gaps were previously treated mainly as efficiency problems, addressed through management and manual coordination. As agents capable of making decisions and executing actions enter physical sites, the problem changes. Agents cannot replace missing data with experience and phone calls. Sites need infrastructure designed for them.

1.4 Compliance requirements encourage evidence-based site records

Compliance requirements point in the same direction. In China, occupational safety standard AQ 3064.3-2025 sets construction, technical and data-exchange requirements for personnel positioning in hazardous-chemical enterprises and chemical parks. Personnel, key areas, work permits and unusual dwell times need continuous, traceable compliance evidence. Such evidence supports regulatory oversight and also provides the inputs agents need for on-site safety management.

CHAPTER 02

2. From displays for people to infrastructure for agents

2.1 Definition

IoT4Agent (Internet of Things for Agents) is IoT designed with agents as its primary users. It continuously turns people, equipment, places and assets into evidence-backed spatial facts, exposes them through machine-readable interfaces, and provides channels to execute, acknowledge and verify agent-issued commands.

People remain the site's decision-makers and accountable parties. IoT4Agent changes the division of work: agents collect data, organise facts, perform routine assessments and track tasks, while people focus on authorisation, exceptions and responsibility for results.

2.2 Traditional IoT and IoT4Agent compared

TRADITIONAL IoT

Site	Platform	Dashboard / report	Duty operator
Project-specific coverage	Readings and charts	Display data to people	Assess screens; dispatch by phone

Commands: phone and work order. Results: manual feedback.

IoT4Agent

Readable site	Trusted spatial facts	Agent	Action transaction
People, equipment, places and assets continuously connected	Objects, relationships and events with source and version	Real-time responses + reasoning	Command → acknowledgement → verification

Acknowledgements and verification update spatial facts. People retain authorisation, exceptions and accountability.

Figure 2-1. From displays for people to infrastructure for agents

Table 2-1. Traditional IoT and IoT4Agent compared

Dimension	Traditional IoT: for people	IoT4Agent: for agents
Primary users	Operators and managers	Agents and the people working with them
Coverage	Project-specific sensor placement, with key equipment connected	People, equipment, places and assets continuously connected through comprehensive sensing and long-running, low-power devices
Data form	Point readings and charts	Spatial facts organised as objects, relationships and event streams across time and space
Trust	Visibility is enough; people judge missing or delayed data	Every fact carries its source, time, location and version; gaps and conflicts are explicitly flagged
Meaning	Point names and device IDs; meaning resides in human experience	A unified device model with meaning that software can query
Interfaces	Dashboards, reports and push notifications	Standard interfaces for queries, subscriptions, commands and acknowledgements
Decisions and execution	People assess dashboards, assign work by phone or work order, and confirm results manually	Two engines share assessment work; commands carry permissions and results are checked against site data
Improvement	Data is archived; reviews require manual compilation	Events become experience and rules, allowing the system to learn with continued operation

2.3 Three capabilities: readable, trustworthy and actionable

- **Readable (Read).** Site objects remain continuously connected through comprehensive, low-power sensing. Agents can query any object's location and status at any time or subscribe to changes.
- **Trustworthy (Trust).** Spatial facts carry a source, time, location and version, organised by site objects, relationships and event streams across time and space. Each agent assessment can be traced to evidence. When a fact changes, affected assessments can also be identified.
- **Actionable (Act).** The real-time engine handles deterministic responses in milliseconds. The reasoning engine handles interpretation and trade-offs in minutes. Outcomes are verified against on-site data.

These capabilities depend on one another. Without readability there are no facts; without trustworthiness actions have no sound basis; without actionability sensing and assessment remain on a screen.

2.4 Design principles

1. **Facts before intelligence.** Establish trusted spatial facts before introducing models and agents. Models evolve while the factual foundation continues to accumulate.
2. **Evidence throughout.** Carry evidence from the original signal to the action's outcome, so every assessment and action can be replayed.
3. **Determinism first.** Safety-related, time-critical responses with clear criteria use deterministic rules, execute in milliseconds and do not rely on large language models. Models handle assessments that require trade-offs.
4. **Actions as transactions.** Check every command's permissions and prerequisites, then verify the result against on-site data.
5. **People in the loop.** Authorise actions by risk level and require human confirmation for high-risk actions. Record confirmations, rejections and corrections as evidence too.
6. **Open and replaceable.** Open protocols, interfaces and data contracts allow third-party devices and business systems to connect. Models, devices and deployment environments can be replaced.

CHAPTER 03

3. Reference architecture: the five layers of Lindos OS

3.1 Overview

Lindos OS is a reference implementation of IoT4Agent. From bottom to top, its layers are CELL sensing devices, GRID network backbone, PORT spatial facts interface, CORE AI computing and situational intelligence, and AXIS digital twin and real-time dispatch. They form four functional stages: sensing, facts, computation and action.

For people: one map, work orders and insights. **For agents:** spatial facts and dispatch commands.

Layer / stage	Main role	Core proposition
AXIS / Act	Digital twin and real-time dispatch; spatial interaction, insight console, multi-device use, APIs and BMS/ERP/MES integration	Sense to Act CORE + AXIS Actionable
CORE / Compute	Spatial semantics; real-time engine and reasoning engine	Sense to Act + Source of Truth
PORT / Facts	Protocol adaptation (Lindos, Modbus, MQTT, HTTP); standardisation; unified device model; device management	Source of Truth PORT + CORE Trustworthy
GRID / Sense	Indoor/outdoor/vehicle receivers; relays; multi-hop and self-healing topology; edge preprocessing; firmware updates	Sensing as a Service CELL + GRID Readable
CELL / Sense	Location sensors, environmental sensors and electronic-paper displays; in-house 2.4 GHz protocol	Continuous low-power sensing

Data ↑ Commands ↓ Third-party devices can report through Lindos and receive device-level commands; third-party sensing connects through PORT.

Figure 3-1. Lindos OS: five reference layers and three core propositions

The principle is simple: **first turn the site into trusted spatial facts, then involve AI in on-site decisions and actions. Lindos OS learns as it operates over time.**

3.2 What each layer provides to agents

CELL - Sensing devices. Location devices include badges, wristbands and tags with indoor beacon and outdoor BeiDou modes. Sensors cover temperature, humidity, pressure, gas, vibration, light and door status, using energy harvesting or long-life batteries. Display devices use electronic paper. All share an in-house 2.4 GHz proprietary protocol for unified device access and identity management: varied physical forms share a common protocol and data core. When state changes, the network backbone sends new content to update displays.

GRID - Network backbone. Indoor, outdoor and vehicle-mounted receivers and wireless relays support multi-hop relaying, self-healing topology and edge preprocessing. Receivers use Wi-Fi or Ethernet uplinks and DC or PoE power. They also deliver firmware over the air and support staged OTA updates, allowing large device fleets to evolve after deployment.

PORT - Spatial facts interface. Protocol adaptation connects Lindos devices and third-party sensors over Modbus, MQTT and HTTP. Data standardisation cleans, aligns and semantically maps data into a unified device model. Device management supports staged OTA and batch configuration. PORT is the standard entry point through which agents read the site.

CORE - AI computing and situational intelligence. The spatial semantics engine models physical spaces as zone topology and binds business meaning to time and location, creating operational event streams. The real-time engine uses deterministic rules to respond in milliseconds to geofence events, SOS calls and unusual dwell times, triggering vibration, sound, light and push notifications without a large language model. Harness orchestrates the reasoning engine's large language models to assess trends, anomalies and improvements in minutes, continually feeding insights back into the real-time rule library.

AXIS - Digital twin and real-time dispatch. The spatial interaction engine provides AI modelling, real-time 3D rendering, semantic scene annotations, multi-source spatial overlays and intelligent zone topology. The insight console provides AI insights, spatial focus on unusual events, workflow triggers and

coordinated response plans. Multi-device collaboration and open APIs cover desktop coordination, mobile inspection, IOC situational views and APIs/webhooks, with BMS, ERP and MES integration. Third-party devices can report data and receive device-level commands through the Lindos protocol.

Table 3-1. What the five layers provide to agents

Layer	Capability provided
CELL - Sensing devices	Object identity, location and status; outputs at the site edge through electronic-paper displays, vibration, sound and light
GRID - Network backbone	Reliable upstream data and downstream command channels
PORT - Spatial facts interface	Queryable, subscribable objects and facts in a unified device model
CORE - AI computing and situational intelligence	Operational event streams across time and space, deterministic responses, assessments and recommendations
AXIS - Digital twin and real-time dispatch	Dispatch commands, workflows, coordinated response plans and business-system integration

3.3 Mapping the three propositions

Each proposition maps to one IoT4Agent capability. CORE appears in the latter two because it both turns data into trusted facts and makes decisions from them. Facts and decisions therefore share the same evidence structure.

Table 3-2. Propositions mapped to IoT4Agent capabilities

Proposition	Layers	Capability	Meaning
Sensing as a Service / Continuous sensing	CELL + GRID	Readable	Deliver sensing as an ongoing service, showing connectivity by object and area. Protocols and networking support battery life and delivery.
Source of Truth / Trusted data foundation	PORT + CORE	Trustworthy	A unified device model provides a common factual source for agents and business systems.
Sense to Act / Dual-engine decisions and actions	CORE + AXIS	Actionable	Two engines assess the situation; commands reach people, equipment and systems, with acknowledgements and verification.

CHAPTER 04

4. Key technologies

4.1 In-house ultra-low-power 2.4 GHz proprietary protocol

Badges, asset tags and environmental sensors are numerous, dispersed and mostly battery powered. The burden of replacing batteries determines whether a system can be sustained. Agents must also send commands back to the site, making a badge vibrate in an emergency or refreshing an electronic-paper display. Battery life and timely downlink delivery must both be supported. That is the starting point for Lindos's protocol.

The same protocol stack supports millisecond delivery of emergency commands and over-the-air firmware updates. The first carries real-time responses to on-site devices; the second supports staged upgrades and continued improvement after large-scale deployment. The protocol stack is a core asset of

Lindos's hardware system. Workload data from real use cases informs the protocol and power strategy, which in turn guide module and hardware selection and configuration.

Table 4-1. Measured system-level power consumption

Metric	Value	Conditions
Sleep current	0.36 µA	Measured at system level on the complete device, including the IMU
Average current	7.72 µA	Measured at system level on the complete device, including the IMU
Peak current	3.91 mA	Instantaneous transmit/receive current in the same test
CR2032 coin-cell battery life	≥ 3.3 years [1]	Estimated from measured average current; actual life depends on environment, temperature and settings

Test subject: in-house 2.4 GHz proprietary-protocol firmware running on real hardware, with a complete device including the inertial measurement unit (IMU). System-level measurements taken on 27 May 2026.

[1] Theoretical estimate: nominal CR2032 capacity of 225 mAh divided by measured average current of 7.72 µA.

4.2 Spatial facts

A sensor reading says only that a point reported a value. An agent needs a spatial fact: which object, where, when, in what state and based on what evidence? These facts have three elements:

- **Objects:** people, equipment, assets, materials, spaces and loads, each with a unique identifier.
- **Relationships:** located in, supplies power to, belongs to, qualified for, responsible for and depends on. These connect objects into a network that supports reasoning.
- **Event streams across time and space:** entry, exit, dwell, limit exceedance and recovery, derived from raw signals by the spatial semantics engine and linked to business meaning.

Object	Relationship	Connected object
C substation	Supplies power to	B production line
Transformer C-3	Located in	C substation
Zhang	High-voltage qualification authorises work on	Transformer C-3
Robot #02	Infrared inspection of	Transformer C-3
Spare JT-120	Compatible replacement for	Transformer C-3

EVENT STREAM / DEMO TIMELINE

12 s	22 s	30 s
Temperature exceeds limit	Restrict access and evacuate	Open C-3
47 s	71 s	76 s
Hotspot: 112 °C	Retest: 41 °C	Close switch and lift restrictions

Figure 4-1. Spatial facts example from the core design, using the chapter 5 simulation

Spatial and business modelling. Spatial modelling organises physical space into a topology of parks, buildings, floors, zones, routes and geofences. Business modelling binds operational meaning to time and space. For example, a person leaving a work area is a spatial event. Combined with the facts that a

critical process is underway and that person is responsible for it, it becomes an operational event: the responsible person has left their post. The real-time engine can then notify and escalate directly.

4.3 Dual-engine architecture

On-site responses operate on two time scales. Entry into a hazardous area, SOS calls and unusual dwell times require immediate assessment and device responses. Equipment trends, competing objectives and explanations for anomalies require more information and conclusions on a minute scale. Lindos OS assigns these to two engines.

Division of work. The real-time engine applies deterministic cloud or edge rules to geofences, SOS calls, unusual dwell and threshold events. It does not depend on large language models and can trigger vibration, sound, light, push notifications and linked controls. Harness orchestrates the reasoning engine's large language models to assess trends, anomalies and improvements, producing anomaly explanations, dispatch recommendations, task orchestration and workforce reports.

Input: spatial semantics engine - operational event streams across time and space, plus past experience.

Real-time engine	Reasoning engine
Milliseconds. Deterministic rules. No LLM required. Geofences, SOS and dwell → rule matching → vibration, sound, light and push notifications. Cloud/edge rules handle safety-related, time-critical responses with clear criteria.	Minutes. Large language models orchestrated by Harness. Combine sources → assess trends, anomalies and improvements → explanations and recommendations. Four runnable outputs: anomaly explanations, dispatch recommendations, task orchestration and workforce reports.

Learning loop: people confirm findings, which become real-time rules. Lindos OS learns as it operates.

Execution: AXIS sends commands and coordinates response plans across people, equipment and systems. Reasoning recommendations execute after human confirmation or authorisation in a response plan.

Figure 4-2. Real-time responses and minute-scale reasoning

Reasoning becomes rules. Patterns found by the reasoning engine are confirmed by people and converted into deterministic rules in the real-time library. For example, past work orders and operating data may show that a batch of connectors overheats under high load. An early-warning rule can then target that batch, letting the real-time engine handle similar events directly. Capability grows with operating experience: **Lindos OS learns as it operates over time.**

Fallback and degraded operation. The real-time engine runs independently of large language models. If the model service is unavailable, responses such as geofences and SOS continue, preserving the baseline safety functions.

4.4 Open connectivity

IoT4Agent becomes more valuable as the range of connected objects grows. Lindos OS remains open towards both devices and applications.

The core design organises agent-facing interfaces around four operations: **query** objects, locations, status and relationships; **subscribe** to event streams; **issue** commands with permissions; and **receive**

acknowledgements and verification results. These operations are planned as agent tool calls so different frameworks can use them consistently.

Table 4-2. Open connectivity options

Direction	Method	Scope
Device side: Lindos protocol	Third-party devices connect to the Lindos protocol	Report data through the Lindos protocol and receive device-level commands
Device side: common protocols	Third-party sensors connect to PORT	Modbus, MQTT and HTTP data is standardised into the unified device model
Application side: agents	Spatial facts interface and dispatch commands	Spatial fact queries and dispatch commands; event subscriptions and acknowledgements are being developed under the core design
Application side: business systems	API / Webhook	Integrate BMS, ERP, MES and other systems, push events and write results back
Multi-device collaboration	Human interfaces	Desktop coordination, mobile inspection and IOC situational views

CHAPTER 05

5. End-to-end example: an overheating incident in a business park

Simulation notice. This chapter uses the Lindos OS 3D demonstration, On-site agent dispatch. It is a simulated scenario. The park layout, device IDs, temperatures, loads, personnel, distances, work orders and experience counts are all simulated data used to explain the IoT4Agent process and architecture. They are not operating results from a real project.

5.1 Setting and roles

The simulated park has five zones: A warehouse, B production workshop, C substation, D coordination centre and E logistics/loading, with 46 sensing channels. The **Lindos OS Coordination hub** is the lead agent. It continuously observes site-wide data, assesses spatial facts, plans tasks and issues commands; all insights and commands originate here. It orchestrates six specialist agents: **sensing** (aggregate data and detect anomalies), **safety** (assess risk, restrict access and evacuate), **inspection** (close-range checks by robots and aircraft), **personnel** (match qualified staff, navigate work orders and stream AOV body-camera video), **energy** (switching, backup power and load allocation), and **review** (summarise by site object and generate reports).

5.2 Seven stages

The incident begins at a transformer in the zone C substation. Figure 5-1 shows how the architecture layers participate in seven stages. Data flows upwards into CORE, assessment takes place in CORE, and commands travel through AXIS and GRID back down to the site. Table 5-1 gives the detail for each stage.

Stage / time	Layer participation and flow	Lead agent
Sensing / 12 s	CELL temperature → GRID multi-hop uplink → PORT unified model → CORE anomaly detection	Sensing
Assessment / 15.5 s	CORE assesses spatial facts	Coordination hub
Dispatch / 22 s	AXIS commands → GRID downlink → CELL access restriction and switching	Safety, inspection, personnel, energy
Inspection / 44 s	CELL infrared → GRID uplink → PORT thermal record → CORE root cause	Inspection
Work / 52 s	CELL AOV video → CORE compliance and retest → AXIS voice reminder	Personnel
Recovery / 76 s	AXIS recovery command → GRID downlink → CELL switching on and lifting restrictions	Energy, safety
Review / 88 s	CORE incident review → AXIS report; experience becomes rules	Review

Times follow the demonstration timeline. Data flows upwards; commands flow downwards. Assessment and decisions take place in CORE, with evidence connecting each stage.

Figure 5-1. Incident stages, architecture layers and agents (simulation)

Table 5-1. Agents, data, assessments, actions and layers by stage (simulation)

Sensing

Agent	Sensing agent
Data used	Time series from 46 channels; C-3 temperature 78 °C, 36 °C above its operating baseline
Assessment	Temperature is above the threshold and still rising: an anomaly
Action taken	Report event EVT-0928-01
Architecture layers	CELL, GRID → PORT → CORE real-time engine

Assessment

Agent	Coordination hub
Data used	Relationships: C-3 → zone C substation → power supply to zone B production line; 12 similar work orders
Assessment	High risk: continued heating may exceed the insulation limit in about 8 minutes; connector degradation suspected
Action taken	Split the response into 4 parallel tasks
Architecture layers	CORE reasoning engine

Dispatch

Agent	Safety, inspection, personnel and energy agents
Data used	Zone B load 1.2 MW; backup power 1.0 MW; badge locations and qualification records; robot route; smoke readings and aerial thermal images
Assessment	A 0.2 MW shortfall requires load reduction. Li, 28 m away, holds only a low-voltage qualification; assign Zhang, high-voltage qualified and 42 m away. A parked lorry requires a detour; smoke is present without open flame.
Action taken	Restrict zone C and evacuate 3 people; reduce zone D air-conditioning and lighting loads by 25%; send Zhang the work order; open C-3 and transfer to backup power; dispatch robot and aircraft
Architecture layers	CORE → AXIS → GRID → CELL

Inspection

Agent	Inspection agent
Data used	Robot #02 infrared images
Assessment	A 112 °C connector hotspot is traced to connector degradation; replacement part JT-120 is needed
Action taken	AGV-03 brings the spare part from zone A
Architecture layers	CELL → CORE → AXIS

Coordinated work

Agent	Personnel agent
Data used	Zhang's first-person AOV body-camera video and robot infrared images
Assessment	The earthing lead is not connected, so work cannot proceed. After the reminder, earthing, voltage testing and lockout/tagout are completed. Retest: 41 °C and secure fastening.
Action taken	Give a voice reminder to connect the earthing lead first; oversee the repair throughout
Architecture layers	CELL → CORE → AXIS

Recovery

Agent	Energy and safety agents
Data used	Retest results; personnel and load status in the affected zones
Assessment	Recovery conditions are met
Action taken	Close C-3 → restore zone D loads → lift access restrictions; robot and aircraft return, personnel resume work
Architecture layers	AXIS → GRID → CELL

Review

Agent	Review agent
Data used	Evidence from the whole incident and past experience
Assessment	Other connectors in the same batch may carry the same risk
Action taken	Generate a report; recommend checking 12 connectors from the same batch and including them in predictive maintenance
Architecture layers	CORE reasoning engine → rule library

5.3 Three conflicts and the evidence behind the decisions

The simulation includes three on-site conflicts. Together, they show that an agent's dispatch quality depends on the facts it can read.

- **Backup power shortfall.** Opening C-3 transfers the zone B production line to backup power, but its 1.2 MW load exceeds the 1.0 MW backup capacity. Using circuit load and supply capacity, the energy agent reduces non-critical zone D loads - air conditioning and lighting - by 25%, closing the shortfall while keeping production running.
- **Reassignment for qualifications.** Li is the nearest engineer but holds only a low-voltage electrician qualification. The personnel agent checks the qualification register, rejects that candidate and assigns Zhang, who holds a high-voltage qualification, is available and is 42 m away. Distance alone would select the wrong person; authorisation facts prevent that mistake before dispatch.
- **Blocked route and smoke reading.** A temporarily parked lorry blocks the inspection robot, so the inspection agent plans a detour. A trace smoke reading in zone C is cross-checked against aerial thermal images, which show mildly overheated insulating oil smoking without open flame. The safety agent keeps firefighting resources on standby. Cross-checking multiple sources avoids an excessive response.




Figure 5-2. English storyboard of the On-site agent dispatch 3D simulation. AOV video analysis and automatically generated review reports are capability directions.

5.4 Incident review

After the incident closes, the review agent summarises the data gathered, assessments made, past experience consulted and actions taken, organised by site object. Unlike a single duration metric, the review retains the complete causal chain. Each item links back to evidence and provides material for new rules.

Table 5-2. Incident review (simulated example)

Dimension	Count	Content
Site objects	12	Transformer C-3; zone C substation; robot #02; aircraft #01; 3 evacuated people; zone B production line; zone D load; Li; Zhang; backup power; spare part JT-120; AGV-03
Data	6 types	46-channel time series (temperature, vibration, smoke and access control); C-3 infrared images; overhead thermal images of zone C; Zhang's AOV video; badge locations and qualification records; circuit loads and backup capacity
Assessments	9	Abnormal temperature; high risk; power shortfall requires load reduction; qualification mismatch requires reassignment; blocked road requires a detour; smoke without flame; connector degradation identified; missing earthing lead requires a reminder; power-restoration conditions met
Past experience	17 records	12 similar overheating work orders, 3 records for the same transformer model, 1 high-voltage work SOP and 1 replacement record for connectors from the same batch
Actions	11	Access restriction and sound/light notification; evacuation of 3 people; robot and aircraft dispatch; reassignment to a qualified electrician; zone D load reduction; switching off and transferring to backup power; robot detour; AGV delivery of spare parts; earthing reminder; switching on and lifting restrictions; report sent to the maintenance group

The review recommends checking 12 connectors from the same batch and including them in predictive maintenance. Once confirmed, this recommendation can become an early-warning rule in the real-time engine.

5.5 Mapping to four operational AI capabilities

Four Lindos OS AI capabilities are connected to models and can run: **anomaly explanation, dispatch recommendations, task orchestration and workforce reports**. Their corresponding stages in the example appear below. Other capabilities in the example are simulated demonstrations of directions supported by the architecture.

Table 5-3. Four operational AI capabilities mapped to the example

AI capability	Corresponding stage	Status
Anomaly explanation	Assessment: relate the C-3 temperature anomaly to object relationships and similar work orders, explaining the affected object (zone B production line) and supporting evidence	Model-connected and runnable
Dispatch recommendations	Dispatch: non-critical load reduction, reassignment of qualified staff and spare-part allocation	Model-connected and runnable
Task orchestration	Assessment through recovery: split the response into 4 parallel tasks with responsibilities, dependencies and recovery order	Model-connected and runnable
Workforce reports	Review: summarise arrivals, check-in/check-out and work-compliance checks	Model-connected and runnable
Multi-agent collaboration lanes, AI analysis of AOV video and automatically generated review reports	Dispatch, coordinated work and review	Simulation; capability directions supported by the architecture

Model-connected and runnable means the capability is connected to a model and can operate. The stages listed are its counterparts in the simulation; they do not establish that it has been deployed with this workflow in a real project.

CHAPTER 06

6. Typical use cases

The following six use cases draw on Lindos OS applications and solution design. Each describes the site problem, the IoT4Agent approach and the actions available to agents.

6.1 Workplace safety

Site problem

Access to controlled areas, work-permit execution and unusual dwell times rely on manual patrols. High visitor turnover makes complete retrospective records difficult.

IoT4Agent approach

Positioning and sensing cover people, equipment and assets. Controlled areas, permits and qualifications form relationships between site objects. The real-time engine detects boundary crossings, excessive dwell and unusual gatherings immediately.

Actions available to agents

Check authorisation on entry. Continuously verify that workers and supervisors are present and hazardous areas are controlled. For two-person work, confirm that both authorised people are in position. Notify and escalate violations, and automatically generate incident reports and compliance evidence.

6.2 Asset management

Site problem

Registers and physical stock disagree, requiring item-by-item manual checks. Finding valuable mobile equipment takes time, while idle capacity and shortages coexist.

IoT4Agent approach

Managed assets have queryable live locations. Geofences control valuable assets. Use frequency, duration and turnover produce objective data, with complete records for receipt and transfer.

Actions available to agents

Complete stocktakes through automated comparison; locate the nearest available equipment by area and recommend allocation; immediately report unauthorised removal in the risk register; generate transfer and procurement recommendations.

6.3 Contract performance oversight

Site problem

Attendance and coverage for outsourced inspection, security and cleaning services are difficult to verify, often causing monthly settlement disputes.

IoT4Agent approach

Staff wear positioning badges while spatial anchors broadcast continuously. Badges are collected at the start of a shift and returned to a smart charging case afterwards. Attendance duration, patrol routes and area coverage are generated automatically.

Actions available to agents

Identify and notify missed patrols, unattended posts and overdue tasks. Summarise hours, compliance and the basis for deductions using contract definitions. Build monthly supplier performance records to support renewal assessments.

6.4 Work-order operations

Site problem

Dispatch relies on experience, arrivals and completion are confirmed verbally, and acceptance and assessment lack objective evidence.

IoT4Agent approach

Work-order locations are linked to staff locations, skills and availability. Arrivals, dwell times, routes and work areas remain traceable. Completion time and location provide acceptance evidence.

Actions available to agents

Match the nearest available qualified person by location; reassign or escalate overdue arrivals; verify and archive completion-time and location evidence; build baselines to improve staffing and service levels.

6.5 Emergency management

Site problem

During an incident, the coordination centre struggles to establish personnel distribution and status promptly. Muster checks and reviews depend on verbal accounts.

IoT4Agent approach

Everyone's location and status are visible in real time, with automatic density counts in key zones. Response plans are linked to zone topology and personnel roles. Locations, times and actions are fully recorded.

Actions available to agents

Dispatch nearby staff under the response plan, check arrivals at assembly points and coordinate public address, access control, sound and light devices. Feed incident and drill data back into the plan library to improve plans, processes and staffing.

6.6 Municipal services

Site problem

Waste collection and municipal inspections cover many dispersed sites. Vehicle and staff dispatch rely on experience, and work is difficult to verify.

IoT4Agent approach

Vehicle, personnel, station and facility status remains connected. Routes and area coverage can be checked. Both directed and undirected inspections are managed within one task system.

Actions available to agents

Arrange collection on demand and divert queues between stations; allocate inspection tasks and respond to emergencies; automatically count trips and coverage to support assessment and settlement.

CHAPTER 07

7. Safety, trust and governance

When agents act on site, system errors can have direct physical consequences. IoT4Agent includes safety and governance from the architecture stage. This chapter describes Lindos OS design requirements, to be implemented progressively through engineering validation.

7.1 Data minimisation and permission domains

- **Collect for the task.** Set positioning accuracy, sampling frequency and retention periods according to the use case. If zone-level location is sufficient, do not collect finer movement trails.
- **Define permission domains.** The design gives each agent, user and external system a clear domain specifying the objects it may read and the actions it may execute.
- **Tiered storage and masking.** Sensitive data such as identity and physiological readings is stored and displayed by sensitivity level. External interfaces mask sensitive information by default.

7.2 Evidence chains, replay and audit

The Lindos OS core design requires evidence chains to support the following:

- Trace any assessment or action back to its original records.
- Record capture time and arrival time separately so replay order reflects events on site.
- Retain rule versions, model versions and runtime settings with records to reproduce the conditions of an assessment. Propagate fact revisions through dependencies, flagging and updating affected assessments and reports.

7.3 People in the loop and risk-based authorisation

Agent actions are authorised by risk level (Table 7-1). Items receiving no response within the time limit escalate automatically. Human confirmations, rejections and corrections are recorded as evidence and inform rule improvements.

Table 7-1. Risk-based action authorisation (recommended approach)

Level	Action type	Examples	Authorisation
Level 1	Informational actions	Reminders, notifications, reports and recommendations	Agents execute automatically with a complete record
Level 2	Reversible on-site actions	Sound/light notifications, enabling geofences and reducing non-critical loads	Execute automatically under the response plan; duty staff may cancel at any time
Level 3	High-risk or irreversible operations	Electrical switching, lifting restrictions and restoring power	Execute after human confirmation. In emergencies explicitly authorised by a response plan, notify people immediately after execution.

7.4 Replaceable models

Large language models and specialist classification models provide callable computing capabilities. Lindos OS maintains common input/output contracts, evidence structures and operating policies. When model versions or deployment environments change, higher-level business contracts stay stable, while the real-time engine continues to provide safety-related fallback responses.

7.5 Privacy and compliance

- Personnel positioning involves personal information. Before deployment, clearly explain the purpose, scope and method of collection. For projects in China, apply data-minimisation and purpose-limitation principles under the Personal Information Protection Law and other applicable laws and regulations.
- Limit collection to working hours and work areas. Badges are collected at the start of work and returned to their charging case afterwards, linking positioning to working time.
- Spatial facts and evidence chains can form compliance records for requirements in sectors such as hazardous chemicals. China's AQ 3064.3-2025, for example, addresses construction, technical and data-exchange requirements for personnel positioning.
- Data storage location, retention periods and external exchange scope are determined by the compliance requirements of the project's industry and jurisdiction.

CHAPTER 08

8. Implementation path and open ecosystem

8.1 Start with one use case

IoT4Agent need not be built all at once. Start with a frequent use case that has a clear problem and measurable outcomes, then follow the five steps in Table 8-1.

Each step creates value on its own: the first makes the site visible; the second makes data usable; the third starts reducing manual responses. By the time agents are introduced at step four, facts and rules are ready.

Table 8-1. Five steps from one initial use case

Step	Main work	Completion criterion
1. Deploy the sensing network	Deploy sensors, receivers and relays for the use case, connecting existing devices and systems	Key objects stay connected and data is delivered completely
2. Establish spatial facts	Model spaces and business relationships, unify the device model and connect the evidence chain	Spatial facts can answer the core operational questions
3. Configure deterministic responses	Turn management requirements with clear criteria into real-time rules linked to devices and work orders	Rules automatically respond to common anomalies
4. Introduce reasoning and agents	Connect anomaly explanation, dispatch recommendations, task orchestration and workforce reports; progressively extend multi-agent collaboration	Recommendations are adopted and findings become rules
5. Replicate and expand	Develop industry templates and rule packages for reuse across more sites and use cases	New sites deploy from templates and reuse cross-site experience

8.2 Open ecosystem

- **Open device access.** Third-party devices can report through the Lindos protocol and receive device-level commands. Specialist equipment and environmental sensors connect over Modbus, MQTT and HTTP.
- **Open interfaces.** Business systems use APIs and webhooks, including BMS, ERP and MES integration. Agent-facing query, subscription, command and acknowledgement interfaces and SDKs are being developed under the core design.
- **Shared capabilities.** Industry templates and rule packages can be reused and developed jointly. Partners can deliver industry solutions on Lindos OS.
- **Stable contracts.** The core design plans to expose evidence structures, event formats and acknowledgement formats as stable contracts, enabling different agent frameworks to connect.

CHAPTER 09

9. Conclusion: spatial intelligence, from insight to action

How quickly agents enter the physical world depends on how accessible that world becomes to them. Earlier IoT connected devices to networks and delivered data to dashboards. Its next stage must make sites readable, trustworthy and actionable for agents. Lindos OS offers three lessons:

- **Facts enable intelligence.** Continuous, low-power sensing and evidence-backed spatial facts determine how comprehensively agents see and how accurately they assess.
- **Determinism establishes trust.** Rules handle clearly defined responses; large language models handle assessment and synthesis. The system can be both reliable and able to improve.
- **Verification establishes value.** A command has genuinely changed the site only when its outcome is acknowledged and verified.

We invite industry partners to advance IoT4Agent together:

1. **Use spatial facts as a common language.** Describe objects, relationships, events and evidence consistently, connecting devices from different vendors and agents from different frameworks.
2. **Base action on evidence and acknowledgements.** Every agent action on site should have a basis, an acknowledgement and a verifiable, replayable record.
3. **Connect the ecosystem through open protocols and interfaces.** Devices, systems and agents connect through open interfaces, while templates and rules share reusable experience.

Spatial intelligence. See ahead, act ahead.

Sense to Act

APPENDIX A

Appendix A. Glossary

Table A-1. Glossary

Term	Expanded or technical name	Definition
IoT4Agent	Internet of Things for Agents	IoT built for agents, making physical sites readable, trustworthy and actionable
Spatial fact	Spatial Fact	A site fact with an object, status, location, time, source and version
Real-time engine	Reflex Engine	A deterministic reflex engine responding in milliseconds without large language models
Reasoning engine	Reasoning Engine	A minute-scale reasoning engine using large language models orchestrated by Harness
Harness	Harness	A runtime framework that orchestrates large language models, tools and data
Action transaction	Action Transaction	A controlled state transition from decision to execution, acknowledgement and verification
Permission domain	Permission Domain	An authorisation boundary defining readable objects and executable actions
Human-in-the-loop	Human-in-the-Loop	A mechanism requiring people to confirm key assessments and high-risk actions
AOV body-worn camera	-	A wearable duty camera that streams first-person video; AI video analysis is a capability direction supported by the architecture
Three core propositions	Sensing as a Service · Source of Truth · Sense to Act	Continuous sensing (CELL + GRID), trusted data foundation (PORT + CORE) and dual-engine decisions and actions (CORE + AXIS), corresponding to readable, trustworthy and actionable

APPENDIX B

Appendix B. Data and test conditions

Table B-1. Data and test conditions

Data	Value	Conditions and definitions	Type
Sleep current	0.36 μ A	In-house 2.4 GHz proprietary-protocol firmware on real hardware, complete device including IMU; system-level test on 27 May 2026	Measured
Average current	7.72 μ A	As above	Measured
Peak current	3.91 mA	Instantaneous transmit/receive current in the same test	Measured
CR2032 battery life	$\geq$ 3.3 years	Estimated from measured average current: 225 mAh $\div$ 7.72 μ A. Actual life depends on environment, temperature and settings.	Theoretical estimate
Chapter 5 example figures	46 sensing channels, 78 $^{\circ}$ C, 12 work orders, 0.2 MW and other values	Lindos OS 3D demonstration: On-site agent dispatch	Simulation

This white paper uses the data categories above and states their conditions alongside the figures. It does not cite other performance, market or commercial figures; standard identifiers and years identify their sources.

This white paper explains the IoT4Agent concept and the Lindos OS reference architecture. Product capabilities depend on the current Lindos OS version. Content marked as a simulation or capability direction does not represent a released feature. Specific functions and metrics depend on project configuration and measurement.