Hiroshima AI Logistics Practical Guide 2026

Bringing international principles into logistics management and operational decisions.

Go beyond introducing AI: define the conditions for its use, check the supporting grounds, and connect its use to operational results.

Turning AI proposals into decisions that businesses can use with confidence.

A practical starting point for business leaders, manufacturing and logistics managers, operational staff, and local government officials. No specialist knowledge is assumed. The guide covers preparation before introduction, verification, operation, and collaboration between companies. Its central approach is to consider business results, safety and trust, and real-world implementation together.

How to read this guide

Start with the overall picture	Business results and investment payback, six points before introduction, the Hiroshima AI Process	Pages 2–4
Learn how to check	Public PoCs, the scope of formal verification, designing connections between companies' decisions	Pages 5–7
Learn from operational examples	Practical manufacturing and logistics cases (with sample entries)	Pages 8–9
Get started in your company	Roadmap, audit checklist, record templates	Pages 10–12
Review and take the next step	Practical issues from the conference, references, contacts for consultation	Pages 13–14

Three tools to take back and use

Roadmap and allocation of roles	Define the conditions for advancing at each stage and who is ultimately accountable	Page 10
Audit checklist	Use critical conditions to decide whether to begin trials or move to production	Page 11
Decision and Incident Response Record Templates	Record the grounds for decisions and the process from stopping to restarting	Page 12

Manufacturing AI Logistics Conference 2026 in Hiroshima

Wednesday, October 14, 2026 International Conference Center Hiroshima

Editor and publisher / Conference organizer ONZALINX Co., Ltd.
Technical cooperation AI Assurance, Inc.
Related activities Hiroshima AI Assurance Council (HAAC)

October 2026 | Conference issues are summarized on page 13

01 / Practical Guide

Design AI logistics around business results

The starting point is not “Which AI should we buy?” but “Which decisions should we change, and whose workload should we reduce?”

This practical guide was compiled by conference organizer ONZALINX Co., Ltd., which sees logistics as a foundation of business management. It organizes themes connecting management, operational data, technology, governance, and implementation into points that companies can take back and check. Alongside published results, it explains how to get started in practice.[1]

First, define the target operation and how to measure results

Choose one operation to address first, such as revising delivery plans, adjusting dock appointments, replenishing inventory, or checking quality. Review current procedures and exception handling, and separately write down the results you want to improve and the conditions that must not be compromised.

Desired results	Examples of metrics to set in your company	What to check alongside them
Reduce operational workload	Staff working time; number of checks and repeated data entries	Has work simply been transferred to another department or business partner?
Improve logistics quality	On-time delivery, waiting time, inventory shortages and surpluses, time required for quality checks	Have improvements in speed or cost come at the expense of safety or quality?
Make decisions easier to check	Number of cases where reasons or records could be checked; time to resolve items returned for review	Have exceptions requiring checks decreased only because they were overlooked or left unrecorded?

Compare before and after introduction under the same conditions

Align the comparison periods, cargo types, number of items processed, and workload levels. Include time added after AI introduction for inspection, correction, and training, and assess whether the workload has decreased across the operation as a whole. Where comparison conditions cannot be aligned, leave the effects in that scope unconfirmed.

Estimate benefits and investment payback in the same table

Review financial benefits and changes in safety, quality, and ease of working in the same table. The example below estimates the effects of measures to reduce waiting time.

Item	Calculation basis (assumed)	Annual amount (assumed)
Less coordination work	40 fewer hours/month adjusting appointments × JPY 3,000 × 12 months	JPY 1,440,000
Less overtime for on-site work	60 fewer overtime hours/month caused by waiting × JPY 3,500 × 12 months	JPY 2,520,000
Operating costs	Service fees and maintenance, plus additional inspection and training time after introduction	−JPY 1,800,000
Net benefit and payback period	Initial cost of JPY 3,000,000 (data preparation and trial) ÷ net benefit of JPY 2,160,000	JPY 2,160,000 / approx. 1.4 years

Gradually expand the scope entrusted to AI

Start with AI generating candidates and people making decisions under existing procedures. For operations with established checks and exception handling, proceed to automatic adoption and execution within defined conditions. Design human involvement to include not just an approval field, but also authority to stop, the information needed for checks, and time to respond.

The estimates use assumed figures. The first step is to select one target operation, one outcome indicator, and one condition to uphold (see pages 8–9). The introduction procedures, sample metrics, and checklist items in this guide are practical arrangements developed by the editorial team. They do not assume the introduction of any particular product.

02 / Practical Guide

Decide six things before introduction

Making AI ready for use requires a purpose, conditions, human roles, supporting grounds, exception handling, and records.

With reference also to the AI Guidelines for Business from the Ministry of Internal Affairs and Communications and the Ministry of Economy, Trade and Industry, this guide sets out six practical starting points. These are questions for designing your own operations, not official certification criteria.[2]

What to check	What to decide or check first
01 What will it be used for?	Define the target operation, purpose of use, people and business partners affected, and what will and will not be entrusted to AI.
02 Under what conditions will proposals be adopted?	Define in advance the conditions to be met for quality, safety, contracts, deadlines, authority, and other matters. Distinguish unmet conditions from insufficient information.
03 Who decides, and who stops it?	Clarify the people responsible and their authority for adoption, execution, returning proposals for review, exception decisions, and stopping and restarting. Also designate deputies.
04 What grounds support the decision?	Check data sources, what the data relates to, collection times, units, missing data, model and rule versions, and permission to use them.
05 Can it stop and be reviewed when anomalies occur?	When information is insufficient, communications fail, inputs are abnormal, or proposals are unexpected, do not proceed with processing that cannot be checked. Put it on hold, refer it for human review, or use alternative procedures.
06 Can it be checked again later?	Link the conditions, evidence, versions, determinations, and execution results used so they can be traced in response to inquiries or incidents.

Do more than collect data: align what it represents

Even “arrival time” means different things for arrival on site, check-in, and the start of unloading. Align cargo, site, and task identifiers, times, units, update frequencies, and handling of missing data. RFID or sensors alone do not automatically align operational meanings or ensure correct data.

Adjust the rigor of checks to the criticality of the operation

AI providing reference information and AI automatically controlling equipment or deciding whether shipments may proceed need different criteria. Adjust verification, approval, and monitoring to the potential for wrong decisions to cause injury, quality defects, or business interruption. Also define external transmission of personal or confidential information, viewing permissions, and a contact point for reporting errors alongside the category.

Category and examples	Verification required before introduction	Approver	Monitoring during operation
Category 1 Reference information Indicative demand forecasts, draft reports	Sample checks and identification of error patterns	Operational staff	Review accuracy and usage monthly
Category 2 Execution with human approval Delivery plans, dock appointments, replenishment proposals	Check conditions, compare trials under the same conditions, and review exception procedures	Business owner	Review metrics weekly; record each exception
Category 3 Automatic execution / high impact Equipment control, quality determinations, shipment authorization	Independent testing, limited operation, and drills for stopping and switching to manual operation	Management and the quality and safety manager	Continuous monitoring; automatic hold based on stopping criteria

These categories are examples developed by the editorial team. When in doubt, use the more stringent category. Leave conditions that cannot be checked as unconfirmed, rather than marking them “no problem.”

03 / Practical Guide

Making the Hiroshima AI Process usable in practice

From international principles to company rules.
From company rules to verifiable decisions.

The Hiroshima AI Process began with the 2023 G7 Hiroshima Summit and developed into a framework that includes international guiding principles for all AI actors and an international code of conduct for organizations developing advanced AI systems.[3]

Through the activities of the Hiroshima AI Assurance Council (HAAC), ONZALINX is working to connect these principles to operational decision conditions and evidence. AI Assurance, Inc.'s "Technical Platform for the Hiroshima AI Process" shows the relationships between principles, operational obligations, evidence, verification scope, and unresolved issues.[4][5]

International principles
Guidelines and code of conduct
→
Company policies
Rules and responsibilities
→
Operational decision conditions
Adopt, hold, stop
→
Evidence and review
Records, monitoring, reassessment

Examples of translating principles into practice

Approach	Example operational checks	What to retain
Prioritize safety	Do not proceed unchanged with deliveries or handoffs that fall outside permitted conditions	Conditions, measurements, determinations, reasons for holds or returns for review
Improve transparency and accountability	Make it possible to trace the grounds and authority for each decision	Evidence references, person responsible, point in time, rule version
Address risks continuously	Use changes to data or systems and the occurrence of anomalies as triggers for rechecking	Changes, scope of impact, retesting, restart decisions

The table gives examples developed by the editorial team. It is not a clause-by-clause translation of the international guiding principles or a mapping of all requirements.

The roles of AI assurance and Responsibility OS

AI assurance is an evidence-based approach to checking the appropriateness of AI-related design, implementation, operation, and decisions. It combines formal verification, testing, operational evaluation, and other methods. This guide focuses on the adoption and execution of logistics decisions.[5][6]

Responsibility OS is a technical platform developed by AI Assurance, Inc., with formal verification at its core. It links requirements, evidence, responsibility information, and verification scope to conditions for adopting and executing decisions, connecting condition checks and records to existing operations. Its scope includes not only generative AI, but also optimization, rule processing, human proposals, and other approaches.[5][6][10]

Generate candidates
AI, optimization, people
→
Verify adoption conditions
Evidence, authority, subject, version
→
Decide on adoption
Adopt / Hold / Do not adopt
→
Execute and record
Authorized processing only

Conceptual diagram: verification results, the organization's adoption decision, and authorization to execute are distinct. Proposals on hold or rejected do not proceed to execution; they return to the necessary checks.

Distinguish a "good proposal" from a "decision that may be executed."

Combine the ability to generate proposals with the ability to check whether they can be used in operations. Clarifying conditions provides a foundation for expanding the scope of AI use, rather than a reason not to use it.

This guide and the technologies it introduces do not indicate accreditation or certification by the G7, the Japanese government, or the OECD, or proof of conformity with the Hiroshima AI Process as a whole.

04 / Practical Guide

See how decisions are checked in a public PoC

Connecting conditions, evidence, determinations, and re-verification through a pharmaceutical cold-chain example.

ONZALINX, together with AI Assurance, Inc. (then GhostDrift Mathematical Institute, Inc.), published its first PoC on pharmaceutical cold-chain handoff decisions. The code, sample data, and reproduction procedures are publicly available. A PoC is a proof of concept for testing an implementation approach.[6]

Processing checked with samples

The input format is inspected, and evidence items for transfer, seals, temperature loggers, and temperature records are checked. Metrics are recalculated from temperature records and checked against predefined conditions, and the determination and rule version used are recorded. The following four categories are specific to this PoC.[6]

PoC determination	Meaning	Default PoC behavior
RELEASE Eligible for adoption	Required evidence and defined metric conditions are satisfied	Proceed to evaluation of the next segment
HOLD On hold due to insufficient evidence or similar reasons	Records are missing or evidence conditions are unmet	Stop evaluation of subsequent segments
QA_REVIEW Additional checks	Metrics fall within the defined range requiring review	Stop evaluation of subsequent segments
REJECT Conditions not met	Metrics exceed the defined upper limits	Stop evaluation of subsequent segments

Invalid input is handled as an input error, separately from the quality determinations above. Actual actions, such as disposal, are decided under operational rules and authority.

General readers: first look at three differences

Compare a normal case, a case with missing evidence, and a case outside the temperature conditions. The key points are that a case does not pass if evidence is missing, even when “the temperature looks fine”; that subsequent evaluation stops at the segment where a problem is detected; and that the determination can be checked again with the same inputs, rules, and implementation. Technical staff can rerun the PoC using the referenced procedures.[6]

Situation	Conventional operation (example)	Operation checked through conditions and evidence
Checks at handoff	Visually check the temperature display and accept the delivery if no problem is found	Assemble temperature records and evidence of seals and transfer, then check them against defined conditions before adoption
Some records are missing	Decide based on staff experience and proceed to the next process	Place on hold as “unconfirmed” and refer to the quality manager for review

What has been checked, and what needs further checks in production

The public PoC demonstrates condition-based determinations and rerunning with sample data. Stopping real equipment, the authenticity of sensor information, integration with production systems and intercompany APIs, and interoperability between different implementations are outside the scope checked by this PoC.[6][8]

2–8°C is the setting used in this PoC. It is not a temperature condition common to all pharmaceuticals. In actual operation, check criteria and thresholds for each product and transport condition. See reference [6] on page 14.

05 / Practical Guide

Keep formal verification and implementation connected

Rule properties, program behavior, and conditions on site.
Check how these three correspond.

Formal verification mathematically describes a subject and checks that its properties hold under specified assumptions. Lean 4 is a proof assistant that expresses proofs in a form a computer can check.[7]

AI Assurance, Inc. is developing a public technical platform that starts from the official Hiroshima AI Process documents and connects, through Responsibility OS, to selected technical requirements of the EU AI Act. Evidence is reused while preserving its subject, purpose, role, version, and unconfirmed items; any additional conditions required at the destination are checked separately.[10]

Layer to check	Key question	Example verification materials
Formal verification Model and logic	Do the required properties hold under the defined assumptions?	Propositions, assumptions, proof files, checking environment and results
Implementation testing Program and connections	Can the target process handle normal, abnormal, and boundary conditions?	Version tested, test inputs, expected results, actual results
Operational checks Individual tasks and executions	Are execution conditions met for the current subject, evidence, authority, and point in time?	Valid records, personnel and authority, decisions, execution results

Checks also cover correspondence between the public PoC and Lean

Public automated verification through GitHub Actions on October 10, 2026 (commit dc67c7b) successfully checked Lean 4 (v4.26.0) proofs, ran 18 Python tests, connected to the fixed version of the logistics PoC (commit 1437b21), and checked its execution results against the formal model. A later update that day (a479ee9) changed only the README; the code under verification was unchanged.[10]

In terms for management and on-site staff

Formal verification checks that, if the specified assumptions are correct, the rules always behave as stated. Whether those assumptions match your equipment, data, and operations is checked separately through on-site checks and operational monitoring. Passing a finite set of tests is also distinguished from proof covering all cases. This guide treats the distinction between the scope of proof and real-world conditions as a basic principle throughout. The version, assumptions, procedures, results, and unconfirmed items are made public, providing a way for third parties to recheck them.

Even if the original decision passes, missing additional conditions prevent it from proceeding.

Even when the logistics PoC returns RELEASE, the decision is put on hold if documents or monitoring plans required at the destination are missing. Only when all conditions are met does it proceed to the adoption process in the reference implementation.[10]

This does not demonstrate full compliance with the EU AI Act, official certification, or the safety of an entire real-world environment. Checking individual PoC executions is distinguished from formally proving the entire implementation. This is the verification scope of the public example and does not limit the fields of application of Responsibility OS.

06 / Practical Guide

Connect decisions across companies

Pass on not only goods and data, but also the conditions under which work may proceed.

The Physical Internet is a concept for sharing and coordinating logistics resources through common mechanisms.[9] ONZALINX and AI Assurance, Inc. have jointly announced a Verifiable Physical Internet concept that allows companies to participate while retaining their own objectives and existing systems.[11]

Translate different objectives into conditions for joint execution

Translation of objectives means converting each company's objectives into conditions that another company's AI or staff can assess, without unifying all objectives into one. The aim is coordination that verifies the necessary conditions without requiring blanket disclosure of detailed costs, customer information, or similar data.[11]

Supporting element	Role in joint execution
Beacon	Conditions shared so that each company can adopt collaborative logistics. These define delivery deadlines, quality, burden limits, consent, and other requirements.
Logistics decision packet	Links a plan to requirements, evidence, responsibility information, approvals, versions, deadlines, and adoption status, and passes them on as a unit of decision-making.
AI assurance platform	Verifies that conditions hold. Distinguishes verification results, organizational adoption decisions, and permission to execute individual operations; deficiencies are returned for correction or replanning.

Set fairness conditions before collaborative delivery by three companies (example)

Sustaining collaborative delivery requires agreement before participation on how costs, benefits, burdens, and losses will be shared. The following is a completed example in which parts manufacturers A and B and carrier C jointly deliver in the same direction.

What to decide	Description	Completed example (assumed)
Acceptance conditions	Minimum conditions for each company to continue participating	A: delivery within ±30 minutes of the scheduled time / B: zero quality incidents / C: compliance with laws on operating hours
Burden limits	Maximum additional burden each company will accept	Additional work for waiting and transshipment: up to 10 hours per company per month
Consent	Who consented, and to what scope	Each company's logistics manager agrees to the routes, period, and shared data covered
Cost and benefit allocation	How savings and additional costs are shared	Transport cost savings are allocated in proportion to weight and distance. Additional costs are borne by the company that causes them
Withdrawal and review conditions	How unmet conditions are handled	Renegotiate if acceptance conditions are not met for two consecutive months. Withdrawal is allowed with 30 days' notice

The completed examples use assumptions for illustration. Agreement is based on verifying that conditions hold, without requiring blanket disclosure of detailed costs or customer information.

Physical Internet Verified Scalability formalizes in Lean 4 matters such as the sufficiency of evidence needed to decide whether execution may proceed and the connections between responsibility records (these are conditions within the formal model, not empirical demonstrations of performance or cost). The plan is to reflect findings from demonstrations in HAAP, which HAAC is developing.[8][11][12]

This does not replace existing laws, contracts, or organizational authority. Logistics is one application; the technology's scope is not limited to logistics and also includes manufacturing, quality management, local government operations, and other fields.

07 / Practical Guide

Manufacturing Operations Case: AI proposals for parts replenishment

A completed example shows the flow: Purpose → Data → AI proposal → VERIFY → Human decision → Execution → Measurement of results.

In this example, AI proposes replenishment of parts beside an assembly line using production plans and actual inventory data. Arranging and transporting parts, previously reliant on replenishment staff's experience, becomes a structured process with conditions checked along the way. The scope is limited to one line and a group of parts prone to stockouts.

Outcome indicator	Current state (assumed)	Target (assumed)
Line stoppages due to parts stockouts	4 per month	0 per month
Inventory beside the line	6 days' supply on average	4 days' supply on average
Checks and transport planning by replenishment staff	90 minutes per day	45 minutes per day

Step	What to do	Completed entries for this example
① **Purpose**	Define the outcomes to improve and the conditions that must not be compromised	Eliminate stockout stoppages. Maintain inventory limits and safe aisle widths
② **Data**	Align the meaning, timing, and units of the data used	Production plan (daily), bill of materials, actual inventory data (every 2 hours), expected arrivals, defect and rework records
③ **AI proposal**	Have AI generate candidates	Propose replenishment times and quantities for each part, and the transport sequence
④ **VERIFY**	Check adoption conditions	Inventory limits, storage capacity, production plan version, data freshness (within 2 hours), and transport staff workload
⑤ **Human decision**	Decide whether to approve, put on hold, or return for review	The team leader approves. Immediately after plan changes or when data is old, put the proposal on hold for the coordinator to check
⑥ **Execution**	Execute only within the permitted scope	Issue only approved replenishment instructions. Do not write automatically to equipment or ordering systems
⑦ **Measurement of results**	Compare under the same conditions and record the results	Compare stockout stoppages, inventory days, work time, and returns for review across periods with similar production volumes and part-number mixes

Example criteria for success

Compare periods with similar production volumes and part-number mixes. There should be no stockout stoppages, inventory limits should be respected, and aisle safety and burdens on other processes should not worsen. Treat this operation as Category 2 (execution with human approval) and keep records in the format on page 12.

Figures are assumptions for illustration, not actual demonstration results.

08 / Practical Guide

Logistics Operations Case: Reduce waiting time

A completed example of rescheduling dock appointments at a logistics center using AI proposals and on-site decisions.

AI proposes changes to truck appointments based on expected arrivals, dock availability, and receiving capacity. The on-site manager checks the conditions and approves the changes, and results are measured under the same conditions.

Outcome indicator	Current state (assumed)	Target (assumed)
Average waiting time per truck	60 minutes	30 minutes (50% reduction)
Overtime for receiving work	60 hours per month	30 hours per month

STEP 1 | AI proposes
Propose appointment changes from expected arrivals, dock availability, and work capacity

→

STEP 2 | VERIFY
Check safe staffing levels, equipment capacity, vehicle constraints, and contract terms

→

STEP 3 | On-site decision
Approve if conditions are met. If unclear, put on hold and recheck

→

STEP 4 | Execute and measure
Update appointments and record waiting time, overtime, and returns for review

Example criteria for success

Compare under similar arrival volumes and workload levels: waiting time should fall without worsening work safety or other companies' waiting times.

Figures are assumptions for illustration, not actual demonstration results.

Apply the same flow to other operations

Writing down the conditions to uphold is the common starting point for every operation. Try one representative example and standardize the record format before expanding.

Field	Example AI application	Conditions and checks
Manufacturing	AI proposes when to replenish parts	Stockout prevention, inventory limits, consistency with production plans
Warehousing	AI optimizes picking sequences	Prevention of shipping errors, work safety, equipment capacity
Transport and delivery	AI replans delivery routes	Laws, operating hours, delivery deadlines, vehicle constraints
Quality management	AI detects temperature deviations and anomalies	Product specifications, evidence, quality manager approval
Intercompany coordination	AI systems coordinate collaborative delivery proposals	Each company's burden limits, confidentiality, consent to participate

09 / Practical Guide

Implementation roadmap and responsibilities

Start small and advance only when each stage's conditions are met.

Stage	Main tasks	Conditions for proceeding
0 Preparation About 1 month	Define the target operation, outcome indicators, and conditions to uphold; align data meanings	The six items on page 3 are completed and the risk category is determined
1 Trial (parallel operation) About 1–3 months	Generate AI proposals, but use existing procedures for actual work; compare proposals with actual results	Accuracy, deviation rates, and exceptions are recorded; no critical condition remains unconfirmed
2 Limited operation About 3 months	Limit the scope, hours, and number of cases; execute after human approval	Results are confirmed under the same conditions; stopping and switching to manual operation have been practiced
3 Production operation	Operate within the defined scope and manage using monitoring indicators and stop criteria	The accountable person has approved the transition and records have been retained
4 Expansion and intercompany coordination	Expand to other operations, sites, and business partners; coordinate under common conditions	Each participating company's acceptance conditions, burden limits, and consent have been checked

Durations are indicative. For Category 3 operations, allow longer at each stage and add independent testing.

Simplified RACI: Who executes and who is accountable

Activity	Management	Business owner	Quality and safety manager	System administrator	External vendor
Decide objectives and investment	A	R	C	C	I
Decide adoption conditions and risk category	A	R	C	C	C
Conduct trials and verification	I	A	C	R	R
Approve transition to production	A	R	C※	C	I
Stop and restart during operation	I	A	R	R	C
Reassess when changes occur	I	A	C	R	C
Report incidents and prevent recurrence	A	R	R	R	C

R = Responsible for execution; A = Accountable (one person per activity); C = Consulted; I = Informed. ※Category 3 also requires approval from the quality and safety manager.

Operating with a small team

People may hold multiple roles. However, identify the people with accountability (A) and authority to stop by name, not just job title, and ensure Category 3 decisions are not approved solely by the proposer. Anyone should be able to request a stop for a safety anomaly, regardless of role. Even when execution (R) is assigned to an external vendor, accountability (A) remains within your company.

10 / Practical Guide

AI governance and audit checklist

Check the critical conditions marked ★ first, then decide whether to proceed at each stage.

Target operation: Risk category: 1 / 2 / 3 Accountable person: Date checked:

Check item	Critical	Status	Evidence, corrective action and deadline
01 Purpose and scope Are the outcomes, those affected and the scope not delegated to AI clear?	★	□Done □Improve □Unconfirmed	
02 Adoption conditions Are the conditions for adoption, hold and non-adoption, and their versions, defined?	★	□Done □Improve □Unconfirmed	
03 Roles and authority Are the persons in charge of adoption, stopping and restarting, and the accountable person, clear?	★	□Done □Improve □Unconfirmed	
04 Data and evidence Can sources, times, units, missing data and scope of use be checked?		□Done □Improve □Unconfirmed	
05 Performance evaluation Are accuracy, failure rates, deviation rates and operational impacts measured?		□Done □Improve □Unconfirmed	
06 Information security Are transmissions to external AI, access rights and connection destinations controlled?	★	□Done □Improve □Unconfirmed	
07 Stopping when abnormalities occur Are stop criteria defined, and have stopping and manual switchover been practiced?	★	□Done □Improve □Unconfirmed	
08 Records and re-verification Are records linked, and have tamper prevention and retention periods been defined?		□Done □Improve □Unconfirmed	
09 Change management Is reassessment performed when AI, data, rules or other elements change?		□Done □Improve □Unconfirmed	
10 Laws and contracts Have applicable laws, contracts and internal rules been checked?	★	□Done □Improve □Unconfirmed	

Transition stage	Conditions for proceeding (status: done / needs improvement / unconfirmed)	If conditions are not met
Trial (parallel operation)	May begin even if some ★ items remain unconfirmed, provided proposals are not used in actual operations	Do not use proposals for execution
Limited operation	All ★ items are confirmed. Other items needing improvement have corrective action plans with assigned persons and deadlines	Do not begin; recheck after correction
Production operation	All items have been checked. The accountable person records risk acceptance for any remaining items needing improvement	Continue limited operation and set a recheck date

Do not decide by total score. Results do not signify legal compliance, safety certification or authorization for production operation. Reference frameworks [2] [13]–[15] (no claim of conformity).

11 / Practical Guide

Templates for AI decision and incident response records

Record the grounds for decisions and the steps from stopping to restarting in a consistent format.

A. AI decision record

Item	Example entry (case on page 8; hypothetical)	Entry field
Record number, date and time	R-0412 / October 12, 9:40	
Target operation and object	Line 2: parts replenishment (part number group A)	
AI and rule versions	Replenishment proposal model v1.3 / Adoption conditions v2	
Evidence references	Production plan (October 12 version), inventory records (as of 9:30)	
Unconfirmed matters	Some scheduled arrivals have not been updated	
Verification result and adoption status	4 of 5 conditions met; 1 unconfirmed → On hold	
Decision-maker and authority	Team leader (Category 2 approval authority)	
Actual action and result	Approved and executed at 10:05 after the coordinator checked scheduled arrivals	

B. Procedure for responding to incidents and abnormalities

Stage	Action	What to record
1 Detection	A monitoring indicator exceeds a stop criterion, or an issue is reported from the workplace	Detection date and time, indicator value, reporting person
2 Containment	Stop AI proposals and automatic execution, and identify the affected scope	Person who stopped it, time stopped, number of affected cases
3 Manual recovery	Switch to existing or alternative procedures and continue operations	Switchover time, alternative procedures, additional work
4 Reporting	Report to the accountable person, quality manager and relevant business partners	Recipients, time, information reported
5 Cause analysis and retesting	Identify whether the cause lies in the data, AI, rules or operations, and retest after correction	Cause, corrections, test results
6 Restart approval	Check the restart conditions; the accountable person approves restarting	Approver, restart conditions, recurrence prevention measures

Examples of monitoring indicators: proposal adoption and return-for-review rates, number of deviations from conditions, missing-data rate, processing time. Triggers for reassessment: worsening accuracy or deviation rates, changes in data meaning or collection methods, updates to external AI services, changes to rules or contracts, or additional target operations or sites. Repeat testing and approval whenever any of these occurs. The existence of records does not prove that a decision was appropriate. Combine records with independent testing, validation in the workplace and operational monitoring.

12 / Practical Guide

Practical issues identified at the conference

Turn the day's discussions into issues and next actions.

Event: Manufacturing AI Logistics Conference 2026 in Hiroshima
Wednesday, October 14, 2026 / International Conference Center Hiroshima [1]

The following discussion points are organized under ONZALINX's editorial responsibility, based on the topics covered in the day's sessions and panel discussions. They are presented as a summary connecting the discussions to the practices in this guide, rather than as a transcript or matters agreed by all participants.

01 Issues facing companies and workplaces

Discussion Differences in logistics requirements and business practices among companies, along with barriers to sharing confidential information, make cross-company collaboration difficult.

Issue A framework is needed to reconcile differing requirements and interests among companies while ensuring that all participating companies benefit from collaborative logistics.

Next action Identify each company’s needs and constraints, and start with small-scale pilot projects, such as utilizing return trips and reducing empty truck runs.

02 Conditions needed for safe AI use

Discussion The success of AI depends not only on the performance of AI models, but also on data quality, business rules, and verification capabilities.

Issue AI decisions cannot be safely implemented without clear data definitions, quality standards, exception handling rules, and approval responsibilities. A system is needed to record the basis for decisions and enable re-verification when changes occur.

Next action Select one target business process and clearly define the data, decision criteria, responsible personnel, and exception handling procedures.

03 Issues for future real-world implementation

Discussion Real-world AI logistics requires not only technical validation but also proven safety and economic viability in actual operations.

Issue Clear criteria, responsibilities, and cost-benefit sharing are needed to move from pilot projects to full-scale operations. Coordination and sustained collaboration among companies remain key challenges.

Next action Use Hiroshima as a testing ground. Start with small-scale pilots, set shared goals and responsibilities, and expand gradually. Share results and lessons to build a sustainable collaboration model.

To our readers

Choose one question in this guide that applies to your company, and return to the roadmap on page 10 and the checklist on page 11. Reorganizing the conditions and evidence for existing operations, as well as introducing new AI, is a step toward putting this guide into practice.

Next steps (information from the publisher)

This guide does not presuppose the adoption of any particular product. If you would like to discuss putting it into practice at your company, the publisher, ONZALINX, offers consultation at the following stages.

① AI Logistics Ready Assessment: review current operations, data, operational arrangements and risks ② PoC: trial and effect measurement for one selected operation ③ Intercompany demonstration (HIROSHIMA FLOW): phased execution with conditions and burdens coordinated among multiple companies

For inquiries, use the contact point on the official conference website [1].

Individual contributors' messages are covered in the separate contributions volume, AI HIROSHIMA Messages. The discussion points organized in this guide are not a substitute for contributors' comments or endorsement.

This does not represent the official views of speakers, participating organizations or government agencies, their support or recommendation of this guide or the technologies introduced, or a joint statement.

13 / Practical Guide

References and public technical platforms

Use the numbers in the text to access primary sources and starting points for re-verification. Versions and review dates are also provided.

[1] ONZALINX | Official conference website
https://www.onzalinx.co.jp/conference2026/

[2] Ministry of Internal Affairs and Communications; Ministry of Economy, Trade and Industry | AI Guidelines for Business, Version 1.2 (March 31, 2026)
https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20260331_1.pdf

[3] Ministry of Foreign Affairs | AI Diplomacy (July 1, 2026)
https://www.mofa.go.jp/mofaj/ecm/st/pagew_000001_01551.html

[4] AI Assurance, Inc. | Technical Platform for the Hiroshima AI Process
https://github.com/GhostDriftTheory/hiroshima-ai-process-assurance-mappings

[5] ONZALINX | Hiroshima AI Assurance Council: Establishment and Activity Policy (October 8, 2026)
https://prtimes.jp/main/html/rd/p/000000015.000169775.html

[6] Joint PoC | Pharmaceutical Cold Chain Handoff Quality Assessment
https://prtimes.jp/main/html/rd/p/000000005.000182721.html
https://github.com/GhostDriftTheory/coldchain-handoff-assurance

[7] Official Lean documentation | Theorem Proving in Lean 4
https://lean-lang.org/theorem_proving_in_lean4/Introduction/

[8] HAAC | Hiroshima AI Assurance Protocol (HAAP)
https://hiroshima-ai-assurance.com/protocol

[9] Ministry of Land, Infrastructure, Transport and Tourism; Ministry of Economy, Trade and Industry | Physical Internet Roadmap
https://www.mlit.go.jp/report/press/tokatsu01_hh_000594.html
https://www.meti.go.jp/shingikai/mono_info_service/physical_internet/20260717_roadmap.html

[10] AI Assurance, Inc. | Hiroshima AI Process × Responsibility OS × EU AI Act
https://github.com/GhostDriftTheory/hiroshima-ai-process-eu-ai-act-bridge
Lean-verified version dc67c7b (the logistics PoC was checked against pinned version 1437b21; the subsequent a479ee9 updates only the README) / CI execution results (October 10, 2026)

[11] Joint announcement by both companies | The Verifiable Physical Internet concept
https://prtimes.jp/main/html/rd/p/000000014.000169775.html

[12] AI Assurance, Inc. | Physical Internet Verified Scalability
https://github.com/GhostDriftTheory/physical-internet-verified-scalability

[13] ISO/IEC 42001:2023 (AI management systems) https://www.iso.org/standard/81230.html

[14] ISO/IEC 23894:2023 (AI risk management) https://www.iso.org/standard/77304.html

[15] NIST AI Risk Management Framework https://www.nist.gov/itl/ai-risk-management-framework

Publication information and where to find this guide

Title: Hiroshima AI Logistics Practical Guide 2026 English Edition
Editor and publisher / Conference organizer: ONZALINX Co., Ltd.
Technical cooperation: AI Assurance, Inc.
Related activities: Hiroshima AI Assurance Council (HAAC)
References reviewed on: October 10, 2026 (versions of the public technical platforms [4] [6] [10] [12] rechecked on October 20, 2026)

Edition and revision history: Advance editorial draft (October 10, 2026, prepared for the conference with fields for additions on the day) → Revised second edition (October 2026; practical cases, implementation roadmap, allocation of roles, risk categories, audit checklist, record and incident response forms, etc. added. Descriptions of verified versions checked against public repositories)

Inquiries, publication and correction notices: ONZALINX official conference website / HAAC official website

This guide is a private-sector document for voluntary reference, edited and published by ONZALINX. It does not represent approval by all speakers and participants, a joint statement, or certification or recommendation by government agencies. For each operation, separately check the applicable laws, contracts, rules and workplace conditions. Ownership and terms of use for technologies, code and intellectual property are governed by the notices in each public resource.