From Stone and Thread to Shared Intelligence

A Credit-Card-Sized Starting Point for Local Making and Global Participation

HSIMC White Paper - English Master
Benny (Yan) Liao - Founder, HSIMC
W0-011 - P4 English Master Consolidation - 12 November 2026

Human Authority · Local Capability · Global Participation
No One Left Behind in the GenAI Physical Era

Publication note

This is an independent HSIMC white paper. It does not claim endorsement, certification, partnership, standard status or approval by the United Nations, OpenAI, the Model Context Protocol project, ISO, IATA, GHG Protocol, Hong Kong International Airport, or any other organization cited for external context.

The white paper separates proposals, documentary baselines, simulations, remote simulation records, physical reference results, delivered physical configurations and operational results. Where a claim is not yet supported by the relevant evidence level, it remains a proposal or a design target.

The standard credit-card footprint is used as a familiar planar starting point. It is not a claim that the full proposed computing, communications, power, sensing and interface system has already been integrated into an ISO/IEC 7810 compliant card.

Founder’s Foreword - Start with What Is Already Here

I have spent much of my working life close to the point where an idea stops being a drawing and has to survive the physical world.

That experience includes standards work, machine vision, industrial projects, supplier coordination and the less glamorous work that comes after a promising concept: checking dimensions, finding the right connector, making a fixture, correcting a drawing, explaining the same requirement to different engineering teams, and discovering that a result that looked obvious in software is not obvious at all on the factory floor.

The arrival of powerful GenAI makes this gap more visible, not less visible.

A model can now produce code, drawings, analyses, translations and alternative designs at extraordinary speed. But the physical world still asks questions that cannot be skipped. What material is actually available? Which machine can process it? Who is allowed to use that machine? Which version is being built? Is the interface electrically compatible? Who accepts the result? What happens if the network disappears after a physical action has already started? Who receives the value when a task is completed?

I became convinced that one of the most important opportunities of the GenAI era is not simply to build more autonomous machines. It is to make existing capability easier to understand, combine, verify and reuse.

That capability is everywhere. It may be a sophisticated UWB team in Shenzhen. It may be a T736 computing module and a supplier’s development work. It may be an old printing shop in Kwun Tong that can make a precise mark or a jig before a distant engineering team has finished a meeting. It may be an embroidery machine that can place conductive thread. It may be a local worker who knows how to handle a difficult handover in a hotel. It may be a small workshop with a laser marker, a camera, a roll of copper foil, a stack of packaging material, a sewing machine, or one reliable way to make one thing repeatedly.

The challenge is not to pretend that all of these things are the same. The challenge is to give each contribution a clear boundary: what it can do, under what conditions, who has authority, how another team requests it, and what evidence proves the result.

This is why the white paper begins with stone and thread.

Stone represents material that already exists - sometimes valuable, sometimes discarded, sometimes poorly described. Thread represents familiar skill and ordinary machinery that can acquire a new function when combined with conductive materials, electronics and a defined interface. The credit-card-sized starting point represents a physical scale that people can immediately understand. It is small enough to hold, print, carry, attach and discuss, while the intelligence behind it can remain distributed across local surfaces, electronic modules, edge computing, existing equipment and cloud services.

The white paper then moves outward. It asks how people keep authority when models and tools become more capable; how MCP’s stateless protocol direction can support collaboration without erasing real-world state; how hardware and software teams can begin together before final hardware ships; how one function can be reproduced by multiple makers; how value should be measured without double counting; and how a successful reference result might mature into a shared specification without casually calling itself a standard.

The principle that holds these parts together is simple:

A defined result, at a named place and time, under agreed conditions, completed by a responsible party and verified by evidence.

I believe this matters economically, technically and socially.

Economically, it can shorten the distance between a GenAI plan and something a customer can use or buy. Technically, it lets different specialists work against the same boundary instead of waiting for one monolithic product team. Socially, it creates a path for people and small organizations to contribute useful functions without first becoming software platforms or fully automated factories.

That is the meaning of “No One Left Behind in the GenAI Physical Era” in this document. It is not a claim that technology automatically creates inclusion. It is a design requirement: participation should be understandable; authority should be visible; contribution should be attributable; evidence should travel with the result; and value should not disappear into a system that the contributor cannot inspect.

My role as Founder of HSIMC is not to claim that every piece is already solved. It is to help define the interfaces, bring complementary capability together, insist on evidence boundaries, and keep moving from plan to tested result.

The invitation is therefore deliberately practical:

Bring your existing capability.
Deliver one reproducible function.
Measure one useful outcome.
Make GenAI plans become real value - now, here, globally. - Benny (Yan) Liao
Founder, HSIMC
12 November 2026

Executive Summary

Generative AI can now produce designs, code, instructions, analyses and alternatives at extraordinary speed. The physical world moves differently. A useful result still depends on materials, machines, tools, people, space, time, permission, safety, manufacturing limits and a clear way to decide whether the result is acceptable.

OpenAI's 3 November 2026 publication *Accelerating engineering for the physical world* describes advances across mechanical design, electronics, software, simulation and manufacturing, and explicitly invites engineers, engineering firms and software vendors to help shape this work around real tools and real problems. HSIMC responds with an independent proposal: make existing real-world capabilities easier to discover, combine, verify and reuse across teams, networks and locations.

The basic unit is deliberately simple: **a defined result, at a named place and time, under agreed conditions, completed by a responsible party and verified by evidence.** That result may be a sample, a processed material, an inspected part, a stored item, a completed handover, a released software build, or a controlled physical action. The objective is not to automate everything. It is to shorten the path from a GenAI plan to a useful and accountable change in reality.

The white paper starts with what people already have. The planar footprint of a standard credit card is used as a familiar physical reference, not as proof that a complete computer has already been compressed into a bank-card form factor. Intelligence can be distributed across a printed or textile surface, a small electronic module, an edge computing node, an existing machine and cloud services. Different contributors can therefore deliver different functions against the same interface and acceptance conditions.

This is why familiar materials and tools matter. Construction-material streams can be explored for new component uses. Conductive thread, sewing and embroidery can become controlled placement processes when combined with appropriate electronics and test. Printing can produce registration marks, labels, electrodes, interfaces and functional surfaces. A local workshop can contribute a jig; an engineer can contribute an algorithm; another team can contribute a camera adapter, a UWB profile or a test procedure. The contribution should be small enough to understand and strong enough to reproduce: **one function, with clear inputs and outputs, conditions, responsibility, version and acceptance evidence.**

HSIMC organizes collaboration around **Resource, Authority and Evidence**. A Resource is the real-world person, place, machine, material, inventory position, service or other capability being discussed. Authority describes who may discover, read, contribute, reserve, configure, execute, stop, accept, release or administer a capability. Evidence describes what was observed, tested, delivered, accepted or rejected, with enough version and provenance information for another party to understand the result.

These business concepts are distinct from Model Context Protocol terminology. In MCP, tools are callable operations and resources provide data or context to a client. The 28 July 2026 specification moved MCP toward a stateless protocol core in which requests are self-describing and application state that must persist is referenced explicitly rather than hidden inside a transport session. HSIMC adopts the engineering lesson without pretending that the physical world has no state. Equipment status, Work reservations, custody, Authority, operation receipts, payment and Evidence remain persistent.

This distinction enables a useful principle: **stateless request; persistent real-world responsibility.** A new engineer can join from another client, device or network without relying on hidden conversational memory, while the physical work remains tied to explicit identities and records. A retry should not create a second physical effect. A historical operation receipt should not be mistaken for a fresh observation. A public QR or NFC entry should not silently grant device-control authority. A cancellation request should not be reported as a physical stop until the physical executor confirms the stop.

The current shortest hardware route uses the T736 module and the engineering assets already accumulated around it. The route is: **cloud simulation and application development -> parallel hardware and interface work -> remote reference-board validation -> configuration freeze -> shipment -> local binding, recovery and continued operation.** This allows software, hardware, process and local making to begin in parallel against common contracts.

The engineering evidence boundary remains explicit. C0-005 is a documentary hardware baseline. C0-007 is a source-verified simulator. C0-008 records 14/14 Azure remote simulation checks from 10 November 2026. The physical T736 was not reachable during that recorded test, and H0-001 did not perform a new physical probe. Card-sized hardware remains a design target. These records support engineering continuity; they do not constitute physical-board acceptance or a finished-product claim.

A proposed first shared function is sample alignment and inspection. One team can produce a versioned printed sample or fixture, another can develop image processing against fixed inputs, a later adapter can bind a real camera and T736 target, and a second site can reproduce the same function. The product metric is not only whether one demo works, but how quickly another team can reproduce it and how much hidden knowledge must be clarified.

The same collaboration method can be used across **Make, Move and Assist** while keeping physical safety cases separate. A small factory may need a low-frequency transfer of material from storage to a machine, then to a temporary holding location and a third-party logistics handover. Airport and hotel concepts can separate RoomReady, BagReady and GuestReady instead of treating delivery as a single event. Empty trolleys can be studied as a space-clearance problem. Assisted mobility can reuse some interfaces and evidence methods, but a cargo result never validates a wheelchair or other person-carrying application.

Value is measured through separate ledgers rather than one inflated number. Cash results record realized revenue, verified avoided cash expense and incremental operating expense. Released-capability results record time, throughput, waiting, setup and the ability to complete additional work. User and service outcomes record timing, choice, access and other service effects. Environmental and material results record actual flows, energy and life-cycle boundaries. Governance results record Authority, acceptance, custody, evidence quality, disputes and recovery. Network-level accounting then removes internal transfers so that a payment from one participant to another is not counted twice as newly created value.

The 17 United Nations Sustainable Development Goals provide an outcome lens, not a badge collection. The strongest initial relevance for this white paper is to skills and learning, decent work and small enterprise, inclusive industry and innovation, reduced inequalities, responsible consumption and production, and partnerships. Other Goals become relevant only in the applications where the necessary evidence exists. **SDG relevance is not SDG achievement. ESG intention is not ESG performance. Evidence from one Work is not evidence for the whole platform.**

“Leave No One Behind” becomes an operating question: who can participate, who cannot, why, who has Authority, who contributes unpaid effort, who receives cash value, who carries new risk, and whether alternative participation paths exist for people without the preferred language, device, account or skill.

Finally, the white paper explains how reference work can mature toward shared specifications. One successful demo is not a standard. The maturity path is: **observation -> reference method -> versioned profile -> conformance method -> independent implementations -> shared specification -> formal standardization, when justified.** Formal standards require an appropriate standards process and multi-stakeholder consensus; HSIMC reference methods and compatibility profiles should not be labelled as international standards simply because they are useful.

The network can begin with one small contribution. A participant may bring a material sample, a sewing or printing process, a machine hour, a fixture, a camera image, an adapter, a local test site, an algorithm, an inspection method or an accepted service window. People without APIs can still participate. The minimum useful contribution can be one function, one sample, one process note, one acceptance method, one Authority statement and actual time/cost evidence.

The invitation is:

Bring your existing capability.
Deliver one reproducible function.
Measure one useful outcome.
Make GenAI plans become real value - now, here, globally.

How to Read This White Paper

This white paper is designed for several kinds of readers who may enter at different points.

If you are a business or operations owner, begin with Part I and Part V. They explain the gap between a digital plan and a real result, the Work unit used throughout the paper, and the way cash, released capacity, ESG effects and distributional outcomes should be measured.

If you are a maker, small manufacturer or local service provider, begin with Part II and Appendix G. They show how ordinary materials, machines and skills can be described as bounded capabilities without requiring a full software platform or factory-wide automation project.

If you are a software, AI or MCP engineer, begin with Part III and Part IV. They distinguish Human Authority from protocol authorization, HSIMC Resource from MCP resource, fresh observations from operation receipts, and stateless protocol behavior from persistent business and physical state.

If you are a hardware engineer or supplier, begin with Part IV and Appendices C-E. They describe common contracts, evidence gates, T736 as the current shortest reference route, replaceable adapters, configuration freeze and second-site reproduction.

If you are a public institution, ESG team, university or development partner, begin with Part V, Appendix A and Appendix B. They show how all 17 SDGs can be mapped without claiming achievement where evidence does not exist, and how environmental, social and governance evidence can be kept separate from marketing language.

If you are considering joining the network, begin with Part VI and Appendix G. You do not need to join with a complete product. One real task or one reproducible capability is enough to start.

Evidence labels used throughout the book are defined in Appendix F. Version and publication status are recorded near the end of the book so that a reader can distinguish a proposal, a simulation, a remote simulation, a physical reference result, a delivered configuration and an operational result.

Unified Table of Contents

Front Matter - Founder’s Foreword - Start with What Is Already Here - Executive Summary - How to Read This White Paper

Part I - The Opportunity: Bringing Intelligence into Everyday Making - Chapter 1 - The Distance Between an Answer and an Outcome - Chapter 2 - A Credit-Card-Sized Starting Point - Chapter 3 - Leaving No One Behind in the GenAI Physical Era

Part II - The Starting Materials: What People Already Know and Have - Chapter 4 - A Second Life for Stone - Chapter 5 - From Thread to Interface - Chapter 6 - From Craft to Reproducible Capability

Part III - The Collaboration Model: People, Tools, Resources and Evidence - Chapter 7 - People Hold Authority - Chapter 8 - Tools, Resources and Shared Contracts - Chapter 9 - Stateless Requests, Accountable Outcomes

Part IV - The Engineering Path: Develop Together, Verify in Reality - Chapter 10 - Build Together Before Hardware Ships - Chapter 11 - The T736 Reference Path - Chapter 12 - One Function, Many Makers

Part V - The Value: Useful Outcomes and Measurable Participation - Chapter 13 - Real-World Journeys: Make, Move and Assist - Chapter 14 - Measuring Value without Double Counting - Chapter 15 - Seventeen Goals, Specific Contributions

Part VI - The Invitation: A Global Network of Existing Capabilities - Chapter 16 - An Open Network of Complementary Capabilities - Chapter 17 - From Reference Results to Shared Standards - Chapter 18 - An Invitation to Make AI Real

Appendices - Appendix A - The 17 SDGs: Relevance, Contribution, Evidence and Limits - Appendix B - ESG Measurement Framework - Appendix C - Common Capability Contract - Appendix D - Physical Interface and Material Profiles - Appendix E - Reference Engineering and Acceptance Roadmap - Appendix F - Evidence Register and Source Notes - Appendix G - Participation Guide - Appendix H - Multilingual Editorial Rules - Glossary - Bibliography and Source Register - Version and Evidence Register

Part I - The Opportunity: Bringing Intelligence into Everyday Making

Part I at a glance

Generative AI can produce plans, code, drawings, analyses and alternatives at a speed that is increasingly difficult for the physical world to match. A material still has to be found. A machine still has to be available. A fixture has to fit. A person must have authority to act. A result must be accepted, repaired, handed over, paid for or rejected. The opportunity described in this white paper is to make those physical capabilities easier to discover, combine and verify without pretending that the real world has become frictionless.

Part I establishes three propositions.

First, the distance between an answer and an outcome is not a temporary inconvenience. It is a design space. Engineering, manufacturing and service work become more valuable when AI can work with explicit physical constraints, responsible parties and evidence rather than with abstract plans alone.

Second, a familiar credit-card-sized footprint can serve as a practical starting point for participation. It does not mean that every function must be compressed into one card. It means that a small, portable and understandable entry can connect printed matter, NFC or QR, local electronics, a computing node, an existing machine and cloud services. Intelligence can be distributed across layers while the interface remains understandable.

Third, broad participation requires more than access to technology. It requires authority, meaningful contribution and a fair way to recognize value. The United Nations promise to leave no one behind provides an important human principle: identify who is excluded, understand why, enable meaningful participation and keep accountability visible. HSIMC applies that principle to a new question for the GenAI physical era: **who is able to help turn plans into useful capability, under what conditions, and who shares in the value created?**

The chapters that follow are therefore not a product specification. They are the foundation for a collaboration model that will be tested through materials, making, engineering, logistics and real-world execution.

Chapter 1 - The Distance Between an Answer and an Outcome

1.1 The physical work behind a digital answer

A model can answer a question in seconds. It can propose a circuit, create code, compare materials, generate a drawing, suggest a process plan or produce a set of tests. Yet none of these outputs is, by itself, a physical result.

A useful physical result still depends on conditions that are often missing from the prompt: - the actual material grade, thickness, moisture, finish or damage; - the real machine, controller, firmware and tooling available at a location; - the dimensions and tolerances of a fixture or connector; - the skill and availability of the person who will perform the operation; - the space, access window and safety rules at the site; - the authority to use the resource, modify the object or move the asset; - the current status of inventory, power, network or transport; - the method by which the result will be inspected and accepted.

This distance between an answer and an outcome is easy to underestimate because digital work hides much of the physical effort. A PDF can contain thousands of pages before anyone has handled a sheet of paper. Thousands of QR codes can be generated before anyone has cut, laminated, punched or attached a card. A PCB can be drawn before a fabricator has confirmed stack-up, copper weight, drill limits or component availability. A warehouse plan can be optimized before anyone has confirmed that the aisle is clear or the pallet is structurally sound.

The gap is not evidence that AI is failing. It is evidence that the next engineering problem is different.

The problem is no longer only, **“Can a model produce a plausible plan?”** It is increasingly, **“Can a plan be connected to real resources, real authority, real constraints and a verifiable outcome?”**

That is the opportunity HSIMC addresses.

A result is more than an action

This white paper uses a deliberately concrete unit of work:

A defined result, at a named place and time, under agreed conditions, completed by a responsible party and verified by evidence.

A result might be: - a sample cut to an agreed geometry; - a printed marker placed within a defined alignment tolerance; - a conductive textile pattern whose electrical properties are measured; - a material transferred from a storage location to a machine; - a finished item placed into an identified temporary location; - a baggage handover completed within an accepted time window; - an empty trolley moved away from a constrained passenger area; - an inspection completed and recorded against a versioned method.

The definition is intentionally neutral about whether the work is performed by a person, a machine, a robot, a service provider or a combination of them. Automation is one possible capability. It is not the definition of value.

The value lies in the accepted change of state.

1.2 The November engineering invitation

On 3 November 2026, OpenAI published *Accelerating engineering for the physical world*, describing advances in how AI models can support mechanical design, electrical engineering, software, simulation and end-to-end system design. The publication emphasizes that engineers must bring together specialized tools, manufacturing requirements and changes across disciplines, and it invites engineers, engineering organizations and software vendors to help shape the work around real problems and real tools.

HSIMC treats that publication as an important engineering signal, not as an endorsement of HSIMC or of the proposals in this white paper.

The complementary opportunity is clear. Frontier models can become increasingly capable at design and analysis, while many of the resources needed to make a result real remain local, fragmented and difficult to describe in a form another team can safely use.

A small workshop may have: - a laser cutter but no standardized capability profile; - a sewing or embroidery machine but no agreed electronic test method; - a printer capable of producing a functional mark but no versioned alignment procedure; - a CNC machine whose operator knows the true limits better than its public brochure; - a box of reusable labels or materials whose current state is known only to one person; - a useful jig that exists only as a physical object; - a reliable local supplier that does not expose an API.

A global engineering team may have excellent code, simulation or analysis but no fast path to those capabilities.

The opportunity is therefore not only to make AI better at engineering. It is also to make engineering capability more legible to AI and to other people.

From hidden capability to a usable contribution

For a capability to become useful across teams, at least five questions must be answerable:

1. **What can it do?**
   The function, input, output and operating conditions.
2. **Where and when is it available?**
   The place, time window, preparation requirements and current status.
3. **Who is responsible?**
   The person or organization that can accept the work and respond when conditions change.
4. **Who has authority?**
   The party allowed to view data, reserve capacity, operate equipment, change settings, release results or cancel work.
5. **How will the result be checked?**
   The evidence, test, measurement, inspection or handover that turns an activity into an accepted outcome.

These questions are as relevant to a person with a sewing machine as they are to a large automated factory.

This is why HSIMC's working model is built around **Resource, Authority and Evidence**. - **Resource** describes what exists and under what conditions it may be useful. - **Authority** describes who may decide, operate, share, accept or withdraw. - **Evidence** describes what was actually observed, tested, delivered or rejected.

The model is deliberately compatible with imperfect reality. A resource can be registered but unavailable. A machine can be capable in principle but unverified in a particular configuration. A file can exist while its physical implementation

remains untested. A human expert can be available only during a narrow time window. A local supplier can contribute through a structured handoff even if it has no sophisticated software stack.

That flexibility matters because much of the physical world does not begin as a clean digital twin.

1.3 The first design principle: make the next step executable

The most useful AI output is often not a complete master plan. It is the next step that another party can actually perform.

Consider a simple example: a printed sample must be aligned and inspected.

A global model can propose an alignment pattern. A local printer can produce the sample. A fixture maker can define the mechanical reference. A vision engineer can write a locate_marks function. A camera integrator can connect a real device. A second workshop can reproduce the test. None of those contributors needs to own the entire system.

The collaboration becomes valuable when their outputs meet through a common contract: - known input image or real camera frame; - known pattern version; - coordinate system and units; - expected output fields; - acceptable error cases; - calibration status; - software version; - evidence file or receipt.

The first workshop may deliver only the printed target and fixture. That is enough if the contribution is understandable and reproducible.

This is the deeper meaning of the project invitation:

Bring your existing capability. Deliver one reproducible function.

It lowers the threshold for useful participation while increasing the quality of the handoff.

Why "now" matters

The physical world has always contained distributed capability. What changes in the GenAI era is the speed at which design, explanation, code, translation and coordination can be produced around that capability.

A local workshop no longer needs to prepare a complete international product catalogue before a remote engineer can understand its contribution. A model can help translate technical notes, generate interface examples, structure a test plan and compare alternatives. A remote team can begin with a simulator while hardware is still being prepared. Evidence can be returned in a form another team can inspect.

This does not eliminate local work. It raises the value of local work that is clearly described.

The competitive variable begins to shift from **"Who can manufacture the cheapest unit?"** toward a richer question:

"Who can provide the right capability, at the right place and time, with a reliable handoff and rapid iteration?"

Low cost remains important in many markets. But for prototypes, repairs, customization, field integration, urgent replacement, small batches and local service, speed and coordination can be economically decisive.

This is especially visible in dense industrial environments, where mature local facilities may already have printing, laser processing, assembly, packaging and repair capabilities. The opportunity is to connect them to new forms of engineering demand without requiring each participant to become a full-stack technology company.

1.4 What this white paper does not claim

A credible physical-world initiative must say what is not yet proven.

This white paper does not claim that: - every local workshop can produce electronics safely; - every recycled material is environmentally preferable; - every sewing or printing process can become an electronic manufacturing process; - every QR or NFC interaction authorizes physical control; - UWB positioning automatically makes motion safe; - a simulator result proves the behavior of a real device; - a successful cargo movement validates a passenger-carrying system; - a credit-card-sized concept is already a complete card-thickness computer; - a global network exists simply because a website lists potential resources.

Instead, it proposes a method for moving from **proposed -> documented -> simulated -> measured -> independently reproduced -> commercially accepted**, while preserving the evidence level at each step.

That evidence discipline is essential to the opportunity. It allows ambition and engineering credibility to coexist.

Chapter 2 - A Credit-Card-Sized Starting Point

2.1 A familiar footprint and an explicit design target

A standard credit card is one of the most familiar engineered objects in daily life. The ISO/IEC 7810 ID-1 format is nominally 85.60 mm by 53.98 mm. The full standard also addresses physical characteristics beyond planar size, so this white paper uses the footprint as a **design reference**, not as a claim that every proposed device already meets the complete card standard.

Why begin there?

Because a familiar physical scale is powerful.

People can immediately imagine: - carrying it; - attaching it to another object; - placing it in a wallet, package or fixture; - printing on it; - adding a QR code; - integrating an NFC inlay; - stacking it with other layers; - using it as an identity and service entry point.

A credit-card-sized starting point turns an abstract infrastructure discussion into something a person can hold.

But the design principle is not **“put everything inside one card.”**

It is **“make the entry small enough to understand, while allowing the capability behind it to be distributed.”**

That distinction prevents the form factor from becoming a gimmick.

2.2 The card is an entry, not the whole machine

The physical stack can be understood as several layers.

Layer A - the visible entry

This may be a printed card, label, tag, patch, plate or textile surface. It can carry: - human-readable identity; - QR entry; - NFC entry; - alignment marks; - instructions; - a product or resource reference; - a version or batch marker.

It can be manufactured with mature local processes and does not need to contain the full computing system.

Layer B - the functional surface or interface

This may add: - conductive thread; - printed electrodes; - a simple antenna; - a flexible conductor; - a sensor contact; - a connector or breakout; - a mechanical locating feature.

The purpose is not to romanticize craft. The purpose is to determine whether a familiar process can deliver a measurable function.

For example, conductive embroidery can form an antenna geometry, but the result still requires a compatible chip or module, electrical matching and testing. Printing can produce a marker or electrode, but its electrical or optical performance still depends on material, geometry and process control.

The value of the maker’s contribution becomes clearer when the function is separated from the rest of the system.

Layer C - the electronic module

Existing electronic modules can provide functions that a local maker should not need to reinvent: - NFC or RFID interface; - UWB ranging or positioning; - camera or sensor input; - USB or serial conversion; - secure storage; - wireless connectivity; - power conditioning.

A good architecture makes these modules replaceable and testable rather than permanently entangling them with a single craft process.

Layer D - the local computing node

The current HSIMC engineering route uses the T736 family as the shortest reference path for a portable computing and physical-interface node. The existing project evidence supports a documentary hardware baseline and a simulator-based engineering route; it does not yet prove a finished credit-card-sized product.

The intended role of the node is broader than “a phone.” It may operate as: - a portable team node; - a desktop local server; - an embedded interface node; - a gateway between existing equipment and cloud services; - a local evidence collector; - a runtime for adapters and constrained functions.

Its value comes from the common interface and engineering workflow, not from the label attached to the enclosure.

Layer E - existing equipment and cloud capability

The card or local node may connect to: - an existing printer; - a laser machine; - a sewing or embroidery machine; - a camera; - a barcode or RFID reader; - a UWB module; - a machine controller; - a warehouse location; - a remote engineering workspace; - a model or service in the cloud.

This is where a small physical entry can unlock a much larger capability.

2.3 Why distributed intelligence matters

The physical world is full of mature machines that were never designed for modern GenAI workflows.

Replacing all of them would be economically and environmentally unrealistic.

A more practical strategy is to separate: - the stable physical capability; - the replaceable interface; - the local compute and control boundary; - the remote model or workflow; - the human authority to use them.

This allows new intelligence to be added incrementally.

A 20-year-old machine may continue performing the motion it already performs well. A new adapter can expose a small set of safe operations. A T736-class node can translate data, hold local state or connect a camera. A remote model can help generate settings or inspect results. A human operator remains responsible for the site-specific authority and exception handling.

The architecture therefore treats **existing things as potential capability**, not as obsolete infrastructure waiting to be replaced.

This is particularly important for small enterprises and local makers. Their competitive advantage is often not a proprietary digital platform. It is proximity, practical experience, equipment access, responsiveness and the ability to improvise within known constraints.

The challenge is to turn that tacit capability into something another team can rely on.

2.4 From stone to a useful component

“Stone” in the title is not a claim that construction waste can directly become electronics. It represents the most basic question of local making:

What useful function can an available material perform, after we understand its source, properties and processing conditions?

A material stream from construction or manufacturing may contain mineral, polymer, wood, metal or mixed fractions. Each requires its own characterization and appropriate processing route. A responsible project does not label an item “recycled” and assume the environmental case is complete.

The useful workflow is:

identify the source and composition as far as practical;

define the intended function;

process a sample using a versioned method;

measure relevant performance;

record yield, waste, energy, labour and cost;

compare against the baseline alternative;

decide whether the material is suitable for that function.

A material may be useful as: - a fixture body; - a non-electrical structural plate; - a display or information surface; - a protective component; - a sample substrate; - packaging; - a temporary jig; - a composite ingredient.

The important point is not the number of possible uses. It is that each use can become an explicit capability with evidence.

That is how an “ordinary” material becomes part of an intelligent production network.

2.5 From thread to a functional layer

“Thread” represents a different path: the transformation of a familiar human skill into a measurable technical contribution.

Sewing and embroidery are highly distributed capabilities. They exist in large factories, small workshops, schools, homes and community settings. Conductive thread makes it possible to explore electrical patterns using equipment and skills that are already widely available.

The contribution might be: - an embroidered antenna geometry; - a conductive path; - a wearable contact; - a stitched sensor interface; - a repeatable placement method for a small electronic module.

But the white paper deliberately separates **making the pattern** from **proving the electronic performance**.

A maker may be able to reproduce stitch density and geometry with high consistency. An RF engineer may be needed to characterize impedance and antenna performance. An electronics partner may supply the compatible chip. Another team may test durability after bending or washing.

This is exactly the kind of multi-party engineering that GenAI can help coordinate.

The network does not require one person to know everything. It requires a shared definition of the result and a reliable handoff between contributors.

2.6 The card as a contract with the physical world

The most important function of the card may ultimately be neither computing nor identification.

It may be **coordination**.

A physical card or tag can point to a persistent resource identity. From there, an authorized user or agent may learn: - what the object is; - who is responsible for it; - which capabilities are available; - which documents are current; - what evidence exists; - what actions can be requested; - what permissions are required; - what work is already in progress.

The physical entry is simple. The responsibility behind it is explicit.

This is why QR and NFC are useful but not sufficient. They solve discovery and proximity. They do not, by themselves, solve authority, freshness, safety or accountability.

The card therefore becomes the **front door** to a larger capability, not the capability itself.

Chapter 3 - Leaving No One Behind in the GenAI Physical Era

3.1 A human principle for a technical transition

The United Nations 2030 Agenda describes 17 Sustainable Development Goals and 169 targets as integrated and indivisible, and includes the pledge that no one will be left behind and that the furthest behind should be reached first. United Nations development guidance further frames Leave No One Behind as a central transformative promise, emphasizing the need to identify who is excluded and why, enable meaningful participation and maintain accountability.

This white paper does not claim that HSIMC represents the United Nations or that every project activity advances every SDG.

Instead, it adopts a practical question inspired by that principle:

As GenAI enters the physical economy, who gains the ability to participate - and who may be left outside the new production system?

The question matters because technological transitions can concentrate capability even while lowering the cost of information.

A person may have access to a powerful AI model but still lack: - a machine; - a safe workplace; - a recognized identity in a supply network; - capital for tooling; - language access; - reliable connectivity; - a way to prove quality; - a channel to receive paid work; - authority over the data or equipment involved.

Conversely, another person may have valuable physical capability but no easy way to expose it to a global engineering workflow.

A seamstress may understand repeatable textile production but never describe the capability as an engineering service. A retired technician may know a machine better than its documentation. A small print shop may be able to produce a same-day functional sample that is more valuable to an engineer than a cheaper overseas batch arriving weeks later. A local operator may know that a pallet, corridor or fixture is unsafe even when the digital plan says otherwise.

These people are not “non-technical.” They hold forms of capability that the digital system often fails to recognize.

3.2 Participation is not the same as access

Giving someone a login does not create meaningful participation.

Meaningful participation requires at least four things.

1. The person can understand the work

The task must be expressed in a language and form appropriate to the participant. Technical detail should be available when needed, but the first interface should make the expected result understandable.

This is one reason the project uses layered publication: card, brief, summary and full white paper. Different people need different depths of explanation.

The same principle applies to engineering work. A machine operator needs the acceptance condition for the operation. A hardware engineer needs electrical limits. A software engineer needs the function contract. A customer needs to know the result, timing, price and responsibility.

2. The person can decide whether to participate

Authority must remain visible.

A resource being discoverable does not imply that it may be used. A person being listed does not imply that they have accepted a job. A machine being technically capable does not imply that remote control is permitted. A document being readable does not imply authority to change the physical system.

This is why HSIMC separates **Resource** from **Authority**.

3. The contribution can be recognized and verified

A contribution should leave evidence that can be attributed to the responsible party and version.

That evidence may be: - a sample; - a measurement; - a photograph; - a machine log; - a code commit; - a test report; - a handover receipt; - an inspection result; - a signed acceptance; - a reproducible procedure.

Recognition matters because invisible work is easily undervalued.

4. The participant can share in the value

Participation without fair value distribution can simply create a new form of extraction.

The project therefore distinguishes: - paid work; - donated or sponsored contribution; - research participation; - in-kind materials; - shared infrastructure; - intellectual contribution; - future commercial rights.

The terms do not need to be identical in every project. They need to be explicit enough for participants to understand what they are contributing and what they may receive.

3.3 Stateless collaboration, persistent responsibility

The 28 July 2026 MCP specification introduced a stateless protocol core. In practical terms, each request can carry the information necessary to process it rather than relying on hidden transport-session state. If an application needs state across requests, it can use an explicit handle or identifier.

This is a powerful foundation for global engineering collaboration.

An engineer can reconnect from another client. A request can be routed through scalable infrastructure. Tool and resource catalogs can be discovered and cached. Workflows do not need to depend on a single fragile long-lived session.

But stateless transport does not remove real-world responsibility.

A physical system still needs persistent records for: - resource identity; - equipment state; - reservations; - work orders; - custody; - operation receipts; - authorization; - evidence; - payment; - exceptions and recovery.

The design principle is therefore:

Stateless request; persistent real-world responsibility.

This separation helps inclusion because participation is no longer tied to one hidden conversation or one client session. A new team can pick up a clearly identified work item, subject to authority, and understand what has already happened.

It also protects participants. A historical receipt should not be presented as a fresh observation. A retry should not create a second physical action. An expired permission should not be silently reused. A public resource description should not grant device-control rights.

This is how a scalable protocol can support accountability rather than weaken it.

3.4 Human Authority · Local Capability · Global Participation

HSIMC summarizes its participation model in three phrases.

Human Authority

People and organizations retain the right to define the permitted use of their resources, data, equipment and work.

Human Authority does not mean that every action requires manual clicking. It means that automation operates inside an explicitly delegated boundary and that responsibility can be traced.

Local Capability

Useful physical capability exists in specific places.

A machine, room, operator, material batch or delivery service has a location, availability window, condition and owner. Global AI cannot erase those facts. It can help coordinate them.

The local nature of physical capability is not a limitation. It is the source of much of its value.

Global Participation

Design, analysis, translation, software, simulation, quality review and business collaboration can come from anywhere, provided that contributors work against a common interface and evidence boundary.

Global participation does not require all data to be public. It requires a way to expose the right information to the right participant under the right authority.

Together, these principles support a network in which a small contributor can participate without pretending to be a global factory, and a global team can use local capability without pretending that physical execution is merely an API call.

3.5 The SDGs as a lens, not a marketing badge

The 17 Sustainable Development Goals provide a broad framework for considering social, economic and environmental outcomes. This white paper will map all 17 goals, but relevance will differ by work item.

The strongest early connections are expected to include: - **SDG 4 - Quality Education:** learning through real, reproducible engineering work; - **SDG 8 - Decent Work and Economic Growth:** paid participation, small-enterprise capability and productive work; - **SDG 9 - Industry, Innovation and Infrastructure:** enabling smaller participants to connect to modern engineering and production systems; - **SDG 10 - Reduced Inequalities:** reducing barriers created by geography, language, capital, age or access to infrastructure; - **SDG 12 - Responsible Consumption and Production:** repair, reuse, material evidence and lifecycle awareness; - **SDG 17 - Partnerships for the Goals:** practical multi-party cooperation around technology, knowledge and implementation.

A responsible SDG mapping requires a claim boundary.

For every proposed contribution, the project should ask:

What changed compared with the baseline?

Who benefited?

Who carried additional cost or risk?

What evidence supports the result?

What remains unknown?

This prevents the SDGs from becoming decorative icons attached to ordinary business activity.

3.6 ESG as an operating discipline

ESG is useful in this project when it shapes decisions, not when it is added after the work is finished.

Environmental

The project should measure relevant material and energy changes, including: - new material introduced; - waste avoided or created; - repair and reuse; - process energy; - transport; - device and battery lifecycle; - end-of-life handling.

A recycled input is not automatically a lower-impact outcome. The full process matters.

Social

The project should observe: - who can participate; - who receives paid work; - whether new skills are learned; - whether physical burden or waiting is reduced; - whether access improves for people with different needs; - whether workers gain or lose agency; - whether the system shifts risk onto less powerful participants.

Governance

The project should preserve: - resource ownership or responsibility; - authority boundaries; - version history; - acceptance evidence; - payment and contribution records; - escalation and correction paths; - privacy and data boundaries.

The same evidence that helps engineering reliability can therefore help ESG credibility.

3.7 The smallest meaningful global action

A global ecosystem can sound impossibly large. The project deliberately defines a much smaller starting action.

A participant may bring: - one material sample; - one manufacturing process; - one machine hour; - one jig; - one camera image; - one software adapter; - one test method; - one local site; - one logistics window; - one real customer need.

The participant then describes: - what can be provided; - where and when; - who is responsible; - what use is permitted; - how the result can be checked.

That is enough to start.

The network becomes real when a second party can use the contribution, produce an accepted result and return evidence.

The next level is independent reproduction at another place.

The level after that is repeatable economic value.

This creates a ladder of participation that does not require every contributor to become a software company, own expensive infrastructure or understand the whole system before they can add value.

3.8 Part I conclusion - from inclusion to capability

The opportunity described in Part I is larger than hardware and smaller than ideology.

It is larger than hardware because the central problem is not a single card, computer, robot or protocol. It is the coordination of people, tools, materials, machines, authority and evidence around useful physical outcomes.

It is smaller than ideology because the proposal can be tested one function at a time.

Can a local process produce a sample that another team can verify?

Can a remote engineer work before final hardware ships?

Can an existing machine gain a useful new interface without being replaced?

Can a person with an ordinary tool contribute a measurable function?

Can the value, cost and responsibility be understood by everyone involved?

If the answers become yes, the idea gains substance.

Part II will move from the opportunity to the starting materials themselves: **stone, thread, printing and other familiar processes**, and the engineering discipline required to turn them from interesting possibilities into reproducible capability.

Evidence boundary for Part I

This is an **independent HSIMC white-paper proposal**. It does not claim endorsement by the United Nations, OpenAI, the Model Context Protocol project, ISO or any named ecosystem partner.

Current inherited engineering evidence remains unchanged: - **C0-005:** documentary T736 hardware baseline; - **C0-007:** source-verified simulator; - **C0-008:** historical Azure remote simulation record, 14/14 on 10 November 2026; - **H0-001:** no new physical T736 probe; the physical reference board remained unbound in that handoff; - the credit-card-sized hardware remains a design target rather than a finished-product claim.

Exploratory mobility, materials and local-manufacturing examples are used to define questions and collaboration patterns. They are not presented as validated product specifications unless a specific evidence level is stated.

Selected source notes

OpenAI, *Accelerating engineering for the physical world*, 3 November 2026. https://openai.com/index/accelerating-engineering-for-the-physical-world/

Model Context Protocol, *The 2026-07-28 Specification*. https://blog.modelcontextprotocol.io/posts/2026-07-28/

Model Context Protocol, specification dated 28 July 2026. https://modelcontextprotocol.io/specification/2026-07-28/

United Nations, *Transforming our world: the 2030 Agenda for Sustainable Development*. https://sdgs.un.org/2030agenda

United Nations Sustainable Development Group, *Leave No One Behind*. https://unsdg.un.org/2030-agenda/universal-values/leave-no-one-behind

ISO/IEC 7810, Identification cards - Physical characteristics; ID-1 nominal dimensions are 85.60 mm x 53.98 mm. https://www.iso.org/obp/ui/#iso:std:iso-iec:7810:ed-3:v1:en

HSIMC, *Card Phone and Cloud Development - Cross-Project Hand Off H0-001*, 11 November 2026.

HSIMC, *White Paper Executive Summary W0-004*, 11 November 2026.

Part II - The Starting Materials: What People Already Know and Have

Part II at a glance

Part I argued that the distance between an AI-generated answer and a physical outcome is not a temporary inconvenience. It is the place where materials, tools, people, authority, time and evidence must be brought together.

Part II begins with a practical question: **what if participation starts not from a new factory, but from what people already know and already have?**

A small workshop may have stone offcuts, packaging stock, a sewing machine, an embroidery machine, a printer, a laser marker, copper foil, hand tools, a phone, a camera, a jig, a box of old tags, or simply a reliable way of making one thing repeatedly. None of those assets, by itself, is “GenAI manufacturing.” Yet each can become part of an AI-enabled production network when its function, conditions, interfaces and evidence are made explicit.

This Part develops three propositions.

First, a material should not be valued only by its original market category. A discarded mineral panel, a piece of construction material or an offcut can acquire a second life if its source, condition, processing limits and new function are understood. The same principle applies to paper-like mineral composites, recycled plastics, metals and other material streams. Reuse claims must be tied to actual material identity and measured performance rather than to attractive labels.

Second, familiar craft processes can become functional engineering processes without ceasing to be craft. Thread can carry current. Embroidery can position conductive paths. Printing can place conductive, dielectric or visual layers. Copper foil can become a quick prototype conductor. A sewing or printing process does not become an electronics factory merely because conductive material is added; the electronic function still depends on geometry, resistance, matching, bonding, insulation, curing, durability and test.

Third, the real transition is from **craft to reproducible capability**. A useful contribution is not simply “I can sew,” “I have a printer,” or “I can cut this material.” It is a versioned function that another person or system can request, understand, inspect and reproduce under declared conditions.

That transition is the foundation for local making and global participation.

Chapter 4 - A Second Life for Stone

4.1 From a material source to a defined use

“Stone” in this white paper is deliberately broad. It can mean natural stone, mineral-filled sheet, construction offcut, concrete fragment, ceramic piece, tile, mineral composite, or another material whose dominant character is inorganic and whose original use may have ended.

The point is not to collapse these materials into one category. The point is to show how a local material becomes useful when its identity and next function are described well enough for another person to act.

Construction and demolition waste illustrates the scale of the opportunity. The European Commission notes that this waste stream includes concrete, bricks, wood, glass, metals and plastics, and accounts for more than a third of waste generated in the EU. It also emphasizes that some fractions have high resource value while others can be reprocessed into new products, provided hazardous substances and poor sorting do not undermine reuse and recycling.

This is an important distinction for HSIMC: **a discarded material is not automatically a resource. It becomes a resource only when its condition, risk and possible use are understood.**

A useful material record therefore starts with facts such as: - source and previous use; - approximate composition or product identity, when known; - dimensions, thickness and mass; - visible damage, contamination, moisture or surface treatment; - whether the material has been cut, painted, laminated or chemically treated; - who owns it and who may authorize reuse; - where it is stored and when it is available; - what tests have been performed; - what uses are explicitly not yet validated.

A photograph can help, but a photograph is not a material specification. A product name can help, but a product name is not a current condition report. A previous building use can help, but it does not prove suitability for a new load, electrical, hygienic or fire-related function.

The value of a Resource record is therefore not that it turns uncertainty into certainty. It makes uncertainty visible.

A local example: from discarded surface to functional sample

Consider a small mineral offcut recovered from a renovation site. The most useful first question may not be, “Can this become a new commercial product?” It may be, “Can this become a stable sample base, jig plate, mounting surface, sign, spacer or fixture for a defined test?”

A modest second use can be economically and environmentally meaningful if it displaces a new purchase and if the preparation effort is low enough. The same piece may be unsuitable for a structural use but perfectly adequate as a non-structural test fixture.

The capability is then not “recycling stone.” It is something more precise, for example:

Prepare a flat mineral sample base, 85.60 x 53.98 mm nominal footprint, within an agreed thickness and edge condition, suitable for a specified non-structural test.

That statement creates a pathway to evidence:

identify the source material;

inspect for contamination and defects;

cut or machine the sample;

measure final dimensions;

record mass, thickness and surface condition;

perform the intended test;

retain a versioned result.

A maker with simple cutting and finishing tools can contribute to this workflow. A remote engineer can define geometry and acceptance criteria. A local camera or measurement tool can generate evidence. A second workshop can reproduce the same sample using a different material and compare results.

The intelligence is not inside the stone. The intelligence is in the **shared definition of what the material is being asked to do, and the evidence that it did it.**

4.2 Stone paper, composites and construction waste: distinct pathways

The phrase “stone paper” is attractive because it appears to connect mineral content with a familiar paper-like product. It is also a source of confusion.

Commercial mineral-filled paper-like sheets are typically engineered composites with a defined formulation and manufacturing process. Construction mineral waste is a heterogeneous waste stream. Natural stone is another category again. These pathways should not be silently merged.

For HSIMC, the correct approach is to treat each as a separate Resource and process profile.

Path A: existing commercial mineral-filled sheet

If a commercially produced mineral-filled sheet is available, the useful questions include: - declared composition and manufacturer data; - thickness and dimensional stability; - surface energy and printability; - heat sensitivity; - laser, blade or punch compatibility; - adhesion of inks, laminates, NFC inlays or labels; - fold, tear and abrasion behavior; - recycling or disposal route in the actual market.

A design can then test whether the material is useful for a card, label, sleeve, gift insert, jig, packaging component or another defined function.

Path B: locally formulated mineral composite

If the objective is to formulate a new sheet or molded part from mineral powder and binder, that is a materials-development project. It requires a declared recipe, particle characteristics, binder system, mixing and forming process, curing conditions, mechanical testing and repeatability data.

The attractive phrase “from stone to card” is not enough. A formulation that works once in a hand sample is not yet a manufacturing capability.

The minimum useful evidence is closer to: - input material identity and batch; - mass fractions or formulation range; - processing method and equipment; - temperature, pressure, time and other relevant process variables; - sample thickness and density; - dimensional change after processing; - strength, flexibility, abrasion or other property relevant to the intended function; - failure modes; - number of samples and variation observed.

Path C: direct reuse of construction material

A third pathway avoids making a new composite at all. A recovered tile, panel, stone offcut or other piece can be cut and reused directly.

This may be the fastest local route because it minimizes transformation. It also places greater importance on inspection, sorting and adaptation to the irregularity of the source material.

The environmental logic is different in each pathway. Reusing an existing offcut may avoid new material production but require transport and cutting. Producing a new composite may consume binders, energy and finishing steps. A commercial mineral-filled sheet may have different upstream impacts and end-of-life options.

Therefore this white paper does not claim that one route is universally “greener.” It proposes that the comparison should be measured for the specific function.

A practical comparison record

For a card-sized sample, a simple comparison could record:

Field

Reused mineral offcut

Commercial mineral-filled sheet

New local composite

Source

recovered local material

purchased product

locally prepared ingredients

Preparation

cut / grind / clean

cut / print / laminate

mix / form / cure / finish

New material added

low to moderate

product already manufactured

binder + additives + energy

Dimensional consistency

source-dependent

usually high

process-dependent

Surface for printing

may require treatment

product-dependent

formulation-dependent

Evidence needed

condition + geometry + use test

datasheet + process test

recipe + process + property test

Current status

must be verified

must be verified

R&D until repeatable

This table is intentionally incomplete. It is a decision framework, not a material certification.

4.3 Measure yield, durability, cost and environmental effects

A material experiment becomes economically useful when it begins to answer four questions:

Yield - how much usable output is obtained from the available input?

Durability - does the result remain fit for the intended function for the required time?

Cost - what does one accepted output actually cost, including failed attempts and labor?

Environmental effect - what new material, energy, transport and waste are introduced or avoided?

Yield

Yield should be defined against an acceptance rule. If ten recovered pieces are collected and only four can produce a flat, clean sample of the required size, the yield is not ten. The rejected pieces are part of the process evidence.

For local making, yield can be more informative than material price. A “free” waste stream with high sorting, cleaning and failure cost may be more expensive than a purchased sheet. Conversely, a local offcut that requires almost no processing may be an excellent resource even if it is available only in small quantities.

Durability

Durability must be tied to the intended use. A sample used for one laboratory alignment test has a different requirement from a card intended to survive years in a wallet, a gift item handled by customers, or a fixture exposed to heat and vibration.

Useful tests can be modest but explicit: - repeated insertion or removal; - bend or flex cycles; - surface abrasion; - moisture exposure; - adhesion of printed or bonded layers; - dimensional change; - edge damage after cutting; - read performance of an attached QR, NFC or RFID element.

The test method and number of cycles matter. “Looks durable” is not an engineering result.

Full cost per accepted result

A useful formula is:

Full cost per accepted result = total material + labor + machine + energy + tooling + inspection + failed-output cost, divided by accepted quantity.

This is especially important in small-batch local manufacturing. The unit price of material may be almost irrelevant compared with setup, alignment and rework.

The same logic helps explain why speed can become more important than low unit cost. A local workshop may be more expensive per sheet but capable of producing a revised sample today, allowing a global engineering team to continue work immediately. The economic value lies in iteration time, not only in material price.

Environmental effects without automatic claims

Environmental claims should be made against a defined baseline.

If a recovered material replaces a newly purchased board, useful questions include: - what new material was actually avoided; - how far the recovered material was transported; - what cleaning, cutting or coating was required; - what waste was generated by preparation; - whether the new use extends life or merely delays disposal; - whether the new combination makes later separation harder; - whether any hazardous or unknown constituent is introduced into a new context.

This is consistent with a broader circular-economy principle: recovery has value, but quality, sorting and safe handling matter.

From environmental story to reproducible evidence

The strongest sustainability story is often not a dramatic one. It may simply be:

This batch used 12 locally recovered offcuts. Eight met the sample criteria. Four were rejected. The accepted samples replaced eight newly purchased boards for this test series. Cutting consumed a recorded amount of machine time and produced a measured mass of waste.

That statement can improve over time. It can be compared with other materials. It can be audited. It can inform the next decision.

It is more useful than an unsupported claim that a product is “zero waste” or “green.”

Chapter 5 - From Thread to Interface

5.1 Conductive thread and textile antennas

A sewing machine is one of the most familiar production tools in the world. It is also an instructive example of how an existing craft capability can acquire a new engineering role.

Ordinary thread is used to join, reinforce and decorate. Conductive thread can also carry electrical current or form part of an antenna, electrode or sensor. Embroidery can place that thread into a repeatable geometry.

Research on embroidered NFC antennas shows the potential and the constraints clearly. A University of Manchester study fabricated NFC antenna prototypes on cotton using conductive yarns coated with metals including silver and stainless steel. The work compared embroidery approaches and measured substantial differences in DC resistance. It also matched an embroidered coil to a specific NFC microchip. The result was not “any metal thread becomes NFC.” The result was that a defined thread, stitch geometry, antenna shape and chip can be engineered together and measured.

That distinction is central to this white paper.

The sewing machine contributes placement capability

The sewing or embroidery machine does not need to understand NFC protocol. Its contribution can be narrower:

Place a specified conductive thread along a versioned path on a specified textile substrate, with declared stitch parameters and inspection evidence.

Another contributor may provide the antenna design. Another may select the chip. Another may model the RF behavior. Another may test read range. Another may design the mechanical attachment to a cable, gift, garment or card sleeve.

This is how an everyday machine becomes part of a global engineering workflow without pretending to perform the entire system function.

The important variables

For a conductive textile function, useful variables include: - thread material, coating and diameter; - measured resistance per unit length; - substrate type, thickness and stretch; - stitch type, pitch, tension and number of passes; - geometry and spacing between conductive paths; - connection method to a component or connector; - insulation and strain relief; - bending, washing, abrasion or moisture requirements; - antenna or sensor tuning, when relevant; - inspection and acceptance method.

A remote engineer can specify these variables. A local maker can report what the machine can actually achieve. The design can then converge through evidence rather than assumption.

Decorative function and electronic function can coexist

Embroidery also demonstrates that electronic function does not require industrial visual language. Research has shown conductive textile antennas integrated into decorative embroidered logos, with conductive thread positioned on the reverse while colored yarn creates the visible design.

For HSIMC, this suggests a valuable design principle:

A functional interface can remain culturally, aesthetically and commercially local.

A hotel gift, team uniform, event badge or cable accessory does not have to look like a laboratory prototype merely because it contains a digital entry point or functional conductor.

That matters for participation. Local producers can contribute not only manufacturing capacity but also material knowledge, aesthetics, cultural meaning and customer familiarity.

5.2 Printed marks, electrodes and functional surfaces

Printing is another mature process that can become part of electronics and sensing when the deposited material has a functional property.

Printed electronics literature describes conductive, semiconducting and dielectric inks deposited on substrates through methods such as screen printing and inkjet printing, often followed by curing or sintering. The performance depends not only on the ink but also on substrate roughness, porosity, wetting, drying and post-processing.

This provides a useful bridge between traditional printing districts and electronic prototyping.

Start with the simplest functional layer

The first functional contribution does not need to be a transistor or a complete circuit. It may be: - a high-contrast visual alignment mark; - a conductive pad; - a capacitive touch electrode; - an antenna trace; - a resistive element; - a connector footprint; - a dielectric separator; - a printed calibration target; - a QR code linked to a Resource identity; - a registration layer that allows a later process to align accurately.

Each function has a different acceptance test.

A QR code can be validated by optical decoding. A conductive trace can be checked for resistance and continuity. A capacitive electrode can be tested with a known controller. An antenna requires tuning and read-performance testing. A registration mark requires geometric measurement.

The local printer does not have to become responsible for the entire final product. It must be responsible for the agreed printed result.

The substrate is part of the process

A design that prints well on coated PET may fail on rough paper, textile or mineral sheet. A conductive ink may penetrate a porous substrate or crack when flexed. A curing temperature suitable for one material may deform another.

Therefore a printed functional layer should always be linked to a substrate profile.

At minimum: - material identity and thickness; - surface treatment, coating or laminate; - print process and equipment; - ink or paste identity and batch; - mesh, nozzle, resolution or other process parameter as applicable; - cure or drying condition; - finished geometry; - electrical or optical test; - environmental and handling limits.

This is where local printing knowledge becomes valuable. Operators often know, through practice, which substrate feeds reliably, which coating causes adhesion problems, how humidity changes drying, and how registration drifts on a particular machine. GenAI can help structure and compare those observations, but the local process still produces the evidence.

Copper foil, film and fast iteration

Copper foil provides another accessible prototyping route. A conductive pattern can be cut, etched, laminated or otherwise formed without committing to a conventional multilayer PCB process.

The value is speed and accessibility. A local team may be able to produce a one-off antenna, electrode, jumper or low-complexity circuit pattern immediately.

The limitation is equally important. A foil prototype is not automatically equivalent to a production PCB. Adhesion, minimum feature size, insulation, oxidation, current carrying capacity, bending, connector reliability, RF behavior and manufacturing tolerance may all differ.

The correct workflow is therefore:

prototype quickly -> measure -> learn -> decide whether to keep the local method or move to a more controlled manufacturing process.

The result of an experiment is not a commitment to a manufacturing technology. It is evidence that narrows the next decision.

5.3 Joining craft, electronic modules and computing

The most powerful pattern in this chapter is separation of roles.

A functional object can be built from layers contributed by different people:

material layer - textile, paper, mineral sheet, plastic, metal or another substrate;

functional surface - printed conductor, embroidered thread, foil trace, marker or electrode;

electronic module - NFC chip, sensor, reader, controller, radio or interface board;

compute node - phone, T736 node, local computer or another runtime;

cloud and model layer - design, analysis, planning, versioning and collaboration;

human authority - the person or organization allowed to use, modify, share or operate the result;

evidence layer - the measurements and records that prove what was actually built and tested.

The advantage of this architecture is that innovation does not require every contributor to master every layer.

A textile workshop can contribute repeatable placement. A printer can contribute functional marks. An electronics company can provide a module. A software team can provide a tool. A local integrator can connect it to an existing machine. A global AI team can help compare designs and generate test plans.

The shared interface is what allows these contributions to meet.

A credit-card-sized example

Imagine a card-sized object whose visible surface is locally printed, whose reverse side contains an embroidered or foil antenna, whose electronic function is provided by a small NFC module, and whose deeper capability is reached through a QR/NFC URI connected to an HSIMC Resource.

The card itself may have no processor.

In another version, the card may connect to a small interface module.

In a later version, a T736-based compute node may provide local processing and physical interfaces.

All three can belong to the same capability family without pretending to be the same hardware.

That is the value of a layered architecture: **the visible object can stay simple while capability evolves behind it.**

A gift can become an entry into capability

This also changes the role of everyday products.

A cable, tag, card sleeve, textile patch, package or gift can carry: - a durable identifier; - a QR/NFC entry; - material and batch information; - a service or configuration link; - a warranty or maintenance action; - a local capability description; - a route to a human or AI support function.

The physical product remains useful on its own. The digital layer does not have to dominate the object. It can simply create a continuing relationship between the object, the maker, the user and the next service.

This is especially relevant to small producers. They do not need to build a global platform. They can produce one well-defined physical layer and connect it to a broader system under agreed authority.

Chapter 6 - From Craft to Reproducible Capability

6.1 A function, an interface and a repeatable result

The most important transformation in Part II is not from “old material” to “smart material.” It is from an informal skill to a capability another party can safely use.

A craft capability often lives in a person’s memory: - “I know how to make this cut cleanly.” - “This machine can handle that fabric if I change the tension.” - “This ink works on this coating.” - “The laser marks well if the focus is slightly changed.” - “This supplier can make the part quickly.”

These statements are valuable. They are also difficult to scale because the conditions are implicit.

A reproducible capability makes the useful parts explicit.

A minimal capability description can include:

Field

Question

Function

What result can be produced?

Input

What material, file, object or state is required?

Output

What physical or digital result is returned?

Conditions

What machine, tool, environment or preparation is required?

Interface

How is the request submitted and how is the result identified?

Authority

Who may request, approve, modify or release the work?

Evidence

What measurement or inspection proves completion?

Version

Which method, design and software were used?

Cost and time

What resources were actually consumed?

Limits

What remains unsupported or unverified?

This description can exist as a human-readable sheet, a machine-readable profile, or both.

The goal is not bureaucracy. The goal is to reduce the amount of hidden knowledge that must be rediscovered every time a new collaborator arrives.

Reproducible does not mean identical everywhere

A second workshop may use a different machine. A different city may have different material. A local operator may choose another fixture.

Reproducibility therefore does not require that every implementation be identical. It requires that the differences are visible and that the result can be checked against the same functional expectation.

For example, two workshops may produce an NFC textile antenna using different embroidery machines. If both record thread, geometry, stitch settings, resistance and read-performance tests, their results can be compared.

The common object is the **function contract**, not the brand of machine.

6.2 Skills, samples, process conditions and acceptance

A practical contribution can be organized into four linked objects:

1. Skill

The human or organizational capability: cutting, printing, sewing, machining, assembling, testing, coding, calibrating, repairing or coordinating.

The skill record should not claim universal competence. It can be narrow and useful.

Example:

“Can embroider conductive thread on cotton and polyester substrates using a specified vector path; can provide resistance measurement and sample photographs; RF tuning is performed by another party.”

That is more trustworthy than “smart textile expert” if the actual contribution is embroidery and measurement.

2. Sample

A sample is evidence that a process has produced a result under known conditions.

A useful sample record includes: - sample ID; - design version; - input batch; - machine and tool; - operator or responsible party; - date and location; - process settings; - inspection result; - photographs or measurement files; - deviations and failures.

A sample can then become part of a growing evidence library.

3. Process condition

Process conditions make the sample interpretable.

Without them, a successful sample may be impossible to reproduce.

For local manufacturing, even a simple record can be powerful: - machine model; - tool or needle type; - feed rate or speed; - temperature or cure time; - fixture version; - environmental condition if relevant; - operator note about a manual adjustment.

GenAI can help turn free-form notes into structured profiles, compare versions and identify missing variables. It should not invent a value that was never measured.

4. Acceptance

Acceptance is what turns “something was made” into “the requested result was delivered.”

Acceptance can be: - dimensional measurement; - continuity test; - resistance range; - QR decode; - NFC read; - image alignment score; - visual inspection against a reference; - successful mechanical fit; - customer sign-off; - a combination of tests.

Different applications need different acceptance rules.

The important point is that acceptance is agreed before scale-up whenever possible.

6.3 Participation across making and use

When capability is described this way, participation becomes broader without becoming vague.

A family workshop can participate by producing a tested textile pattern. A print shop can participate by making a registered conductive or visual layer. A student can participate by writing an inspection tool. A retiree with deep process knowledge can contribute a fixture method or troubleshooting guide. A large supplier can contribute a certified component. A cloud engineer can contribute a simulation. A logistics provider can contribute an accepted transport window.

They do not need to become one company.

They need a way to make their contribution understandable, authorized and verifiable.

Participation is not free labor

A global participation model must distinguish between learning, volunteering, research contribution, paid work, commercial supply and intellectual property licensing.

The same physical action can have very different economic and ethical meaning depending on the agreement.

A contributor should be able to know: - what is being requested; - whether the work is paid, donated, experimental or speculative; - how the result may be reused; - whether the contributor will be named; - who owns the new design or data; - who carries the cost of failure or rework; - when payment is due; - how the contributor can stop participating.

This is where Human Authority becomes a practical economic principle.

Participation should create capability, not dependence

A useful project should leave participants with something they can use again: - a tested process; - a better fixture; - a reusable template; - a machine profile; - a documented skill; - a new customer relationship; - a measurement method; - a software tool; - a reference sample; - a clearer understanding of cost and limits.

The objective is not to make every person a hardware engineer. It is to let people contribute at the level where they already have competence, while providing pathways to learn more.

From local proof to global participation

A local capability becomes globally useful through a sequence:

discover - someone identifies an existing material, skill, tool or process;

describe - the function and limits are recorded;

sample - a small result is produced;

measure - the result is inspected;

version - the method and evidence are frozen;

offer - the capability becomes available under stated conditions;

accept - a responsible party agrees to perform a specific Work;

execute - the work is performed locally;

verify - evidence is returned;

learn - the result improves the next design, process or model.

This loop connects local making to global engineering without requiring the local maker to operate a global software platform.

The first economic question: what changed?

The value of a capability should be measured by the change it creates.

Did the sample arrive faster?

Did the new method reduce failed iterations?

Did the local material avoid a new purchase?

Did a maker gain a new paid task?

Did another team reproduce the result without a long explanation call?

Did a customer receive a better product or service?

Did the process reduce or increase material and energy use?

Did the contributor retain authority over the use of their work?

These questions connect Part II directly to ROI, ESG and the SDGs without turning those frameworks into decoration.

A compact participation package

The white paper proposes a simple first package for a new contributor:

Bring one existing capability. Deliver one reproducible function.

A contribution package can be small: - one function description; - one sample or reference output; - one process note; - one acceptance method; - one statement of authority and reuse terms; - one record of actual time and cost.

That is enough to begin.

It can later connect to a card, NFC/QR entry, HSIMC Resource, machine-readable profile, T736 node, simulator or global collaboration workflow.

The small package is intentional. Participation should not require a small workshop to build a complete digital twin before anyone can work with it.

The bridge to Part III

Part II has focused on materials, familiar tools and the transformation of craft into reproducible capability.

Part III will address the collaboration model that makes those capabilities usable across teams: people who hold authority, tools that perform actions, resources that provide context, contracts that define functions, and evidence that preserves accountability across stateless requests.

The material is local.

The skill is local.

The responsible person is local.

The collaboration can be global.

Source notes for Part II

The following external sources support specific background statements in this Part. They do not certify HSIMC prototypes or future products.

1. European Commission, **Construction and demolition waste** - construction and demolition waste composition, resource recovery potential, sorting and hazardous-material considerations.
   https://environment.ec.europa.eu/topics/waste-and-recycling/construction-and-demolition-waste_en
2. University of Manchester research PDF, **e-Textile Embroidered Wearable Near Field Communication RFID Antennas** - embroidered NFC antenna prototypes, conductive yarns, stitch-method comparison, measured resistance and chip matching.
   https://research.manchester.ac.uk/files/76660584/e_Textile_Embroidered_Wearable_Near_Field_Communication_RFID_Antennas.pdf
3. Gil et al., **Colorful Textile Antennas Integrated into Embroidered Logos**, Journal of Sensor and Actuator Networks - conductive e-thread integrated into embroidered decorative forms.
   https://www.mdpi.com/2224-2708/4/4/371
4. A. Kamyshny et al. and related review literature summarized in **A Review on Printed Electronics: Fabrication Methods, Inks, Substrates, Applications and Environmental Impacts** - functional inks, substrates, printing methods, curing and the importance of ink-substrate interaction and repeatability.
   https://www.mdpi.com/2504-4494/5/3/89
5. Rubio & Bolduc, **Screen Printing for Energy Storage and Functional Electronics: A Review** - screen printing of conductive and functional materials on flexible substrates, with process variables influencing device performance.
   https://www.mdpi.com/2673-3978/6/2/7
6. OpenAI, **Accelerating engineering for the physical world**, 3 November 2026 - background context for AI-supported mechanical, electrical, simulation and manufacturing workflows.
   https://openai.com/index/accelerating-engineering-for-the-physical-world/

Evidence boundary

Part II is an architectural and participation framework. It does **not** claim that: - recovered construction material has been qualified for a structural, food-contact, electrical-safety or fire-rated use; - locally formulated “stone paper” has been developed or validated by HSIMC; - ordinary sewing thread or any arbitrary conductive thread is suitable for NFC; - a sewing or embroidery machine alone can manufacture a complete electronic product; - a printed conductor is equivalent to a production PCB; - any particular conductive ink, foil, textile, substrate, NFC chip or attachment process has been qualified for a commercial release; - a credit-card-sized T736 product has been completed; - environmental benefit has been demonstrated without a defined baseline and measurement.

The purpose of this Part is to define how such claims could become testable, attributable and reproducible through Resource, Authority and Evidence.

Part III - The Collaboration Model: People, Tools, Resources and Evidence

Part III at a glance

Part I defined the opportunity: GenAI can produce plans quickly, but a physical outcome still requires materials, machines, people, time, authority and evidence.

Part II showed that participation can begin with what people already know and already have. A mineral offcut, a sewing machine, an embroidery machine, a printer, copper foil, a camera, a jig or a local workshop can become part of an intelligent production network when a function, its conditions and its evidence are made explicit.

Part III asks the system question: **how can many people, models, tools, resources and organizations work together without turning access into authority, context into control, or a successful request into an unverified claim of completion?**

The answer proposed here has three parts.

First, **people hold authority**. A model can propose, compare and call permitted tools, but the right to read a document, reserve a machine, start a motor, disclose a customer record, accept a commercial obligation or release a payment comes from an accountable person or organization acting within an agreed scope. HSIMC therefore treats authority as a first-class business fact rather than an assumption hidden inside a chat session.

Second, **tools and resources need shared contracts**. MCP tools expose callable operations. MCP resources expose data and context. HSIMC Resources are real-world business objects such as people, places, machines, materials, inventories, work slots and services. These concepts can be connected, but they should not be collapsed into one another. A machine may be represented by an HSIMC Resource, described through MCP resources and controlled through one or more MCP tools. The relationship must be explicit.

Third, **stateless protocol requests do not remove real-world state or responsibility**. MCP 2026-07-28 removes protocol-level sessions and makes each request self-describing. Applications can still maintain state through explicit handles. In a physical workflow, persistent state is unavoidable: a machine can be busy, a pallet can be in transit, a bag can be in custody, a motor can still be moving, an invoice can remain unpaid. HSIMC therefore adopts a simple principle:

Stateless request; persistent real-world responsibility.

The purpose of this Part is not to invent a new replacement for MCP, ERP, MES, WMS, IAM or contract law. It is to define a collaboration layer that lets those systems and real-world participants meet around one verifiable Work.

Chapter 7 - People Hold Authority

7.1 M-Person, organizations and M-Teams

GenAI changes who can plan. It should not silently change who is allowed to decide.

A person may use an AI system to draft a design, inspect a photograph, compare suppliers or prepare an operation request. A team may use several models and tools to work in parallel. An organization may expose selected capabilities through an API or MCP server. None of those facts, by itself, answers a more basic question: **who has the right to authorize the next real-world action?**

HSIMC uses the term **M-Person** as a participation concept: a real person who can work with models, tools and teams while retaining authority over the rights that belong to that person or the organization they legitimately represent. An **M-Team** can include people, models, software services and specialized agents, but it is not itself a magical legal principal. It acts through defined authority relationships.

This matters because the same person can occupy different roles in different Works.

A factory owner may be a buyer in one Work and a capability provider in another. A designer may contribute a drawing but have no permission to start a laser. A hotel may control the room and custody process but not the airline baggage system. An engineer may be permitted to query a T736 reference board but not deploy a new image to a delivered unit. A founder may be able to authorize one company directly but need another director or co-holder for a different company.

The architecture therefore separates: - **identity** - who or what is being referred to; - **role** - what function the participant has in this Work; - **authority** - what decisions or operations the participant is permitted to make; - **scope** - which Resource, time, location and operation the authority covers; - **evidence** - what proves the authority was granted and what happened under it.

A single global profile can be useful for discovery, but authority should be evaluated at the point where a decision matters.

Identity is continuity, not permission

HSIMC preserves existing Runtime and Resource identities wherever possible. A card, QR code, NFC tag or short URL can help a person find the same object again. This continuity is valuable. It does not make the identifier an authorization token.

A printed QR on a machine might open its public capability page. That does not mean anyone who scans it may move the machine. An NFC tag on a luggage item might identify the service record. It does not mean anyone with a phone may change the delivery destination. A Card ID can connect a physical object to a digital profile. It does not replace the identity of the actual T736 device, its operator or its owner.

The distinction is simple:

Identity answers “what is this?” Authority answers “who may do what with it, now?”

This is especially important in a global system where the same public Resource may be discovered by many models and users.

M-Teams should make responsibility clearer, not blur it

A strong M-Team can divide work across specialties: - one model interprets a customer request; - another checks a manufacturing constraint; - a human operator confirms a material condition; - a vision tool measures alignment; - a supplier accepts a machining task; - a local controller executes deterministic motion; - an evidence service records the result.

The advantage is parallel capability. The risk is diffusion of responsibility.

For every Work, HSIMC therefore asks for named responsibility points. At minimum:

Requester - who wants the outcome;

Resource authority - who can commit the relevant machine, material, space, data or service;

Executor - who or what performs the action;

Acceptor - who decides the result meets the agreed condition;

Economic principal - who pays or receives value when relevant;

Evidence custodian - where the authoritative result record is retained.

One person may hold several of these roles. A model may assist one or more roles. The record should still distinguish them.

7.2 Document access, device control and contribution rights

A recurring mistake in connected systems is to treat all access as one permission.

It is not.

Reading a PDF, editing a design file, querying a sensor, reserving a machine and commanding physical motion are different actions with different risks. They should have different permissions even when they appear in the same user interface.

A useful permission ladder is:

Permission class	Example	Typical consequence
Discover	See that a capability exists	low consequence; may still expose commercial information
Read	Read a datasheet, drawing or current observation	information disclosure
Contribute	Upload a model, design, sample result or process note	changes shared knowledge
Reserve	Hold a machine slot, inventory quantity or service window	affects availability to others
Configure	Change parameters, mappings or profiles	may alter later execution
Execute	Start a machine, write a tag, move a cart	physical or transactional side effect
Stop / recover	Stop motion, cancel or recover a process	safety and continuity impact
Accept / release	Accept output, release custody or authorize payment	commercial / governance consequence
Administer	Change identities, permissions or policy	system-wide consequence

This ladder is not a universal standard. It is a practical design tool. A real implementation can map these categories to existing IAM, OAuth scopes, operating procedures or local control roles.

Document authorization is not device authorization

The T736 handoff illustrates the distinction. An engineer may be authorized to receive a versioned hardware profile in OneDrive or another workspace. That delivery can be proven by file identity, version, hash and readback. It does not imply the engineer may control the reference board.

Similarly, a remote client can be given permission to run a simulator without being granted a physical target. The H0-001 handoff explicitly preserves this separation: document destinations and hardware control are different permissions, and a simulated execution target must not be silently promoted to a physical one.

This separation improves collaboration because information can move earlier than control.

A software engineer can begin writing against a common contract. A hardware engineer can design the interface. A local technician can prepare a fixture. The physical actuator can remain unbound until the authority and safety case are ready.

Contribution rights matter as much as access rights

A global capability network also needs to answer a different question: **what happens to the contribution itself?**

If a small workshop develops a better jig, a student writes an adapter, an engineer creates a calibration file, or a supplier provides a tested antenna geometry, the system should not assume unrestricted reuse.

A contribution record can include: - contributor identity or organization; - Work and Resource to which the contribution relates; - version and hash; - permitted use - evaluation, one project, one customer, internal use, resale, open reuse or another agreed scope; - attribution requirement, if any; - payment or revenue arrangement, if any; - confidentiality classification; - expiry or withdrawal conditions when applicable; - evidence showing what was actually used in a delivered result.

This is not a substitute for a legal agreement. It is a way to make the operational facts legible enough that legal and commercial agreements can be executed with less ambiguity.

Least authority, explicit step-up

The safest default is to grant the smallest useful authority and expand it when the Work requires more.

A user may start with public discovery. Reading a private drawing requires authenticated access. Reserving a machine requires a stronger business relationship. Starting motion may require a named operator, location check, current safety state and an execution token or equivalent local authorization.

MCP 2026-07-28 strengthens protocol authorization and routing, including issuer validation and credential isolation. Those mechanisms help answer whether a client is authorized to call a protected protocol endpoint. HSIMC still needs the application-level decision about whether the requested real-world action is appropriate for that Resource, time and Work.

The protocol can carry the request. The business layer defines the authority.

7.3 Responsibility, payment, redress and withdrawal

Authority is incomplete without accountability.

A collaborative system should record not only who was allowed to act, but also what happens when work is late, incomplete, unsafe, disputed or no longer wanted.

Acceptance is a separate act

An operation can finish without being accepted.

A laser can complete a marking cycle and still produce the wrong serial number. A sewing machine can complete a conductive path that fails resistance testing. A delivery can arrive at a hotel but before the guest is ready. A model can generate a valid file that uses the wrong design revision.

Therefore a useful Work lifecycle separates:

requested -> accepted for work -> executing -> execution complete -> evidence available -> accepted / rejected / partially accepted

This allows the system to distinguish “the machine stopped” from “the customer got the agreed result.”

Payment should follow a defined economic event

A global network becomes more inclusive only if contribution can translate into understandable economic value.

Possible payment triggers include: - reservation accepted; - material supplied; - sample delivered; - test passed; - service completed; - customer accepted; - usage measured; - revenue received.

The correct trigger depends on the Work. It should be agreed before execution where possible.

The evidence record should also distinguish cash value from non-cash contribution. A workshop may donate a sample for a pilot. A student may contribute work under an educational arrangement. A supplier may loan a development board. A

company may sponsor testing. Those contributions still have value and should be recorded without pretending they are revenue.

Redress and correction

“No One Left Behind” requires more than access. Participants need a way to challenge an incorrect record or harmful outcome.

Examples include: - a maker says the system attributed a sample to the wrong contributor; - an operator says a task was assigned outside the agreed time window; - a customer disputes an acceptance result; - a person wants a private data link revoked; - a device owner reports that a capability profile is outdated; - a supplier says a model was reused outside the agreed scope.

The system should preserve the original evidence while allowing a correction, dispute or superseding decision to be added. Silent rewriting of historical evidence is usually worse than a visible correction trail.

Withdrawal does not reverse physics

A person can withdraw permission for future use. That does not always reverse an action already completed.

A file may already have been sent to an authorized recipient. Material may already have been cut. A package may already be in transit. A motor may already have moved a cart. A payment may already have settled.

Withdrawal therefore needs temporal scope: - stop new requests; - stop future reuse; - cancel reservations not yet executed; - request stop or recovery for actions in progress; - preserve evidence of what already occurred; - apply return, deletion, recall or compensation rules where those are actually possible and agreed.

This is another reason persistent real-world state matters even in a stateless protocol architecture.

Chapter 8 - Tools, Resources and Shared Contracts

8.1 Tools for actions; resources for context

MCP provides a useful vocabulary for connecting models to external systems, but the vocabulary must be used precisely.

In MCP, **tools** expose callable operations. A tool has a name, description and input schema, and may return structured or unstructured output. The server can list available tools and the client can call a selected tool.

In MCP, **resources** expose data and context, identified through URIs. They can represent text or binary content and can be listed or read by clients according to the server’s capabilities and authorization model.

HSIMC uses **Resource** in a different, broader business sense: a real or organizational object that can participate in Work. Examples include: - a person or team; - a room, factory, hotel, farm or airport zone; - a T736 device; - a sewing machine or laser marker; - a material batch; - an inventory position; - a pallet, wheelchair or luggage cart; - a UWB anchor; - a delivery service; - a time slot; - a verified engineering capability.

These concepts connect naturally, but they are not synonyms.

One real Resource can have many digital surfaces

Consider a UV laser marker in a local workshop.

The **HSIMC Resource** is the actual machine instance with an owner, location, configuration, maintenance state and authority relationship.

MCP **resources** might expose: - operating manual; - supported material profile; - current calibration record; - current job queue; - latest camera alignment image; - approved logo or serial-number specification.

MCP **tools** might expose: - get_status; - prepare_job; - locate_marks; - submit_marking_job; - get_operation; - request_stop.

A shared **Work contract** connects the business request to these interfaces:

Mark this versioned identifier on this batch, using this machine profile, within this time window, under this authority, and return this evidence.

The machine is not the tool. The manual is not the machine. The URI is not the authority. The Work is not the transport session.

Keeping these layers separate is what allows the system to scale.

8.2 MCP resources and HSIMC Resource objects

A global collaboration network needs both **context discovery** and **real-world identity**.

MCP resources are well suited to context discovery. They give a model or client a structured way to obtain the documents or data needed for a task. HSIMC Resource objects are used to maintain continuity across physical and organizational changes.

A Resource object should survive interface changes

A machine may first be controlled manually. Later, a USB adapter is added. Later still, a T736 node provides an MCP interface. The real machine should not receive a new identity every time its software interface changes.

Likewise, a pallet can acquire a QR label, then an NFC tag, then an RFID tag. The tag technology can change while the Resource remains the same physical object or business object.

A useful Resource profile can include: - stable Resource ID; - human-readable name; - organization or owner relationship; - location or location reference; - type and model where applicable; - interfaces currently available; - capability profiles; - current evidence level; - authority bindings; - last verified observation time; - versioned references to manuals, drawings or test reports; - known limitations.

The profile should distinguish **possible capability** from **observed capability**.

A supplier datasheet may state that a module supports a bus. That is documentary evidence. A reference board may show that the bus is wired a certain way. That is another evidence layer. A real device test demonstrating communication with a specified peripheral is stronger evidence for that instance.

The HSIMC T736 handoff already uses this discipline. C0-005 is a document hardware baseline. C0-007 is a verified-source simulator. C0-008 records a remote simulation run. None of those records is silently rewritten as a current physical-board capability.

Fresh observation versus historical receipt

A collaboration system also needs to distinguish two questions:

What happened in operation X?

What is true now?

The answer to the first should be stable. The answer to the second may require a new observation.

For example, an operation receipt may correctly say that a tag was read at 09:00. Replaying the same receipt at 14:00 does not prove the tag is still present. A machine may have reported idle before another team reserved it. A UWB coordinate can become stale in seconds.

HSIMC therefore distinguishes: - **operation receipt** - immutable or append-only evidence about a specific operation; - **observation** - a time-stamped measurement or status assertion; - **effective state** - application logic derived from observations, reservations, accepted Works and known transitions; - **freshness requirement** - how recent the observation must be for the next action.

This distinction prevents idempotency from turning into stale-state reuse.

8.3 Function contracts across hardware and software

A capability becomes globally usable when another team can understand its contract without depending on the original conversation.

A shared function contract should answer at least: - **name** - what function is requested; - **purpose** - what useful result it produces; - **inputs** - parameters, files, Resource IDs and units; - **outputs** - result fields, units and evidence references; - **preconditions** - physical, electrical, authority and freshness conditions; - **side effects** - what may change in the real world; - **exclusive resources** - what cannot be used concurrently; - **timeouts** - when the caller should stop waiting or change strategy; - **idempotency** - how retries avoid duplicate side effects; - **errors** - machine-readable failure reasons; - **recovery** - what to do after uncertain or partial execution; - **version** - contract, adapter, firmware, model and process versions; - **acceptance** - how the business result is verified.

A contract can exist before the final hardware

This is one of the most important benefits of the HSIMC cloud-development route.

Suppose a team is developing a sample-alignment function. Before the T736 camera path is ready, the contract can already specify: - image input format; - marker definition; - output coordinates; - confidence or quality field; - calibration reference; - error cases; - version and evidence output.

A cloud simulator can operate on fixed images. A remote engineer can implement the algorithm. A local maker can print the marker and build a jig. A hardware engineer can design the camera adapter. When the physical board becomes available, the real camera replaces the synthetic input behind the same business-facing function, subject to new device validation.

The contract does not eliminate hardware differences. It gives those differences a place to be recorded.

Effects need explicit identity

For operations with side effects, a retry should not create a second physical action simply because a network reply was lost.

A practical pattern is to give the intended effect a stable operation identity and bind that identity to the effect definition.

For example: - operation ID: op-123; - target Resource: laser-KT-01; - function: submit_marking_job; - design hash: abc...; - batch: B-42; - quantity: 20; - authority context: current authorized caller and scope.

If the same operation ID is retried with the same intended effect, the server can return the existing receipt or continue the same work. If the operation ID is reused with a different effect, the server should reject the conflict.

The operation identity should not depend on a short-lived access token. Authorization is evaluated on each request; effect identity remains stable across token refresh or client reconnect.

Shared contracts are not universal permissions

Publishing a contract means others can understand how a capability works. It does not mean everyone may call it.

A public contract can coexist with private execution. This is useful for ecosystem development: engineers can build compatible clients and simulators while real devices remain protected by local and organizational authority.

This is another way to let collaboration begin before control is granted.

Chapter 9 - Stateless Requests, Accountable Outcomes

9.1 28 July 2026: context carried by each request

The 28 July 2026 MCP specification is a useful milestone because it changes the protocol’s collaboration model in a way that maps well to global engineering.

The release moves MCP to a **stateless protocol core**. The previous initialize / initialized handshake and protocol-level session ID are removed for the modern protocol. Each request carries the protocol version, client identity and relevant client capabilities. A client can use server/discover when it wants server capabilities in advance, but the protocol no longer depends on a long-lived hidden session for ordinary request meaning.

The practical consequence is significant:

A request can be understood from explicit information carried with the request, rather than from an invisible conversational history attached to one connection.

This makes ordinary web-style routing, scaling and recovery easier. It also aligns with the HSIMC objective that another engineer, device or network can continue work when the required identities, versions, authority and evidence are explicit.

Stateless protocol does not mean stateless application

The MCP maintainers make this distinction directly. If an application needs state across calls, it can mint an explicit handle and have the client pass that handle on later requests.

For a physical system, that approach is not optional. Many facts persist: - Work ID; - Resource reservation; - device instance; - operation ID; - task handle; - material batch; - custody record; - calibration version; - payment state; - acceptance result; - recovery state.

The improvement is that the state is explicit and addressable rather than being hidden inside a transport session.

HSIMC summarizes the design principle as:

Stateless request; persistent real-world responsibility.

Header-based routing can support governance

MCP 2026-07-28 also adds protocol headers such as Mcp-Method and, where relevant, Mcp-Name. Gateways can use this information for routing, metering and policy enforcement without having to infer the operation from a long-lived session.

For HSIMC, this suggests a useful separation: - protocol gateway decides whether a request is structurally valid and authenticated; - application authorization decides whether the principal may request this function on this Resource; - local runtime checks current physical preconditions; - deterministic controller executes the allowed physical action; - evidence layer records the outcome.

No single layer needs to pretend it can decide everything.

9.2 Fresh observations, operation receipts and retries

Stateless systems must be especially disciplined about retries because a client can lose a response while the physical action still occurs.

Imagine a laser job:

client submits the operation;

server accepts it;

laser performs the mark;

network connection fails before the success response reaches the client.

If the client simply submits a new operation, the part may be marked twice.

The correct recovery behavior is:

When the result is uncertain, query the original operation before creating a new effect.

This principle applies to manufacturing, logistics and payments.

Suggested operation states

A practical physical-operation lifecycle can include: - prepared - inputs validated, no physical effect yet; - accepted - authority and resources accepted; - running - physical or transactional execution has started; - input_required - execution needs explicit additional input or confirmation; - cancel_requested - a stop/cancel request has been received; - cancel_confirmed - the responsible executor confirms the action has stopped or will not proceed; - completed - execution finished; - failed - execution ended unsuccessfully with known state; - result_unknown - the system cannot yet prove the final physical state; - accepted_result - business acceptor has accepted the delivered result; - rejected_result - result was inspected and rejected.

These are HSIMC application concepts, not a claim that MCP defines this exact state machine.

Cancellation is not physical stop

A digital cancel request can be acknowledged while the machine is still decelerating, while a courier is still holding the parcel, or after an irreversible step has already happened.

Therefore the system should distinguish: - **request to cancel**; - **executor’s acknowledgement**; - **confirmed physical stop or non-execution**; - **recovery or disposition of partially changed material**.

For safety-related motion, local stop behavior must remain deterministic and independent of a remote model’s continued connectivity.

Freshness belongs in the contract

A stateless request can be self-describing and still use stale information.

Therefore the request should identify freshness expectations where they matter. Examples: - inventory observation no older than 5 minutes; - UWB location no older than 500 milliseconds; - machine safety state read immediately before execution; - guest delivery preference valid until a stated time; - calibration valid for a specific fixture and software version.

If the available observation is older than the threshold, the correct outcome may be refresh_required, not a guessed result.

Multi Round-Trip Requests and explicit confirmation

MCP 2026-07-28 introduces Multi Round-Trip Requests for cases where a tool needs additional input during a stateless interaction. A server can return an input-required result, and the client can resubmit the original call with the requested responses.

This is conceptually useful for high-consequence actions: - confirm the quoted cost before a paid manufacturing run; - ask for a missing dimension before cutting; - request approval before releasing a shipment; - confirm a changed hotel delivery time; - request a local operator acknowledgement before activating a physical reference board.

MRTR is a protocol mechanism. Whether a particular real-world action requires confirmation remains an application and governance decision.

9.3 Persistent responsibility across changing clients and roles

The strongest reason to adopt explicit handles, Resources, Work IDs and evidence is not technical elegance. It is continuity.

Global collaboration means the original client may disappear.

The engineer who designed the part may go offline. The phone used to scan a tag may change. The model session may end. A different service provider may take custody. The hardware may be shipped from Hong Kong to another country. A future engineer may need to reproduce a result months later.

A responsible system should still answer: - What Work was requested? - Which version was used? - Which Resource actually performed it? - Who had authority at the time? - What observations were fresh enough to support the decision? - What physical effect occurred? - What evidence proves it? - Who accepted the result? - What remains uncertain?

That is what makes a collaboration portable across teams.

The Work object as the continuity anchor

HSIMC proposes treating **Work** as the business continuity anchor. A Work is not merely a task record. It binds: - desired result; - requester; - relevant Resources; - authority and acceptance roles; - time and place; - capability / tool contracts; - input versions; - reservations; - operations; - evidence; - economic terms; - acceptance outcome.

A Work can outlive a chat session. It can reference an MCP task when a tool chooses to run asynchronously. It can also exist without using MCP Tasks at all.

This distinction is important because MCP 2026-07-28 moves Tasks into an extension. Protocol support should be detected and used when both sides actually support it. The business Work should not disappear simply because a particular client or server does not implement that extension.

From global planning to local execution

A useful architecture for physical work is:

Global / cloud layer - discover capabilities; - compare plans; - assemble contracts; - check business authorization; - schedule and route Work; - maintain versioned evidence references.

Local runtime layer - bind to the actual device; - read current physical state; - enforce local interlocks and exclusive resources; - translate common contracts to device-specific drivers; - keep operation records through network interruption; - expose result and recovery state.

Deterministic control layer - execute timing-sensitive motor, laser, actuator or safety logic; - stop safely on defined faults; - avoid depending on remote language-model latency for hard real-time safety.

This layered model lets GenAI contribute strongly without asking it to become a safety PLC.

Evidence is the bridge between trust and scale

Global participation fails if every new partner must trust every other partner personally.

Evidence reduces that burden.

A maker can show that a sample was produced under a known process. A remote engineer can reproduce the analysis. A machine can return an operation receipt. A customer can accept or reject against a stated criterion. A second location can repeat the procedure and compare results.

The evidence can include: - hashes and version IDs; - photographs and video; - sensor or measurement files; - calibration references; - machine logs; - material batch records; - read/write verification for NFC or RFID; - signed or authenticated acceptance; - timestamps and location references; - financial receipts where appropriate.

Evidence does not eliminate the need for trust, standards or legal responsibility. It makes trust more transferable.

No One Left Behind requires visible authority and visible value

A global system can be technically open and still exclude people economically or socially.

If only large companies can describe capabilities, small makers remain invisible. If the model can use a person's contribution but the contributor cannot see how it was used, participation becomes extractive. If an operator is expected to absorb every failure but has no authority to reject unsafe work, the system is not meaningfully inclusive.

Part III therefore ends with a governance test for every Work:

Can the participant understand the requested result?

Can the participant say yes, no or propose conditions?

Is the authority scope explicit?

Can the contribution be attributed and evidenced?

Can the participant see the acceptance outcome?

Can economic value or non-cash contribution be recorded accurately?

Can a dispute or correction be added without erasing history?

Can another authorized team reproduce the result from the records?

If the answer is consistently yes, Stateless MCP becomes more than a scalable transport pattern. It becomes part of an accountable collaboration architecture in which models can move quickly while people retain authority over real-world consequences.

Part III conclusion - From connectivity to accountable collaboration

The physical GenAI era does not need another promise that "everything will be connected." Connection is already abundant.

The harder problem is coordination with responsibility.

Part III proposes a simple operating model:

People hold authority. Tools expose actions. Resources provide context. HSIMC Resources preserve real-world identity. Shared contracts define functions. Stateless requests carry explicit context. Persistent Work and Evidence preserve responsibility.

This model allows a small maker, a global engineer, a cloud model, a T736 node, a machine controller, a logistics provider and a customer to participate in the same Work without pretending they have the same role or authority.

It also prepares the next Part of this white paper.

Part IV moves from collaboration architecture into the engineering path: how teams develop together before hardware ships, how the T736 reference route supports cloud-to-physical iteration, and how one function can be reproduced by many makers and devices.

Evidence and source notes for Part III

S1 - Model Context Protocol, 2026-07-28 release. The official MCP release describes the stateless protocol core, removal of the initialize / initialized exchange and protocol-level session ID, self-describing requests, server/discover, header-based routing, cacheable list results, authorization hardening, MRTR, Tasks as an extension and deprecations. It explicitly states that applications can remain stateful through explicit handles.
https://blog.modelcontextprotocol.io/posts/2026-07-28/

S2 - MCP 2026-07-28 release candidate / protocol evolution. The release-candidate documentation describes the stateless redesign, Tasks lifecycle changes, authorization hardening and the principle that application state can be carried by explicit handles rather than transport sessions.
https://blog.modelcontextprotocol.io/posts/2026-07-28-release-candidate/

S3 - MCP SDK migration guidance for 2026-07-28. Current SDK documentation describes per-request modern headers, protocol-version handling, resource subscriptions and authorization requirements for the 2026-07-28 era. It is implementation guidance and does not by itself prove compatibility of any HSIMC client or server. https://ts.sdk.modelcontextprotocol.io/v2/migration/support-2026-07-28

S4 - HSIMC Card Phone & Cloud Development Handoff H0-001, 11 November 2026. The project handoff preserves T736 as the shortest current route, cloud simulation -> remote physical reference board -> freeze -> ship -> local binding, Resource / Authority / Evidence continuity, separate document and hardware-control permissions, and the evidence boundary for C0-005 / C0-007 / C0-008. H0-001 did not re-probe the physical T736.

S5 - HSIMC Stone & Thread White Paper, Part I W0-005 and Part II W0-006. These earlier Parts define the physical-outcome gap, credit-card-sized starting point, Leave No One Behind participation principle, material/function evidence discipline and the transition from craft to reproducible capability.

Evidence discipline

This Part distinguishes: - **MCP protocol facts** from official MCP sources; - **HSIMC architectural proposals** such as M-Person, Work, the permission ladder and proposed physical-operation states; - **existing HSIMC project evidence** from H0-001 and the previous white-paper Parts; - **future implementation requirements** that still need client, server, hardware and real-work validation.

No statement in this Part should be read as proof that a particular external MCP client already supports HSIMC physical control, that a real T736 board is currently reachable, or that any machine is authorized for remote execution.

Part IV - The Engineering Path: Develop Together, Verify in Reality

Part IV at a glance

Parts I-III established the opportunity, the starting materials, and the collaboration model. Part IV asks the engineering question: **how can people in different places begin useful work before the final hardware is finished, while still preserving the distinction between simulation, reference validation and delivered physical evidence?**

The proposed answer is a staged engineering path built around common contracts and replaceable execution targets.

A software engineer should not need to wait for a finished custom board before testing the meaning of an operation. A hardware engineer should not need to wait for the final customer workflow before fixing power, connector and recovery constraints. A local print shop should be able to prepare a marked sample or jig while a remote vision engineer develops against a fixed image set. A reference T736 board can later replace the simulator for selected checks. Only after the intended configuration has been validated should the version be frozen for shipment. After shipment, the delivered instance should continue operating under its own identity, configuration and evidence record.

This is not an argument for pretending that simulation is hardware. It is the opposite. The purpose of parallel development is to make the boundary between each stage explicit enough that teams can work earlier without confusing what has and has not been proved.

The current HSIMC handoff provides a concrete reference. T736 is the shortest hardware route because a real module family, supplier documentation, existing design work and a source-verified simulator already exist. The evidence boundary remains strict: - **C0-005** is a documentary hardware baseline. It contains reviewed dimensions, interface groups, pin information and conflicts that still require instance-level confirmation. - **C0-007** is a source-verified simulator and operation ledger. It is useful for contract development and deterministic replay, not for claiming physical I/O. - **C0-008** records a 14/14 Azure remote simulation run from 10 November 2026. It is a historical remote simulation record, not a physical T736 acceptance report. - The real 736 instance was not re-probed in H0-001. Its present reachability, exact BOM, operating system, drivers, tool directory and Azure-to-board route remain to be bound and verified. - A credit-card-sized complete product remains a design target. The known core module and reference board dimensions do not prove that a complete card-sized product has already been built.

The engineering principle is therefore:

Develop against the contract early. Bind physical capability only when the evidence supports it.

This Part develops that principle in three steps. Chapter 10 defines the parallel development route. Chapter 11 maps the route onto the T736 reference platform. Chapter 12 follows one proposed function, locate_marks, from a versioned sample through simulation, local making, reference-board integration and second-location reproduction.

Chapter 10 - Build Together Before Hardware Ships

10.1 Parallel application, hardware and process development

Traditional hardware projects often force work into a sequence:

specification -> board -> firmware -> application -> test -> customer

That sequence can be appropriate when the interfaces are immature or the safety case requires a tightly controlled design. It can also create long idle periods. Software waits for hardware. Hardware waits for a stable use case. Manufacturing waits for drawings. A customer sees value only after many upstream decisions have already become expensive to change.

GenAI creates an opportunity to reorganize some of this work, but only if the project has a sufficiently clear contract.

The contract does not need to describe every internal implementation detail. It needs to define what another team may rely on. For a function, that usually means: - the name and purpose of the function; - required inputs and accepted formats; - output structure and units; - errors and unsupported cases; - expected side effects; - timing or freshness requirements where relevant; - the Resource or execution target to which the function applies; - the Authority required to invoke it; - version and compatibility information; - evidence needed to accept the result.

Once that boundary is stable enough, several teams can begin in parallel.

Application and model work

A remote application team can build client logic, user flows, test vectors and error handling against a simulator or recorded dataset. The team does not need access to a live actuator in order to decide how a result will be represented, how duplicate requests should be handled, or how an unsupported input should be reported.

For example, a vision function can be developed against a fixed collection of images before the final camera is selected. The result is not yet a physical-camera capability. It is an algorithm and interface result under a documented input condition.

Hardware and interface work

A hardware team can work on power, recovery, connector allocation, USB host budget, serial transceivers, camera attachment, thermal constraints and mechanical packaging. Those questions can progress while the application team works on the common function contract.

The hardware team should publish the result as a Profile: what is documented, what has been measured, which interfaces conflict, what remains optional, and which evidence applies to which board revision.

Local process and fixture work

A local maker can develop the physical sample, print, jig, mount, antenna pattern, label position or fixture while the software and hardware teams work elsewhere. The local maker needs a versioned drawing or acceptance condition, not the entire software stack.

This is where local speed can become strategically important. A design revision can be turned into a visible sample in hours or a day, photographed or measured, and returned to the global team. The value is not simply low manufacturing cost. It is shorter feedback distance.

Evidence work

An evidence layer can also begin before final hardware. Test images, sample hashes, dimensions, expected outputs, known failure cases and acceptance criteria can be recorded early. When the real reference board later runs the function, the team can compare new evidence against the earlier contract instead of inventing a test after the fact.

The result is not four separate projects. It is one Work with several execution lanes tied to the same versioned boundary.

A practical parallel matrix

Lane

Can start before final hardware?

Early deliverable

Cannot claim yet

Application / GenAI

Yes

client logic, test vectors, contract handling

physical device behavior

Algorithm

Yes

result on fixed data, performance on declared dataset

live-camera performance

Hardware

Yes

interface Profile, power and connector design

final integrated product capability

Local making

Yes

sample, jig, trace, panel, fixture

repeatability beyond tested sample set

Simulator

Yes

deterministic operation semantics and replay

electrical or timing behavior of real board

Reference board

After binding

selected physical I/O and runtime evidence

delivered-unit identity or field conditions

Delivered unit

After freeze and shipment

instance-specific acceptance and local evidence

capabilities not included in that configuration

The matrix protects the central idea: **parallel work is useful only when the claim follows the evidence level.**

10.2 Remote reference validation and configuration freeze

A simulator reduces waiting. A reference board reduces uncertainty. They are not interchangeable.

The reference board stage exists to answer questions that cannot be settled from documents and host-side simulation alone.

Typical questions include: - Does the selected operating-system image boot reliably on the actual configuration? - Is the intended USB camera enumerated and stable? - Does the serial path use the expected logic level and transceiver? - Is the selected network route available under the intended security policy? - Does the application survive restart and reconnect? - Are operation identifiers preserved across process restarts? - Does the device report a busy state when an exclusive Resource is already in use? - Does cancellation produce a request, an acknowledgement and a separately verified physical stop where applicable? - Does the evidence contain the actual board identity, software version, driver version and timestamp?

The answer should be recorded against a named physical reference instance. A successful test on an x86 Azure host is not a T736 result. A successful test on one T736 reference board is not automatically a result for every board variant. A successful physical-reference run is not automatically a result for the device later delivered to another site.

Replaceable execution targets

A useful contract can target several execution classes: - simulation - physical_reference - physical_delivered

The same high-level operation can exist across those classes, but the capability Profile and evidence differ.

For example:

locate_marks(image_ref, marker_profile, calibration_ref?)

could be executed:

on Azure against a fixed image;

on a T736 reference board using a USB camera;

on a delivered T736 instance using its locally bound camera and calibration.

The application can preserve the operation meaning while the system preserves the executor identity and evidence level.

The freeze is a configuration, not a file name

A project often says “version frozen” when a software archive has been named. That is insufficient for physical delivery.

A delivery freeze should include, as applicable: - hardware module and carrier revision; - BOM options that affect interfaces; - power-input assumptions and external power budget; - connector allocation and pin usage; - image / operating-system identifier; - kernel, driver or device-tree changes needed by the selected peripherals; - Runtime version; - function-contract versions; - application version; - simulator/reference compatibility record; - recovery method; - credentials or provisioning method, without placing secrets in public artifacts; - offline dependencies required at the destination; - acceptance test and known limitations.

The freeze should also say what is **not** included.

This is particularly important for T736 because the documentary baseline includes optional and multiplexed functions. A connector shape or block diagram does not prove that every interface can be used simultaneously. The frozen configuration must identify the actual choices.

Freeze evidence should be readable by the next team

A freeze is successful only if another engineer can understand what was delivered.

The package should therefore contain a short START_HERE, an English collaboration Profile, source references, hashes, a supported-feature matrix, test procedure and result evidence. Large supplier source trees can remain separately indexed and retrieved only when needed.

This supports the HSIMC rule: **move the small, decision-relevant package first; retrieve large engineering sources by exact reference when the work actually depends on them.**

10.3 Shipping, local binding, recovery and continued use

Shipping hardware is not the end of the engineering path. It changes the location of the evidence problem.

Before shipment, the project can control a reference environment. After shipment, the delivered unit encounters a new network, new power conditions, new peripherals, a new operator and possibly a different maintenance model.

The delivered unit therefore needs its own binding record.

Bind the delivered instance

At minimum, local binding should record: - delivered device Resource ID; - hardware/BOM identity; - software and Runtime versions; - owner or responsible organization; - location and local operating scope; - connected peripherals and their Profiles; - allowed remote operations; - local stop/recovery responsibility; - evidence destination; - date and result of local acceptance.

A Card ID, QR or order number may link to this record. It should not replace the device identity.

Local operation should not depend on an always-open cloud session

A physical system should define what happens when the cloud is unavailable.

Possible local behaviors include: - finish a safe atomic operation already accepted; - reject new work while disconnected; - remain in a safe idle state; - allow a local operator to stop or recover; - queue evidence for later upload; - continue selected offline functions that do not require fresh cloud Authority.

The correct behavior depends on the application. A camera inspection station and a moving wheelchair do not have the same risk profile.

The July 28, 2026 MCP stateless core is helpful here because protocol continuity no longer depends on a hidden transport session. Each request can identify its context, while the application carries explicit Work, operation or Resource handles. The physical application still needs its own state machine and safe offline behavior.

Recovery is part of the product

A delivered system needs a known path back from failure.

For a T736-based node, a minimum carrier design should preserve service and recovery paths intentionally. The current documentary roadmap proposes keeping power, OTG/FEL/recovery, UART0 service, one network route, USB Host and selected serial expansion. The exact connector and pin plan remain to be frozen for the actual carrier.

Recovery evidence should answer: - Can the unit be identified without the normal application running? - Can a local operator return it to a known software state? - Can a failed update be distinguished from a hardware fault? - Can the evidence from the failed operation be retained? - Can the next engineer reproduce the recovery method from the package?

A product that works only while the original engineer is present has not yet become a reusable capability.

Chapter 11 - The T736 Reference Path

11.1 The existing module and a minimal common interface

HSIMC is not starting the T736 path from a blank sheet. The H0-001 handoff inherits C0-005 hardware review, C0-006 remote engineering planning, the C0-007 simulator and C0-008 recorded remote simulation result.

The current shortest hardware path uses the **HelperT736C** core module as the reference. The documentary baseline records a module size of approximately **60 x 40 x 3.0 mm**, with thickness tolerance stated as ±0.2 mm, using **188 castellated peripheral pads plus 34 bottom LGA contacts**. The reviewed reference carrier is approximately **150 x 105 mm**. These are documentary dimensions, not a newly measured unit in this Part.

The credit-card plane - approximately 85.6 x 53.9 mm as an HSIMC design target - is therefore a packaging challenge, not a proven current product form.

The engineering question is not “can the reference carrier be shrunk by scaling the drawing?” The question is: **which functions must remain on the common carrier, which can move into replaceable adapter modules, and which belong in the cloud or external equipment?**

A proposed minimum common carrier

The current handoff proposes a small common carrier that preserves, at minimum:

defined power input and protection;

USB0 / OTG / FEL or equivalent recovery path;

reset / boot service functions required for recovery;

UART0 service access;

one deliberate network route;

USB Host with a declared external power budget;

intentionally assigned serial expansion;

the mechanical, thermal and connector conditions needed for the selected product form.

This is a design scope, not a frozen schematic.

It deliberately avoids trying to expose every possible T736 function on a credit-card-sized board. Optional high-speed, display, camera and expansion functions can be added through product-specific carriers or adapters when a real Work requires them.

Power must be specified by configuration

The documentary baseline contains power statements that need reconciliation. The reference carrier accepts 12 V and converts to the module’s 5 V input; the reviewed documents contain conflicting interpretations of a 1.5 A statement, and some PK-bank voltage descriptions also require BOM/configuration confirmation.

Therefore a T736 Profile should not contain a single vague line such as “5 V / 1.5 A supported.” It should separate: - accepted carrier input range; - module supply requirement; - regulator rating and thermal condition; - peak current assumptions; - USB/peripheral power budget; - start-up or transient condition; - actual measured result on the named configuration.

This distinction matters when a camera, UWB radio, storage device or serial adapter is attached. The base module can boot successfully while the complete system still has an inadequate peripheral power budget.

Interface names are not electrical guarantees

A UART signal on the SoC side is logic-level serial. RS-232 and RS-485 are different electrical interfaces and require appropriate transceivers. A connector that resembles miniPCIe does not prove PCIe signaling. The reviewed T736 carrier includes cases where connector form and routed protocol differ. Likewise, multiplexed PHY resources mean that not every block-diagram function should be assumed to operate simultaneously.

The Profile must therefore describe:

logical interface -> electrical layer -> connector -> driver -> intended device -> evidence

That chain is much more useful to a remote engineer than a connector photograph alone.

11.2 Adapters for cameras, tags, positioning and controllers

A common T736 node becomes more useful when device-specific complexity can move into small, explicit adapters.

The adapter principle is simple:

Keep the function contract stable; replace the physical adapter when the local equipment changes.

Camera adapters

For the first vision path, a compatible USB camera is a practical candidate because it reduces the number of unresolved CSI-specific issues during early development. This is a route choice, not a claim that USB is always superior.

A camera adapter Profile should record: - vendor/model and USB identity; - resolution and supported pixel formats; - frame rate used by the application; - lens and focus condition; - mount position; - exposure / lighting assumptions; - capture timestamp behavior; - calibration version when metric coordinates are reported; - driver and kernel dependencies; - observed CPU/memory/latency on the actual T736 configuration.

Switching from a synthetic image to a real camera therefore creates a new evidence gate. The algorithm can remain the same version, but the camera, capture path and calibration become part of the result.

NFC, UHF RFID and UWB adapters

These technologies should not be collapsed into a single “wireless ID” function. - **NFC** can provide a short-range phone or device interaction path. - **UHF RFID** can support item identification at greater range with compatible readers and tags. - **UWB** can provide ranging, angle or positioning observations depending on the specific hardware and geometry.

A T736 node may connect to any of them through USB, serial, network or another supported interface. The actual module, driver, antenna and firmware determine capability.

The adapter should return an observation with its own time, quality and source rather than silently converting an old reading into current truth.

Machine and motion-controller adapters

The T736 should not be required to replace deterministic motor or safety control that is better handled by an existing PLC, servo controller or dedicated motion controller.

A useful architecture is:

GenAI / client -> HSIMC Work -> T736 adapter / local supervisor -> deterministic controller -> physical machine

The T736 layer can translate a versioned operation, check local conditions, collect status and package evidence. The motion controller remains responsible for the real-time behavior for which it was designed.

For old equipment, an adapter may use RS-232, RS-485, Ethernet, discrete I/O or a vendor protocol. Each electrical conversion and timing assumption should be part of the Profile.

The adapter is also a commercial boundary

A local integrator can contribute value without redesigning the entire platform. If a factory has a legacy cutter, a hotel has an existing cart, or a workshop has a laser with a proprietary controller, the local contributor can build and validate the adapter that connects that Resource to a shared Work contract.

This turns system integration into a modular capability that can be attributed, versioned, tested and paid for.

11.3 Documented capability, simulation and physical evidence

The T736 path is useful only if the evidence ladder remains visible.

A simple ladder is:

Evidence level

What it proves

Example

Documented

supplier/reviewed documentation states a capability or constraint

pin group, module dimension, block diagram

Simulated

application contract behaves under declared simulated conditions

C0-007 operation replay

Remote simulated

another client reaches the simulator through a remote path

C0-008 14/14 Azure record

Physical reference

a named reference board performs the selected function

camera enumerates; UART operation succeeds

Physical delivered

the delivered unit passes local acceptance

customer site run on delivered identity

Operational

the capability performs repeatedly under real Work conditions

field evidence over time

The ladder prevents a common failure in AI-assisted engineering: a plausible explanation becomes a design assumption, the design assumption becomes a presentation, and the presentation is later remembered as a tested fact.

C0-007 remains valuable precisely because it is limited

The C0-007 source-verified simulator exposes hsimc_get_profile, hsimc_sim_operation and hsimc_get_operation, plus a simulated Profile resource and SQLite operation records. It supports duplicate handling and exclusive-group conflict behavior for its declared test conditions.

It also reports physical targets as unbound.

That is the correct behavior for a simulator. It allows application engineers to work on operation semantics without inventing a physical result.

The current C0-007 initialization declares protocol version 2025-11-25. Therefore Part III's discussion of MCP 2026-07-28 is a target architecture, not evidence that the simulator has already migrated. A future migration should be implemented and tested, not achieved by editing a version string.

C0-008 is a remote-path record, not hardware acceptance

C0-008 records that an engineering client called the Azure simulator path and passed 14/14 checks on 10 November 2026. That is useful evidence that the remote simulation route worked under the recorded conditions.

It does not prove: - current Azure availability; - current T736 reachability; - hardware I/O; - camera compatibility; - current client compatibility with the 2026-07-28 protocol; - delivered-product behavior.

A white paper gains credibility by saying this plainly.

The next physical-reference gate

When a real T736 reference instance is reachable and authorized, the first physical-reference gate should remain intentionally narrow.

A reasonable sequence is:

identify the board and configuration;

record hardware, OS and Runtime versions;

verify recovery/service path;

verify the selected network route;

bind one permitted peripheral;

run one read-only or low-risk observation function;

persist operation and observation evidence;

test restart / reconnect / duplicate handling;

only then expand to a bounded physical action.

The project does not need the entire supplier BSP to be mirrored into every engineer's workspace before this can begin. The exact source dependency can be retrieved when a driver, device tree or kernel change actually requires it.

Chapter 12 - One Function, Many Makers

12.1 The proposed locate_marks contract

A platform can become abstract very quickly. A small function makes the collaboration model concrete.

HSIMC therefore proposes locate_marks as a reference engineering function for the next phase.

The function answers a narrow question:

Given a versioned image and marker definition, where are the expected marks, and what evidence supports the result?

This is useful in printing, card making, laser alignment, fixture verification, packaging, labels and other local production tasks.

It is deliberately not a machine-motion command.

Why this function is a good first reference

It connects several existing capabilities without requiring the final product to be complete: - a designer can create the mark pattern; - an ordinary printer can produce a sample; - a local maker can build a jig; - a phone or camera can capture the sample; - a remote vision engineer can develop the detector; - Azure can run fixed-image tests; - T736 can later run the same function with a real camera; - an operator can compare the returned evidence against the sample; - a laser or other machine can consume the result only through a separate, authorized action contract.

This separation makes it possible to develop vision and physical process work in parallel without allowing a detection result to become an unintended actuator command.

Candidate input contract

A first candidate request might contain:

Field

Purpose

work_id

links the request to the engineering or business Work

operation_id

makes this specific invocation explicit and queryable

target_resource_id

identifies simulator, reference or delivered executor

image_ref

references the versioned input image or capture

marker_profile

defines which marks should be located

calibration_ref

optional calibration when metric coordinates are required

algorithm_version

identifies the detector implementation

authority_context

identifies the current permitted requester / scope

The exact wire schema remains to be implemented and tested. The important point is that the operation does not rely on hidden conversation memory to know which image, profile or executor is intended.

Candidate output contract

A useful result can include: - operation_id; - observation_id; - executor identity and execution class; - capture or source timestamp; - algorithm version; - detected marker IDs; - pixel coordinates; - confidence or quality values; - calibration reference if metric coordinates are emitted; - input-image hash; - evidence references; - warnings or unsupported conditions.

If no valid calibration is available, the function should report pixel coordinates rather than quietly inventing millimetres.

If the image is stale, damaged or unsupported, the function should say so.

Observation, not control

The result means: **this executor, using this algorithm and this input, observed these marks at these locations.**

It does not mean: - the material is safe to process; - the machine is aligned; - the operator has authorized motion; - the laser should fire; - the result is within the customer's tolerance; - a later image will contain the same marks.

Those are separate decisions and evidence steps.

This is exactly the kind of boundary Part III called a shared contract.

12.2 From synthetic images to a real sample

The first useful implementation can begin before a live camera is bound.

Stage A - fixed synthetic and generated inputs

Create a small versioned dataset: - nominal sample; - rotated sample; - scale variation; - one missing marker; - partial occlusion; - low contrast; - unsupported format; - wrong marker profile.

Each input has a known expected result. The purpose is to validate the function semantics, not to claim field accuracy.

A remote engineer can implement the detector and produce deterministic evidence against that set.

Stage B - locally printed sample and jig

A local print shop or workshop receives the versioned artwork and fixture dimensions. It prints a real sample, places it in the jig, and captures images under declared lighting and camera conditions.

The local deliverable should include: - sample version; - print method / material description; - jig version; - sample photograph; - measured dimensions relevant to acceptance; - image files and hashes; - visible defects or deviations; - operator / location / time evidence where appropriate.

The remote detector then runs on the real image.

Now the team can distinguish two failure classes:

the physical sample or capture does not match the expected condition;

the algorithm fails on an acceptable physical sample.

Without this separation, software and manufacturing teams can easily blame each other without evidence.

Stage C - T736 reference camera path

Once the reference board is reachable, bind a specific camera Profile.

The same function is executed from a live or locally captured frame. The evidence now adds: - T736 Resource identity; - hardware and Runtime version; - camera identity and driver; - capture settings; - timestamps; - processing latency; - CPU / memory observations if relevant; - calibration identity; - result evidence.

The algorithm may be unchanged from Stage A. The evidence level is different because the capture and runtime are physical.

Stage D - bounded downstream use

Only after the observation is trusted should another Work consume it.

For example, a laser-alignment preparation function could transform detected marks into a proposed coordinate transform. A human or local controller could validate the proposal. A separate authorized machine action could then execute under its own operation ID.

This creates an auditable chain:

sample -> image -> locate_marks observation -> alignment proposal -> acceptance -> machine action -> output inspection

Each step can be owned by a different contributor while remaining part of one Work.

12.3 Independent reproduction at a second location

The strongest evidence that a capability is reusable is not that the original team can repeat it. It is that another team can reproduce the function from the package.

A second-location reproduction should therefore be an explicit engineering gate.

What the second team receives

The team should receive a compact package containing: - START_HERE; - function contract; - supported / unsupported matrix; - sample artwork; - jig drawing or construction note; - fixed test images; - algorithm or client package within the permitted distribution scope; - environment / dependency instructions; - expected results; - evidence template; - source references and hashes.

Large supplier archives remain referenced separately unless the reproduction actually needs them.

What the second team must return

The return package should include: - which instructions were sufficient; - which instructions were ambiguous; - environment and tool versions; - the result for each test vector; - real sample or image evidence if produced; - deviations from the expected result; - time spent setting up and running the work; - any local changes required; - a clear statement of what was and was not reproduced.

This produces more than a pass/fail result. It measures the quality of the collaboration interface.

Reproduction time is a product metric

One useful metric for a GenAI-native physical engineering system is:

time from receiving a versioned package to independently reproducing the agreed function.

A second metric is the number of clarifications required.

A third is the amount of hidden local knowledge that had to be rediscovered.

These metrics expose a different kind of engineering debt. A design may work perfectly in the original workshop but still be difficult to transfer. That transfer difficulty is part of the product.

Many makers, one evidence chain

The locate_marks example shows how a simple function can connect many participants:

Contributor

Possible contribution

Evidence

Designer

marker artwork / version

design hash

Printer

physical sample

print / material record

Local maker

fixture / jig

dimensions, photos, version

Vision engineer

detector

algorithm version and test results

Cloud engineer

simulator / service

execution records

T736 engineer

camera adapter

board, driver and capture Profile

Operator

local confirmation

observation / acceptance record

Machine integrator

downstream adapter

action contract and result evidence

Second site

independent reproduction

reproduction report

None of these participants needs to own the entire product.

That is the engineering meaning of the wider HSIMC invitation:

Bring your existing capability. Deliver one reproducible function.

The system becomes more powerful as the functions become easier to combine, but combination should never erase the evidence boundary of each contribution.

From one function to a capability network

After one function is reproducible, the same method can expand: - read_tag - write_tag - measure_range - get_machine_state - prepare_job - reserve_slot - verify_output - request_pickup - confirm_handover

Each function can be implemented by different local Resources and different technologies.

The common requirement is not identical hardware. It is a contract that another participant can understand, an Authority model that limits action appropriately, and Evidence that shows what happened.

Part IV conclusion - The product is the reproducible path

The engineering opportunity in the GenAI physical era is not simply to make one more smart device.

It is to shorten the distance between a plan and a verified outcome while allowing more people and more existing facilities to participate.

That requires a path where: - software can begin before custom hardware is finished; - hardware can stabilize interfaces before every application detail is known; - local makers can turn drawings into real samples quickly; - a simulator can prove semantics without pretending to be hardware; - a reference board can prove selected physical behavior; - a frozen configuration can be shipped with enough evidence for another engineer to understand it; - the delivered unit can be bound locally and continue working without the original project team standing beside it; - a second team can reproduce the function and expose hidden assumptions.

T736 is the current shortest hardware reference for this route, not the destination of the entire ecosystem. The credit-card-sized form is a design challenge, not a prerequisite for useful collaboration. The deeper product is the **reproducible engineering path** connecting people, tools, resources and evidence.

This is the bridge from Part III's collaboration model to Part V's question of value: **when the path becomes reproducible, what economic, social and environmental outcomes does it create, and how should those outcomes be measured without double counting?**

Evidence and source notes for Part IV

HSIMC engineering evidence - H0-001 Card Phone and Cloud Development handoff, 11 November 2026. - C0-005 T736 documentary hardware standardization baseline, 9 November 2026. - C0-006 remote engineering plan, 9 November 2026. - C0-007 source-verified simulator and recorded simulator tests, 10 November 2026. - C0-008 recorded Azure remote simulation result, 10 November 2026: 14/14 checks in the recorded report.

External context - OpenAI, *Accelerating engineering for the physical world*, 3 November 2026. The article describes work across mechanical engineering, electrical engineering, software, simulation and manufacturing, and invites engineering and manufacturing collaborators. HSIMC's architecture and claims remain independent. - Model Context Protocol, *The 2026-07-28 Specification*, 28 July 2026. The release introduces a stateless protocol core, self-describing requests, header-based routing, an extensions framework and related authorization changes. Application state may still be carried through explicit handles.

Evidence discipline

This Part distinguishes proposed architecture, documentary hardware information, simulated evidence, recorded remote-simulation evidence and future physical validation. No statement in this Part should be read as claiming a currently reachable T736, a completed credit-card-sized product, a live-camera acceptance result, or a delivered-unit test unless separately supported by dated physical evidence.

Part V - The Value: Useful Outcomes and Measurable Participation

Part V at a glance

Parts I-IV moved from opportunity, to existing capabilities, to accountable collaboration, and then to an engineering path that allows multiple teams to develop together before final hardware is shipped. Part V asks the next question: **what useful value is actually created, for whom, and how can that value be measured without counting the same benefit several times?**

This is where a white paper can easily become vague. It is simple to say that a project supports sustainability, saves labour, improves inclusion, creates innovation, reduces waste and contributes to many Sustainable Development Goals. It is harder to show which result happened, which participant benefited, which participant carried a new burden, which cash flow actually changed, which material or energy flow changed, and what evidence supports the claim.

HSIMC therefore treats value as a set of observed changes, not as a marketing adjective.

A useful Work should make it possible to answer at least five questions:

Outcome - What changed in the physical or service world?

Participant value - Who gained time, income, capability, access, convenience or reduced risk?

Participant cost - Who contributed cash, labour, equipment time, materials, attention or responsibility?

Evidence - What baseline, receipt, measurement or acceptance record supports the result?

Distribution - Was the benefit shared fairly enough that the system can continue without shifting hidden cost to another participant?

This Part develops that approach in three chapters.

Chapter 13 follows three classes of real-world journeys: **make, move and assist**. These are not claims that every proposed HSIMC application has already been deployed. They are reference patterns for organizing Work, custody, Authority and Evidence. Where adjacent services already exist, such as airport baggage delivery, smart baggage tags and trolley monitoring, the text identifies them as external evidence rather than as HSIMC deployments.

Chapter 14 defines a practical value-accounting method. It separates cash results, released capability, user/service outcomes, environmental effects and governance evidence. It also distinguishes participant-level accounting from network-level accounting so that a payment from one participant to another is not counted twice as new value. For greenhouse-gas claims, it separates organization inventories from project/intervention effects and treats avoided-emission claims as a separate analytical layer rather than subtracting them casually from a footprint.

Chapter 15 maps the work against all 17 United Nations Sustainable Development Goals. The purpose is not to claim that one platform “achieves” 17 Goals. The purpose is to state where a contribution may be relevant, what evidence would be required, and where the relevance remains conditional. The strongest initial connections are expected to be around skills, decent work, inclusive industrial participation, reduced inequalities, responsible production and multi-stakeholder partnerships. Other Goals become relevant only when a concrete application creates the corresponding outcome.

The value principle of this Part is:

Count the change once. Show who received it. Show who paid for it. Keep the evidence close to the Work.

Chapter 13 - Real-World Journeys: Make, Move and Assist

13.1 Make: from available material to accepted output

The most basic real-world journey is transformation.

A material, part, file or object begins in one state. A person, machine or process changes it. Another party accepts the result. The economic value is not created simply because a machine was switched on. It is created when a requested result reaches a condition that someone is prepared to accept and use.

A useful manufacturing Work therefore has more structure than a purchase order line.

A minimal chain is:

Demand -> confirmed input -> preparation -> process -> inspection -> accepted output -> storage or next handover

Each transition can be represented by a different responsible party and a different type of evidence.

For a small printed or embroidered interface, for example: - the demand defines the pattern, quantity, substrate, tolerance and delivery window; - the material record identifies the substrate, thread, ink, film, laminate or other input actually used; - preparation records the file version, machine settings, fixture and operator; - the process record identifies the machine or maker and the time window; - inspection returns dimensions, resistance, image evidence, read/write result or another relevant acceptance measure; - the accepted-output record identifies who accepted the sample and under which version; - storage or delivery records where the accepted items went next.

That structure makes it possible for a very small producer to participate without pretending to be a fully integrated electronics company.

A sewing or embroidery workshop may provide the controlled placement step. A print shop may produce registration marks or a functional surface. A hardware team may provide the NFC or sensor module. A remote engineer may provide image inspection or test vectors. A local integrator may make the connector or fixture. The customer may accept the integrated result.

The commercial unit is therefore not “one factory”. It can be **one accepted transformation** assembled from multiple capabilities.

Demand, inventory and process should remain distinct

A common failure in digital manufacturing systems is to treat three different statements as if they were equivalent: - “the customer wants ten units”; - “the material is available”; - “the machine can produce ten units before the deadline”.

They are not equivalent.

Demand describes a requested future result. Inventory describes a material or object state at a particular observation time. Capability describes what a Resource may be able to do under specified conditions. A commitment is created only when an authorized provider accepts a Work against a sufficiently current view of those conditions.

This distinction becomes especially important in a stateless collaboration model. The client should be able to re-submit the context needed to understand a request, but the material reservation, machine booking, accepted Work and operation receipt remain persistent business facts.

Low frequency can still create a valuable shared network

Part I introduced a recurring observation from small industrial spaces: a single factory may need a movement or production capability only occasionally. That does not mean the capability has no market.

The more useful question is whether many low-frequency users can share a capability without unacceptable conflicts.

For a shared machine, fixture, power module or inspection station, the relevant variables include: - setup and teardown time; - travel or transfer time between users; - minimum useful batch size; - preparation requirements; - operator availability; - maintenance window; - storage requirement; - deadline overlap; - failure and rework probability.

The business opportunity is not merely to sell more hardware. It may be to make existing hardware, skill and space easier to schedule and trust.

A local maker who can return a verified sample in two hours may create more value for a fast iteration than a remote supplier with a lower unit price but a five-day feedback loop. The correct comparison is the **full cost of reaching an accepted result**, including waiting, failed iterations, communication and rework.

A transformation evidence packet

A reproducible manufacturing Work should eventually be able to produce a compact evidence packet containing, as appropriate: - Work ID and accepted version; - responsible requester, provider and acceptor; - input identity and observed quantity; - material batch or source reference; - machine / tool / process Profile; - process parameters that materially affect the result; - start and completion timestamps; - inspection result and uncertainty where relevant; - rejected or reworked quantity; - accepted quantity; - actual direct cost and time; - next custody or storage location.

Not every local maker needs to operate this data model directly. A phone, card, QR/NFC entry or HSIMC assistant can translate a simple local interaction into a structured record.

That is the practical meaning of making local capability globally usable: **the person does not need to become a software platform; the contribution needs a reliable interface and evidence.**

13.2 Move: from warehouse to machine, airport to hotel, and one custodian to another

Movement creates value when an object reaches the right place, at the right time, under the right custody and condition.

A route alone is not enough.

For a warehouse-to-machine Work, the relevant state change may be:

available stock -> reserved stock -> picked -> in transit -> received at machine -> consumed or returned

For finished goods it may be:

accepted output -> stored -> released -> handed to carrier -> received by next custodian

For travel baggage it may be:

passenger has bag -> bag accepted for service -> bag in custody -> bag arrives at hotel or storage point -> guest chooses receiving time -> final handover

Each step has a different Authority and different evidence.

Existing travel services show that the pieces already exist

Hong Kong International Airport currently lists paid baggage delivery services that can move baggage between the airport and a passenger's hotel, home or designated downtown location, depending on the provider. This proves that airport-to-hotel baggage movement is already a purchasable service category; it does not prove the HSIMC proposed timing and handover model has been deployed.

HKIA also operates MyTAG, an RFID smart baggage tag that can notify a passenger through the My HKG app when a bag is ready for pickup. The MyTAG service uses tag registration plus flight information, and its support documentation describes scanning the tag QR code and linking the tag to a flight. HKIA also describes a real-time trolley monitoring and tracking system for keeping baggage trolleys available at appropriate locations.

These examples matter because they show that identity, tracking, booking, baggage movement and trolley availability are already separate operational capabilities. A new service does not need to reinvent every layer. It can focus on the missing handover between existing layers.

The proposed airport-hotel journey

The HSIMC concept separates three states that are often collapsed: - **Room ready** - the hotel room is available; - **Bag ready** - the baggage has reached an agreed receiving or storage point; - **Guest ready** - the guest wants the final handover now or within an accepted time window.

This distinction is operationally important.

A guest may arrive before the baggage. A bag may arrive while the guest is showering, sleeping, in a meeting or simply not ready to answer the door. A hotel may have luggage storage but no bell service at a particular hour. An economy hotel may have minimal staff. A person with heavy work equipment may value assistance even when the room rate is low.

The proposed service therefore does not assume that "bag at hotel" means "deliver immediately to room".

A more respectful sequence is:

the transport provider hands the bag to an authorized hotel or storage receiver;

the bag becomes **Bag ready**;

the guest receives a quiet digital notice;

the guest chooses now, a time window, continued storage, or an agreed pickup point;

a provider accepts the final handover Work;

the bag is handed over and the custody record closes.

The value can then be measured as more than transport distance. It includes fewer unsuccessful room visits, less unwanted interruption, clearer staffing, clearer pricing and potentially better access to luggage assistance in hotels that do not provide traditional bell service.

Identity is not custody

Barcodes, airline baggage identifiers, UHF RFID, QR and NFC can help identify or associate an object. They do not by themselves prove who currently holds the bag or who is authorized to move it.

A service record should therefore separate: - observed tag / barcode identifier; - internal bag object identity for the service; - current custodian; - requested destination; - accepted next handover; - guest or owner Authority; - service status; - completed handover evidence.

The same principle applies to warehouse materials, returned equipment, exhibition crates and shared carts.

The proposed airport trolley journey

Another proposed journey is the empty trolley after a passenger reaches a boundary where the trolley cannot or should not continue.

The first useful outcome is not “perfectly organize every trolley”. It may simply be:

Move the released trolley away from a high-friction handover point quickly enough that the space becomes available to the next traveler.

That outcome can be measured directly: - time from traveler release to cleared space; - number of trolleys blocking the handover point; - worker minutes required per cleared trolley; - secondary congestion created at the temporary holding point; - retrieval time from temporary holding to normal trolley pool.

Automation is only one possible execution mode. Human collection, powered assistance, a clip-on module or a future autonomous unit can all be compared against the same outcome.

This is an important pattern for HSIMC: **define the outcome before deciding which machine should own it.**

AI is already entering the travel planning layer

IATA’s 2026 Global Passenger Survey reports that 49% of respondents had used AI for travel planning, based on more than 10,000 responses from over 200 countries. The report explicitly notes the importance of connecting AI conversations to real booking and service flows.

That makes the physical handover layer increasingly important. An AI may help a traveler decide what they want. The value gap remains whether an authorized local provider can accept the requested Work and complete it.

13.3 Assist: make existing mobility and work easier without collapsing safety boundaries

“Assist” covers a wider class of outcomes than manufacturing or logistics.

The user may be a worker moving a heavy cart, an older person using a wheelchair, a hotel employee moving linen, a researcher moving equipment, or a person who simply needs a task to require less physical effort.

The opportunity is attractive because many existing assets already perform most of their function well.

A wheelchair already supports the person. A hotel bell cart already carries luggage. A trolley already rolls. A pallet already supports goods. A sewing machine already places thread. A print shop already controls registration. The proposed innovation may be an additional capability that can be attached, borrowed, scheduled or removed.

This leads to a general rule:

Upgrade the missing capability before replacing the whole asset.

Assistance must preserve the difference between person and cargo

A cargo cart and an occupied wheelchair may use related positioning, power, communication or evidence infrastructure. They do not share the same safety case.

A successful empty-cart movement cannot be used as evidence that an occupied wheelchair is safe. A UWB position estimate cannot, by itself, authorize movement. A software stop request cannot substitute for a verified braking and recovery design. A clip-on mechanical interface that works on a pallet does not prove compatibility with a wheelchair frame.

Any assistive-mobility branch therefore needs its own: - user and caregiver roles; - mechanical load and stability tests; - braking and hill-hold behavior; - speed and acceleration limits; - obstacle and drop-off handling; - manual takeover method; - battery and charging conditions; - fault response; - environment and route constraints; - local regulatory and operator requirements; - user acceptance and dignity considerations.

The shared HSIMC layer is the contract, Authority, Resource discovery and Evidence method. The physical safety evidence remains application-specific.

Human effort is not automatically a cost to eliminate

An inclusive system should avoid treating every minute of human work as waste.

Some work is skilled, socially valuable and a source of income. Some human presence is a safety control, a service relationship or a form of dignity. The question is not simply “how many people can the machine replace?”

The better questions are: - Which task is physically burdensome or repetitive? - Which task requires judgment, trust or human interaction? - Can a tool allow one person to do safer or higher-value work? - Does released time actually reduce paid overtime or enable more paid service? - Does the change remove an earning opportunity without creating another pathway? - Who decides whether the assistance is desirable?

This connects directly to the “No One Left Behind” principle. A physical GenAI system should not create value for one party by making another party’s work invisible.

Measure dignity and choice where they are part of the service

For baggage, assistance and hospitality, value may include the ability to choose when an interaction occurs.

For mobility, value may include reduced physical strain, greater independent movement or easier caregiver support.

For local makers, value may include access to work that previously required a larger company or a distant engineering team.

These are real outcomes, but they require appropriate measures. A project should not convert every social benefit into a fictional dollar amount merely to make the spreadsheet look complete.

Some value should remain in its natural unit: - minutes of unwanted interruption avoided; - kilograms or meters of manual handling avoided; - percentage of requested time windows met; - number of people newly able to perform a task; - number of makers who can independently reproduce a function; - number of participants receiving paid work; - user-reported ease or confidence under a defined survey method.

Part V therefore uses multiple ledgers rather than forcing every result into one financial number.

Chapter 14 - Measuring Value without Double Counting

14.1 Separate the ledgers before combining the story

A multi-party system can create confusion because the same event appears differently to each participant.

If a hotel pays HKD 100 to a baggage-service provider, the hotel sees HKD 100 of cost and the provider sees HKD 100 of revenue. If a joint report adds both values as “HKD 200 of economic value created”, it has counted the same transfer twice.

If an automation saves 20 worker-minutes but no worker leaves early, no overtime is reduced and no additional paid service is completed, the project has released capacity. It has not yet demonstrated a cash amount of cash savings.

If a recycled material replaces virgin material in one component, the project should not automatically claim the full avoided emissions of an entire product unless the system boundary, baseline and life-cycle effects support that conclusion.

HSIMC therefore uses several ledgers.

Ledger 1 - Cash

Record actual incremental cash flows for each participant.

For participant *j* during period *t*:

Incremental cash contribution(j,t) = new cash revenue + verified avoided cash expense - new operating cash expense

Initial integration or equipment investment is recorded separately.

A simple period ROI can then be calculated as:

ROI(j) = (sum incremental cash contribution - initial investment) / initial investment

when initial investment is greater than zero.

If the investment is zero or primarily in-kind, report net cash contribution and in-kind contribution rather than forcing an undefined ROI ratio.

Cash evidence should come from actual invoices, payroll/overtime changes, supplier invoices, paid service orders or another auditable record. A model estimate can be shown as a forecast but should not be relabeled as realized cash.

Ledger 2 - Released capability

Measure capacity in operational units.

Examples include: - worker minutes released; - machine minutes newly usable; - queue time reduced; - additional jobs completed within the same shift; - faster sample iteration; - shorter time from design freeze to visible sample; - shorter time from release package to second-site reproduction; - more storage positions usable; - fewer unsuccessful handover attempts.

Released capability can later cause a cash effect. Until that effect occurs, it remains an operational result.

Ledger 3 - User and service outcome

Measure the outcome in terms that matter to the user or operator.

Examples include: - percentage of deliveries completed in the requested time window; - minutes a trolley blocks a transfer point; - number of re-deliveries or failed handovers; - user-selected delivery method; - travel or walking distance avoided; - successful accessibility assistance requests; - defect or rework rate; - time to find the correct material or tool; - time to recover from a failed operation.

This ledger prevents the project from optimizing a technical metric that the user does not value.

Ledger 4 - Environmental and material effect

Record actual material and energy flows before making a sustainability claim.

Possible measures include: - mass of reused material; - mass and type of new material added; - failed sample and scrap mass; - packaging material; - electricity consumed in processing; - battery charging energy; - transport distance or service movements; - repaired or extended-life assets; - product or component lifetime; - end-of-life route where known.

The GHG Protocol Product Standard is useful as a reference because it frames product life-cycle emissions across raw materials, manufacturing, transportation, storage, use and disposal. That breadth is important: local reuse may reduce one material flow while adding processing energy, transport or an electronic component elsewhere.

Ledger 5 - Governance and evidence

Record whether the system actually operated under the promised Authority and accountability model.

Examples include: - percentage of physical actions linked to an accepted Work ID; - percentage of deliveries with a completed custody record; - operations rejected because Authority was missing or expired; - duplicate requests correctly resolved without repeating a physical action; - stop/recovery requests with confirmed outcomes; - evidence completeness by Work type; - payment disputes and resolution time; - number of participants able to inspect the evidence relevant to them.

Governance is not a decorative ESG category. It determines whether the other ledgers can be trusted.

14.2 Baseline, attribution and the difference between transfer and creation

Value measurement needs a baseline.

Without a baseline, “saved time” is merely a description of a new process. Without a baseline, “reduced waste” may reflect a smaller order rather than a better process. Without a baseline, “higher productivity” may be caused by a different product mix.

A practical Work should therefore record, where feasible: - baseline period or baseline method; - planned result; - provider commitment; - actual result; - variance; - known external changes; - confidence or uncertainty where important.

The baseline does not always need to be sophisticated. For a small workshop it may be three observed manual runs before introducing an assistance tool. For baggage delivery it may be the historical number of failed or unwanted handovers. For sample iteration it may be the time from design change to measurable sample under the previous process.

Participant accounting and network accounting are different

At participant level, a payment is economically real.

At consolidated network level, a payment between two participants is a transfer. It should be eliminated if the purpose is to estimate net value created by the network as a whole.

A useful consolidated equation is:

Network net cash contribution = external revenue + verified avoided external cash cost - new external operating cost

Internal payments between included participants are eliminated in consolidation.

This is similar in spirit to eliminating intercompany transactions in consolidated financial statements, although this white paper does not propose a new accounting standard.

A time saving is not automatically a wage saving

Suppose a worker previously spent 30 minutes pushing a cart and now spends 10 minutes supervising an assisted movement.

The project has released 20 minutes of capacity.

Possible next outcomes include: - the worker finishes the shift earlier and overtime falls; - the worker completes another paid delivery; - the worker performs quality control instead; - the worker waits because no other task is available; - the employer reduces staffing; - the service improves without changing staffing.

Those are materially different social and financial outcomes.

The system should therefore preserve the operational result first, then observe what happens to cash, workload, service quality and employment.

Avoid counting the same environmental benefit twice

GHG Protocol guidance distinguishes organization-level inventory accounting from project/intervention accounting.

Inventory accounting describes emissions associated with an organization’s defined boundary and value chain. Project accounting estimates the effect of a specific intervention relative to a counterfactual baseline. The two methods answer different questions.

This distinction is especially useful for HSIMC.

If a local project reuses an existing pallet, reduces a transport movement or extends the useful life of equipment, the project may estimate an intervention effect relative to an appropriate baseline. That result should not automatically be subtracted from every participating organization’s Scope 1, 2 or 3 inventory. Avoided-emission estimates should be shown separately and with their assumptions.

The GHG Protocol also emphasizes avoiding or transparently handling double counting in comparative and crediting contexts. The practical HSIMC rule is:

Report the physical change, the inventory effect and any estimated avoided effect as separate facts unless a recognized method explicitly connects them.

Attribution should follow the evidence, not the loudest brand

Multi-party Work creates another risk: one organization may claim the whole result even though several parties made it possible.

A more credible record separates: - the party that defined the demand; - the party that supplied the material; - the maker or operator; - the model or software contributor; - the equipment owner; - the site provider; - the logistics provider; - the funder or sponsor; - the acceptor; - the party responsible for measurement.

A published case can then describe contribution without pretending that every participant caused the entire outcome alone.

14.3 Decide what to scale by comparing marginal value with marginal burden

The purpose of measurement is not reporting for its own sake. It is to make better next decisions.

A project can therefore use staged investment gates.

Gate 1 - Is the problem observable?

Can the team measure the current delay, cost, burden, defect or access problem?

If not, the first Work may simply be a baseline observation.

Gate 2 - Can a minimal service improve the result without new custom hardware?

Use existing people, QR, phones, available tags, manual confirmation, borrowed tools or existing logistics where appropriate.

This identifies whether the service concept has value before expensive automation.

Gate 3 - Does an added tool improve the full outcome?

Introduce a fixture, reader, UWB module, clip-on power unit, T736 adapter or software automation.

Measure: - new direct cost; - failure and intervention rate; - time saved; - additional demand served; - user outcome; - safety or governance burden; - maintenance requirement.

Gate 4 - Does sharing improve utilization without creating coordination cost that cancels the benefit?

A shared asset has a second-order cost: reservation, travel, setup, return, charging, cleaning, maintenance and conflict management.

The relevant utilization metric is not simply “percentage of day powered on”. It may be:

useful accepted Work / total capability cost

A low-utilization asset can still be rational if the alternative is extremely costly, urgent or inaccessible. Conversely, a frequently used asset can be a poor investment if failures and supervision are high.

Gate 5 - Can a second site reproduce the value?

Replication should reproduce both the function and the accounting logic.

A second site should be able to state: - what stayed the same; - what changed locally; - what cost changed; - what acceptance measure changed; - whether the original benefit remained; - which new burden appeared.

This is the difference between a demonstration and a scalable method.

A practical value dashboard

A compact Work-level dashboard can therefore show:

Dimension

Example measure

Evidence source

Outcome

accepted units / completed handovers

acceptance receipt

Cash

external revenue, avoided external expense, operating cost

invoice / payroll / supplier record

Capacity

minutes released, jobs added, queue reduced

timestamped workflow

User

time-window success, interruption, accessibility

service record / survey

Material

reused mass, scrap, added material

batch / weight record

Energy

kWh or battery cycles where material

meter / charger / estimate with method

Governance

authorized actions, duplicate handling, recovery

operation ledger

Replication

time to reproduce at second site

release and acceptance records

The dashboard should not be used to compress all dimensions into a single score unless the weighting method and purpose are explicit.

A business can make money while creating a social burden. A project can produce an environmental benefit while losing money. A service can improve user dignity without reducing labour cost. A training programme can create capability before it creates revenue.

The value model should make these differences visible rather than hide them.

Chapter 15 - Seventeen Goals, Specific Contributions

15.1 Use the SDGs as a relevance and evidence map, not as a badge collection

The United Nations 2030 Agenda describes 17 Sustainable Development Goals and 169 targets as integrated and indivisible. It also commits to leave no one behind and to reach those furthest behind first. Goal 17 explicitly calls for multi-stakeholder partnerships that mobilize and share knowledge, expertise, technology and resources, as well as better timely and reliable data.

Those ideas fit naturally with a network of local capability and global collaboration.

But fit is not proof.

A white paper should not display seventeen goal icons and imply that every Work advances every Goal. Relevance depends on the application, baseline and actual result.

HSIMC therefore proposes three levels of SDG connection: - **Direct initial relevance** - the collaboration model itself can reasonably generate evidence related to this Goal. - **Application-dependent relevance** - the Goal becomes relevant only when a particular Work creates the corresponding outcome. - **Conditional / guardrail relevance** - the Goal mainly identifies an effect that must be checked or a future application area, not a current positive claim.

The following matrix is a starting point for measurement, not a certification.

SDG	Relevance to HSIMC approach	Example contribution path	Evidence needed before a positive claim
1 No Poverty	Conditional	paid micro-work, access to productive capability, resilient local income	net income change, participant baseline, persistence, displacement effects
2 Zero Hunger	Application-dependent	farm logistics, local food processing, reduced handling loss	food quantity/quality, loss baseline, farmer/work income, energy/material effects
3 Good Health and Well-being	Application-dependent	assistive mobility, reduced physical strain, safer handling	user outcome, safety evidence, adverse events, accessibility data
4 Quality Education	Direct initial relevance	learning by producing a reproducible function; technical/vocational skill formation	pre/post capability, independent task completion, participation and access
5 Gender Equality	Cross-cutting guardrail	reduce barriers to participation; recognize unpaid/care constraints; equal decision roles	

participation, pay, leadership, workload, access data disaggregated appropriately

6 Clean Water and Sanitation

Conditional

water-process monitoring or lower-impact local processing

water use, discharge, contamination and treatment data

7 Affordable and Clean Energy

Application-dependent

energy-aware local production, shared equipment, lower-energy process

measured energy, baseline service equivalent, equipment life and rebound effects

8 Decent Work and Economic Growth

Direct initial relevance

paid local capability, safer work, SME participation, new productive roles

pay, hours, safety, job quality, external revenue, worker outcome

9 Industry, Innovation and Infrastructure

Direct initial relevance

connect small makers and existing equipment to digital/AI engineering

number of capable SMEs, accepted functions, infrastructure use, replication evidence

10 Reduced Inequalities

Direct initial relevance

lower entry barriers by language, hardware, geography, disability or firm size

who participates, who is excluded, affordability, accessibility, benefit distribution

11 Sustainable Cities and Communities

Application-dependent

accessible mobility, local repair/making, better use of shared urban assets

service access, congestion, safety, material and operating outcomes

12 Responsible Consumption and Production

Direct initial relevance

reuse, repair, traceable materials, less unnecessary replacement, verified production

material flow, scrap, added components, life extension, end-of-life and quality evidence

13 Climate Action

Application-dependent

measured energy/material interventions and climate-related capability building

GHG method, baseline, boundary, uncertainty, separate inventory/project reporting

14 Life Below Water

Conditional / guardrail

prevent leakage or harmful waste from materials/processes

discharge, waste route, chemical/material evidence relevant to aquatic impact

15 Life on Land

Conditional / guardrail

responsible material sourcing, reuse, land/agriculture applications

source traceability, land/ecosystem baseline, waste or restoration outcome

16 Peace, Justice and Strong Institutions

Direct governance relevance

transparent Authority, evidence, grievance and accountable decision records

access rules, consent, dispute handling, auditability, participation evidence

17 Partnerships for the Goals

Direct initial relevance

multi-stakeholder sharing of knowledge, expertise, technology and Resources

actual commitments, Work completed, contribution records, reliable data and replication

The strongest initial cluster

For the present white paper, the most defensible initial cluster is:

SDG 4, 8, 9, 10, 12 and 17.

Why?

Because the proposed system directly concerns: - skills that let people perform useful technical work; - productive and potentially paid participation; - inclusive industry and innovation using existing infrastructure; - reducing entry barriers between large and small participants; - reuse, repair and responsible production decisions; - multi-stakeholder partnership and shared technical capability.

Even here, the project still needs evidence. A training session is not automatically a decent job. A reused material is not automatically environmentally better. A small business connected to a platform is not automatically earning more. A partnership logo is not an implemented partnership.

The SDG mapping therefore asks for results rather than affiliation.

Application-dependent Goals should stay application-dependent

The mobility and baggage branches may create evidence relevant to SDG 3 or 11. Farm applications may create evidence relevant to SDG 2, 6, 13 or 15. Energy measurement may connect to SDG 7. Material leakage or waste handling may create a guardrail relevant to SDG 14.

The correct approach is to activate the relevant measurement when the Work enters that domain.

That avoids a common reporting problem: using a broad platform vision to claim outcomes that only a future project might create.

15.2 ESG is an operating lens, not a replacement for the SDGs

The SDGs describe public development outcomes. ESG is used in this white paper as a practical management lens for the environmental, social and governance effects of a Work.

The two can support each other, but they are not interchangeable.

Environmental: what physical burden changed?

For each material or physical Work, ask: - What material entered? - What new material was added? - What scrap or failed sample was created? - What energy was used? - What transport or repositioning was required? - Was an existing asset repaired, reused or kept in service longer? - What happens at end of use? - Did the intervention shift burden to another site, material or energy source?

Where GHG claims matter, state whether the number is: - an organization inventory quantity; - a product life-cycle inventory; - a project/intervention estimate versus a counterfactual baseline; - or an avoided-emissions estimate disclosed separately.

This is more credible than using a generic “green” label.

Social: who gains capability, income, access or dignity?

Ask: - Who can participate who could not participate before? - What skill was gained? - Can the person complete a task independently? - Is the contribution paid? - How long does payment take? - Does the new process reduce physical burden? - Does it increase surveillance or pressure? - Does it remove an income source without an alternative? - Can people with different languages, devices or abilities use the service? - Can a participant refuse or withdraw without losing unrelated access?

The social ledger should include negative outcomes and complaints, not only successful stories.

Governance: who decides, who can inspect, who can correct?

Ask: - Who owns or controls the Resource? - Who can read the relevant data? - Who can authorize a physical action? - Who can stop it? - Who accepts the result? - Who can dispute a payment or quality result? - Can the evidence be reconstructed later? - Are temporary credentials separated from permanent operation identity? - Are participant roles and conflicts of interest visible? - Is a failed or unsafe outcome preserved rather than erased from the story?

This governance layer connects directly to Part III: **Human Authority is a real-world decision relationship; protocol authorization is only one technical mechanism that helps enforce it.**

15.3 No One Left Behind becomes measurable when distribution is visible

The United Nations' Leave No One Behind guidance emphasizes identifying who is excluded and why, meaningful participation, evidence, transparency, privacy and accountability.

That means an HSIMC project should not report only the average result.

If ten firms benefit but the smallest firms cannot afford the required reader, the distribution matters.

If an airport service becomes faster for most passengers but harder for a traveler without a smartphone, the distribution matters.

If a local maker gains orders but must absorb unpaid sampling and rework, the distribution matters.

If an automation reduces heavy lifting but also removes a worker's paid hours, the distribution matters.

A practical LNOB review can ask:

Who was expected to benefit?

Who actually participated?

Who could not participate, and why?

Who received cash or paid work?

Who contributed unpaid time or equipment?

Who gained decision power?

Who carried new risk, monitoring or compliance burden?

What alternative access path existed for people without the preferred device, language, account or skill?

Were complaints and failures recorded?

What changed in the next version because of that evidence?

Relevance is not achievement

The most important reporting rule in this chapter is simple:

SDG relevance is not SDG achievement. ESG intention is not ESG performance. Evidence of one Work is not evidence for the whole platform.

A credible public case should therefore label the evidence stage.

For value outcomes, a useful sequence is:

Proposed -> Baseline measured -> Work accepted -> Outcome observed -> Comparative result -> Independently reproduced

This value sequence complements the engineering sequence from Part IV:

Documented -> Simulated -> Remote simulated -> Physical reference -> Physical delivered -> Operational

Together, the two sequences prevent a common failure of technology storytelling: an idea, a prototype, a deployment and a sustained outcome are all presented as if they were the same thing.

A contribution statement template

A public statement can be short and still be rigorous:

Work: [what was done]
Participants: [who contributed / who accepted]
Baseline: [what happened before]
Observed result: [measured change]
Cash effect: [realized / forecast / none demonstrated]
Capacity effect: [minutes, throughput, access, etc.]
Environmental effect: [measured flow and method]
Social effect: [who benefited / who carried burden]
Governance evidence: [Authority, custody, acceptance, dispute]
SDG relevance: [specific Goals and why]
Limits: [what this case does not prove]

This format is deliberately modest. It does not need a complicated rating to make the result useful.

The value of partnership is an executed interface

Goal 17 is especially important to this white paper because it explicitly recognizes multi-stakeholder partnerships that mobilize knowledge, expertise, technology and resources.

For HSIMC, however, the strongest evidence of partnership is not the number of organizations shown on a slide.

It is: - one partner contributes a real capability; - another contributes a real demand; - Authority is clear; - a Work is accepted; - interfaces are understood; - the Work is completed; - value and burden are measured; - another team can reproduce or improve it.

That is the point where partnership becomes infrastructure.

The invitation after Part V

The global invitation can now be stated with a stronger meaning:

Bring your existing capability.
Deliver one reproducible function.
Measure one useful outcome.
Make GenAI plans become real value - now, here, globally.

The next Part asks how this method can become an open network rather than a collection of isolated pilots: how technology ecosystems, local makers, institutions and future partners can join without losing local Authority, how reference results can mature into shared standards, and how a person or organization can begin with one specific contribution.

Source notes for Part V

The following sources support the external factual context used in this Part. HSIMC application journeys and accounting structures are proposals unless explicitly described as observed or externally deployed.

United Nations, Transforming our world: the 2030 Agenda for Sustainable Development. Source for the 17 SDGs, the integrated and indivisible framing, the pledge that no one will be left behind, and Goal 17 partnership/data language. https://sdgs.un.org/2030agenda

United Nations Sustainable Development Group, Leave No One Behind. Source for meaningful participation, evidence, data, transparency, privacy and accountability in LNOB implementation. https://unsdg.un.org/2030-agenda/universal-values/leave-no-one-behind

GHG Protocol, Product Standard. Source for product life-cycle accounting across raw materials, manufacturing, transportation, storage, use and disposal. https://ghgprotocol.org/product-standard

GHG Protocol, Inventory and Project Accounting: A Comparative Review. Source for the distinction between organization inventory accounting and project/intervention accounting, including separate treatment of avoided-emissions analysis. https://ghgprotocol.org/blog/inventory-and-project-accounting

GHG Protocol, Estimating and Reporting Avoided Emissions. Source for neutral disclosure of positive and negative comparative GHG impacts and the difficulty of credible avoided-emissions claims. https://ghgprotocol.org/estimating-and-reporting-avoided-emissions

IATA, 2026 Global Passenger Survey. Source for the reported 49% global use of AI for travel planning and survey scope of more than 10,000 responses from over 200 countries. https://www.iata.org/en/pressroom/2026-releases/10-28-iata-2026-global-passenger-survey-travelers-biometrics-ai-entry-requirements/

Hong Kong International Airport, Baggage Pick-up and Delivery. Source for current paid airport-to-hotel/home/downtown baggage delivery services by listed providers. https://www.hongkongairport.com/en/passenger-guide/airport-facilities-services/baggage-pick-up-and-delivery

Hong Kong International Airport, MyTAG / MyTAG Support / MyTAG Privacy Policy. Sources for RFID baggage-arrival notification, QR registration/linking and the data associated with the service. https://www.hongkongairport.com/iwov-resources/html/MyTag/en/mytagseries.html

Hong Kong International Airport, About Us. Source for HKIA’s description of real-time trolley monitoring/tracking and other smart-airport capabilities. https://www.hongkongairport.com/en/about-us/?section=our-history

Part VI - The Invitation: A Global Network of Existing Capabilities

Part VI at a glance

Parts I-V built the argument step by step. Part I described the distance between a generated answer and a verified physical outcome. Part II showed that useful participation can begin with materials, tools and skills people already have. Part III defined a collaboration model around people, Authority, Tools, Resources, Work and Evidence. Part IV turned that model into an engineering path that allows cloud simulation, hardware design, local making and reference-board validation to proceed in parallel. Part V asked how useful value should be measured without counting the same benefit several times.

Part VI asks the final question: **how does this become a global network that people can actually join?**

The answer is not to build one giant factory, one closed app, one universal robot or one central company that owns every capability. The opportunity is to let existing capabilities become easier to describe, request, combine, verify and reuse - while keeping the people who own or operate those capabilities in control of their participation.

That network can include a person with a sewing machine, a print shop, a warehouse, a UWB engineer, a device maker, a cloud developer, a hotel operator, a logistics company, a university, a standards expert, a global AI platform, or a manufacturer with an established supply chain. The contribution of each participant may be very different. What makes the network coherent is not uniform ownership. It is a common way to state: - what capability is available; - under which conditions it can be used; - who has Authority to commit it; - what result is expected; - what evidence proves completion; - what value and cost were created; - which version another team can reproduce.

This Part develops that invitation in three chapters.

Chapter 16 describes an open network of complementary capabilities. It explains why HSIMC aims to complement model, cloud, creative, device, manufacturing, logistics and local-service ecosystems rather than replace them. It also defines a low-barrier participation unit: one capability, one accountable owner, one result and one evidence method.

Chapter 17 explains the route from a successful reference result to a shared specification. It deliberately distinguishes a local Profile or HSIMC reference method from a formal standard. A useful standard emerges only when repeated use, independent implementations, compatibility evidence, stakeholder need and a legitimate consensus process justify it. Formal international standardization belongs to recognized standards organizations and their procedures, not to a platform declaration.

Chapter 18 closes the English master with a practical invitation. A participant can bring a real task, a local capability, an engineering function, purchasing or sponsorship, or independent reproduction and review. The first objective is not to join a grand alliance. It is to complete one useful Work in a form that another authorized participant can understand and verify.

The invitation of this Part is therefore:

Bring your existing capability.
Deliver one reproducible function.
Measure one useful outcome.
Make GenAI plans become real value - now, here, globally.

Chapter 16 - An Open Network of Complementary Capabilities

16.1 Complement ecosystems instead of rebuilding them

The physical world already contains enormous capability.

Factories already have machines. Print shops already know materials, registration and finishing. Hotels already have staff, rooms, storage areas and service procedures. Airports already have baggage systems, carts, contractors and operational boundaries. Logistics companies already have vehicles, hubs, tracking systems and commercial terms. Universities already have laboratories, students and specialist equipment. Cloud companies already provide identity, storage, compute and collaboration. Creative ecosystems already provide design tools and content workflows. Device ecosystems already provide phones, cameras, sensors and operating systems. AI platforms already provide models capable of planning, coding, translation, analysis and increasingly sophisticated engineering support.

The challenge is not that all of these capabilities are absent.

The challenge is that they are difficult to combine quickly across organizational, technical and geographic boundaries.

A customer may know the desired result but not the available local machine. A model may generate a drawing but not know the exact fixture, material condition or operator constraint. A supplier may have the right equipment but no structured way to expose a narrow time window to an outside project. A hotel may have a person available for a luggage handover but no way for a travel assistant to know the exact service condition. A small maker may be able to produce a functional sample but have no software system through which a global engineering team can discover and trust that capability.

HSIMC therefore does not need to replace the ecosystems that already work.

Its proposed role is to help **connect a model-level plan to an accountable chain of local capabilities**.

This is complementary by design.

A model provider can keep improving reasoning, coding, engineering and multimodal capability. A productivity cloud can remain the place where a team stores and reviews files. A creative platform can remain the place where a designer develops artwork, packaging or content. A device ecosystem can remain the trusted interface through which a user scans, taps, measures or authorizes. A manufacturer can keep its own ERP, machines and quality system. A logistics company can keep its fleet and operational platform.

HSIMC adds value only if it reduces the friction between them.

The OpenAI engineering signal

OpenAI's 3 November 2026 publication, *Accelerating engineering for the physical world*, describes work spanning mechanical and electrical engineering, simulation, specialized tools and end-to-end system design. It emphasizes the need to bring together mechanical designs, electronics, software and manufacturing requirements, and it invites collaboration to understand where these capabilities can make the biggest difference and what still needs to improve.

HSIMC treats that publication as an important external signal, not as an endorsement of HSIMC and not as evidence of an official partnership.

The complementary question is:

When a model can prepare, analyze or coordinate more of the engineering work, how does the result reach a real machine, a real material, a real operator and a real customer - under explicit Authority and with evidence?

That is where local capability becomes strategically important.

The stronger models become, the more valuable a well-described real-world interface can become. A machine that is difficult for software to understand remains isolated. A workshop whose capability is documented with clear conditions can participate in more workflows. A local operator who can return structured evidence can become part of a distributed engineering team without becoming a software company.

The MCP signal

The 28 July 2026 Model Context Protocol specification introduced a stateless protocol core. At the protocol layer, requests can be self-describing, and servers no longer need hidden session state merely to understand which protocol context is being used. Applications that need state can still use explicit handles and persistent business records.

This is important because it reduces the architectural pressure to force all collaboration into one long-lived connection.

A client can request a tool in one place. Another authorized client can later refer to the same Work, Resource, version or operation through explicit identifiers. A server can expose a tool or resource without pretending that the underlying factory, person or asset is stateless.

The physical world remains stateful. The protocol can become easier to route and scale.

That combination supports a network in which different teams can participate at different times without depending on one hidden conversation history.

Complementarity should be visible in the Work

A collaboration is genuinely complementary when each participant can identify a distinct contribution.

For example, a functional card prototype may involve: - an AI model helping draft a pattern and test plan; - a creative tool producing artwork; - a local print shop producing registration features; - an embroidery workshop placing conductive thread; - an electronics provider supplying a module; - a T736 or another local compute node running an inspection

adapter; - a remote engineer reviewing image or electrical evidence; - a customer accepting the result; - a document space preserving the versioned package.

No participant needs to own the entire stack.

The contribution becomes valuable because it is **composable**.

The same principle applies to movement. One party may own the baggage. Another may provide airport transport. Another may provide hotel storage. Another may provide final delivery. A QR or NFC entry may help the traveler express intent. A model may explain options in the traveler's language. A local operator still has to accept the final Work.

The network is strongest when every link can answer:

What did I contribute? Under what Authority? For which Work? What result did I produce? What evidence did I return? What did I receive in return?

16.2 One capability is enough to enter

A global network often becomes inaccessible because the entry requirement is too large.

A small producer is told to install an enterprise platform, expose an API, build a digital twin, integrate a marketplace, provide a full product catalogue and accept standardized commercial terms before the first useful collaboration can occur.

That approach excludes precisely the people and businesses that may hold valuable local capability.

HSIMC proposes a smaller entry unit.

A participant can begin with **one capability note**.

A useful capability note should answer six practical questions:

What can you do?

Where can you do it?

When is it available?

Who is authorized to accept Work?

What result can another party inspect or accept?

What is the expected cost, time or participation condition?

This can be enough to begin a real conversation.

The structured system can grow around the Work rather than demanding that the participant understand the entire architecture first.

From capability note to Capability Profile

If a capability proves useful, the note can become more structured.

A Capability Profile may include: - provider identity; - responsible contact or role; - location or service area; - input types; - output types; - versioned process description; - accepted units and tolerances; - machine / tool / material Profile; - preparation requirements; - Authority boundary; - booking or availability rule; - expected lead time; - evidence returned; - pricing or compensation method; - known limits; - safety or regulatory conditions; - recovery or dispute path.

The provider does not have to expose every internal detail.

A machine shop may reveal only that a specific process can be provided within a tolerance and size range. A hotel may reveal only that luggage storage and an accepted delivery window are available. A person may expose a two-hour service window without exposing their entire calendar. A manufacturer may publish a capability Offer without exposing inventory quantities until a qualified Work request exists.

Discoverability does not require total transparency.

People without APIs can still participate

A Capability Profile does not imply the provider must run an MCP server.

The provider can interact through: - a card or QR/NFC entry; - a mobile web page; - an email or message translated into structured fields; - a local operator interface; - a partner organization; - an HSIMC assistant; - an existing ERP or service system through an adapter; - eventually, a native MCP tool or Resource if that adds value.

The important point is not the transport mechanism.

The important point is that the Work becomes understandable, authorized and verifiable.

This is essential to the principle of **No One Left Behind in the GenAI Physical Era**. A network should not require a participant to possess the most advanced software stack merely to contribute a physical skill.

Participation can be local while discovery is global

Physical capability has geography.

A laser machine in Kwun Tong cannot cut a sample in Boston without moving either the machine or the material. A hotel bell desk cannot receive a bag in another city. A local sewing machine cannot embroider a physical textile across the network.

But the **description, method, contract and evidence pattern** can travel globally.

This creates two different scaling speeds: - physical capability scales through local partners and logistics; - validated methods and interface Profiles can scale much faster through software and documentation.

A global network therefore does not mean every physical Resource serves every customer directly.

It means that a remote team can discover a relevant local capability, understand its conditions, send a versioned Work request, receive evidence, and - where appropriate - reproduce the same method through another local provider.

A capability can belong to more than one ecosystem

A local printer may serve packaging, electronics, exhibitions, education and hospitality.

A UWB team may serve mobility, warehouse tracking, tourism, robotics and accessibility.

A T736-based node may act as an engineering reference board in one Work, a local interface node in another, and a portable compute node in a third.

The capability should not be trapped inside one vertical label.

This is one reason HSIMC uses a Resource / Capability / Work model rather than a permanent seller category.

A person is not permanently a supplier. A company is not permanently a customer. Roles can change by Work.

That is an application-level participation model. It is distinct from the statelessness of MCP at the protocol layer, but the two ideas complement each other: **explicit context makes changing participation easier to represent without relying on hidden session assumptions.**

16.3 Participation must create fair reasons to stay

An open network fails if participation is easy but value distribution is opaque.

People will not continue to contribute equipment time, design, local labour, test data, customer access or risk unless they can understand what they are receiving and what they are giving up.

A useful Work therefore needs more than a technical acceptance result.

It should also make visible: - who contributed cash; - who contributed unpaid time; - who provided equipment or space; - who accepted operational responsibility; - who contributed intellectual work; - who gained a reusable method; - who earned revenue; - who gained access to a new customer or market; - who carried a new safety, maintenance or support burden; - which rights apply to future reuse.

Part V introduced the principle of avoiding double counting. Part VI applies the same discipline to partnership.

A network cannot claim the same contribution as the exclusive achievement of every participant.

Contribution rights should travel with the Work

The Work record can become the natural place to state contribution rights.

For example: - a maker may allow the sample result to be published but keep the fixture design private; - an engineer may allow an adapter to be reused under specified terms; - a customer may allow anonymized performance data but not images of the site; - a university team may require attribution; - a sponsor may fund a public reference implementation; - a commercial partner may require a defined exclusivity period; - a local operator may require payment for every execution even if the software method is open.

There is no need to force every collaboration into one intellectual-property model.

The requirement is that reuse rights are not left implicit.

Five participant roles can coexist in one Work

A practical global initiative can invite at least five kinds of contribution:

Problem owner - brings a real need and acceptance condition.

Local capability provider - brings people, space, equipment, material or service.

Engineering / model contributor - brings software, analysis, design, sensing or integration.

Buyer / sponsor - funds the Work, the infrastructure or the public reference result.

Independent reproducer / reviewer - tests whether the result can be repeated and whether the evidence is sufficient.

The same organization may hold several roles, but the roles should still be distinguishable.

A project becomes more trustworthy when the person who designed the solution is not automatically assumed to be the only person who can verify it.

Goal 17 is a useful frame, not a badge

United Nations Sustainable Development Goal 17 explicitly calls for multi-stakeholder partnerships that mobilize and share knowledge, expertise, technology and financial resources. It also emphasizes data, monitoring and accountability.

That is highly relevant to a network of complementary capabilities.

But relevance is not achievement.

A partnership should be able to show actual commitments, actual contribution, actual outcomes and actual learning.

For HSIMC, the strongest possible Goal 17 story is therefore not a partner-logo page.

It is a chain of completed Work in which different participants can show how knowledge, technology, resources and payment moved through an accountable interface.

The value of partnership is an executed interface.

Chapter 17 - From Reference Results to Shared Standards

17.1 A successful demo is not yet a standard

The word *standard* is attractive because it suggests stability, compatibility and scale.

It is also easy to misuse.

A working prototype is not a standard. A company specification is not automatically an industry standard. A widely used API is not automatically an International Standard. A test passed once in one laboratory does not establish universal compatibility.

This white paper therefore uses a deliberate progression.

A capability may mature through the following levels:

Reference result -> Versioned Profile -> Conformance method -> Multi-implementation compatibility -> Shared specification -> Formal standardization, where justified

Each level adds evidence and participation.

Reference result

A reference result is a reproducible example under known conditions.

For locate_marks, a reference result might include: - a versioned marker definition; - a known input image; - an expected output; - an algorithm version; - a test fixture or calibration reference; - acceptance criteria; - a receipt showing the result.

A reference result is useful because it gives another team something concrete to reproduce.

It is not yet a claim that every camera, every substrate, every lighting condition or every target device will behave the same way.

Versioned Profile

A Profile defines a supported subset or configuration.

For a physical interface it may specify: - connector and pin assignment; - voltage and current constraints; - mechanical envelope; - protocol or message format; - timing assumptions; - required safety state; - recovery procedure; - known exclusions.

For a process it may specify: - accepted materials; - machine family; - fixture version; - setup steps; - inspection method; - tolerances; - evidence output.

Profiles allow different implementations to be compared without pretending that every optional capability must be supported.

This is especially important for T736. Part IV described a minimum common route - power, recovery, maintenance UART, a chosen network path, USB Host and intentionally assigned expansion - while keeping other interfaces and electrical assumptions dependent on the actual BOM and board configuration. A Profile prevents a block diagram from becoming an unsupported promise.

Conformance method

A shared interface becomes more useful when there is a test method that determines whether an implementation satisfies the stated Profile.

A conformance pack may contain: - test vectors; - expected responses; - required error behavior; - timing limits where relevant; - side-effect rules; - recovery tests; - Authority tests; - version negotiation tests; - evidence format requirements; - negative tests for unsupported operations.

Conformance is not the same as quality certification for an entire product.

It answers a narrower question:

Does this implementation satisfy the requirements of this Profile under the specified test conditions?

ISO’s terminology similarly distinguishes conformity assessment as the process of determining whether something meets the requirements of a standard. Depending on the context, assessment can be performed by the supplier, purchaser or an independent third party.

Multi-implementation compatibility

The next step is not to add more documents. It is to observe whether independent implementations work together.

Two teams should be able to implement the same Profile without sharing hidden assumptions.

Useful evidence includes: - team A client with team B server; - team B client with team A server; - different hardware adapters returning equivalent evidence; - different local makers reproducing the same sample acceptance result; - old and new versions interacting according to the documented compatibility policy; - failure cases behaving in the agreed way.

A compatibility matrix can show what is actually known.

For example:

Profile

Implementation A

Implementation B

Implementation C

Evidence

Mark locate v1

Pass

Pass

Not tested

test pack R1

NFC URI entry v1

Pass

Pass

Pass

device matrix R2

UWB observation v0.2

Simulated

Physical reference

Not tested

separate evidence

T736 serial adapter v0.1

Documented

Not bound

Not bound

hardware evidence pending

This is more informative than the word *compatible* alone.

17.2 Shared specifications should grow from evidence

The objective of a shared specification is not to freeze innovation too early.

It is to stabilize the parts that multiple participants need to depend on.

A good candidate for shared specification usually has several characteristics: - more than one participant needs it; - multiple implementations are possible; - ambiguity repeatedly creates integration cost; - the interface is stable enough to document; - independent testing is possible; - versioning can be managed; - the benefit of compatibility is larger than the cost of standardizing.

Examples may include: - the fields required in a Capability Profile; - the structure of a Work acceptance receipt; - the distinction between operation receipt and fresh observation; - an evidence status vocabulary such as Documented / Simulated / Physical Reference / Operational; - a card or QR/NFC URI structure; - a device adapter Profile; - a reference test pack for a specific function; - a minimum error and recovery behavior.

Keep business semantics separate from transport where possible

A shared Work contract should not unnecessarily depend on one network transport.

For example, the business concept of work_id, resource_id, authority_context, operation_id, evidence_ref and acceptance_status can exist whether the execution request travels through MCP, an HTTPS API, a local queue or a human operator interface.

MCP can be a powerful interoperability layer for AI clients and tools, but HSIMC business semantics should remain clear enough to survive protocol evolution.

This is another reason the 2026-07-28 MCP move toward explicit, stateless requests is strategically useful. The model can see and pass explicit handles rather than relying on hidden transport session state. The application can preserve the real-world Work independently.

Versioning is a governance function

A shared interface without version discipline becomes a source of risk.

Every Profile should define: - version identifier; - effective date; - supported predecessor versions; - migration notes; - breaking vs non-breaking change rule; - deprecation window where appropriate; - test pack version; - evidence schema version.

The MCP 2026-07-28 release itself provides a useful example of why this matters: it introduced major breaking changes alongside a formal deprecation policy so implementers could plan migration instead of discovering incompatibility during execution.

HSIMC can borrow the discipline without pretending to be the MCP standards process.

A physical adapter may need even slower change than software because replacement requires parts, labour, shipping and safety re-validation.

Negative results belong in the shared knowledge

Standardization is weakened when only successful tests are preserved.

A useful compatibility record should also show: - which material failed; - which camera did not meet the required latency; - which NFC tag could not be rewritten; - which connector overheated; - which device firmware broke a previously working function; - which hotel workflow caused repeated missed handovers; - which language version created user confusion; - which economic assumption failed under real pricing.

Failure is not noise.

Failure defines the boundary of the current Profile.

This is one reason Evidence is a first-class HSIMC concept.

17.3 From shared specification to formal standardization

Formal standards development should be treated with respect.

ISO explains that standards development begins in response to a market or stakeholder need, is carried out by relevant experts and stakeholders, and is consensus-based. ISO’s procedures explicitly state that consensus requires substantive objections to be heard and addressed; consensus does not necessarily mean unanimity.

That process is very different from a company publishing a PDF and calling it a global standard.

HSIMC should therefore use precise language.

Possible labels include: - **HSIMC reference method**; - **HSIMC Capability Profile**; - **shared implementation guide**; - **reference test pack**; - **interoperability profile**; - **consortium specification**, if a real consortium approves it; - **industry proposal**, if submitted into a wider process.

The phrase **International Standard** should be reserved for the recognized standardization process that produces one.

A practical standardization ladder

A useful ladder is:

Level 0 - Observation
One team records a problem and a result.

Level 1 - Reference method
The team publishes enough information for another team to reproduce it.

Level 2 - Versioned Profile
Scope, required fields, interfaces, limits and evidence are defined.

Level 3 - Conformance pack
A test method distinguishes conforming and non-conforming implementations.

Level 4 - Independent implementations
Multiple teams implement the Profile and publish compatibility evidence.

Level 5 - Shared specification
Multiple stakeholders jointly maintain the Profile under an explicit governance process.

Level 6 - Formal standardization, when justified
A recognized standards body or appropriate industry standards process considers the work through its own procedures.

This ladder protects both speed and credibility.

A local team can move quickly at Levels 0-3. A wider ecosystem can join at Levels 4-5. Formal standardization can begin only when the need and stakeholder base justify the additional process.

Open reference does not require open everything

A reference implementation can be public while a production implementation remains proprietary.

A test vector can be open while a supplier keeps its manufacturing method private.

A Profile can describe an electrical interface without revealing a company’s internal PCB.

A Work contract can expose acceptance conditions without exposing a customer’s private operational data.

The objective is not compulsory openness.

The objective is enough shared information to support safe interoperability and reproducible acceptance.

Global relevance may still require local options

ISO’s directives explicitly recognize that a globally applicable solution may require justified options to address legitimate differences across markets.

The same principle is important for HSIMC.

A luggage handover method may need different hotel practices in Hong Kong, Bangkok or Boston. A power interface may need different certified components. A mobility application may face different regulatory boundaries. A local maker may use different machinery while still meeting the same output specification.

A strong Profile distinguishes: - what must remain invariant for interoperability; - what may vary locally; - how that variation is declared; - what evidence is needed for the local option.

That is how a network can be global without erasing local capability.

Shared standards should make local differences legible, not invisible.

Chapter 18 - An Invitation to Make AI Real

18.1 Five ways to join

A global initiative becomes real only when participation produces Work.

HSIMC therefore does not need a participant to begin by signing up for a large platform relationship.

A person or organization can begin in one of five ways.

1. Bring a real task

A real task is one where someone cares about the result.

Examples include: - make a small batch of a versioned card or functional interface; - inspect a printed or embroidered sample; - move a defined object between two local points; - prepare a T736 peripheral adapter; - test a UWB observation Profile in a specific space; - clear an airport-style trolley handover point in a simulated or authorized real setting; - arrange an accepted luggage handover at a hotel pilot; - adapt an existing wheelchair or cart only within a properly authorized engineering test scope.

The task should state who will accept the result and what evidence is sufficient.

A problem without an acceptor is still research.

Research is valuable, but it should be labelled honestly.

2. Bring a local capability

A local capability can be: - a sewing or embroidery machine; - a printer; - a laser machine; - a CNC machine; - a test instrument; - a storage area; - a delivery route; - a vehicle; - a UWB kit; - a camera; - a T736 board; - an engineering skill; - a translation skill; - a person who can inspect, pack, deliver or accept.

The first question is not “How advanced is it?”

The first question is:

What useful result can this capability produce, under what conditions, and who can verify it?

3. Bring an engineering or model function

Software and model contributors can join before final hardware exists.

Useful contributions include: - schema and contract design; - simulation; - test generation; - image inspection; - CAD review; - electronics review; - device adapters; - error handling; - multilingual interaction; - data normalization; - scheduling; - evidence analysis; - reproducibility tooling.

The strongest contribution is not a demo that only works in one developer’s environment.

It is a versioned function another team can execute against the same contract.

4. Buy or sponsor a useful result

A network needs buyers and sponsors as much as it needs engineers.

A company may buy a local prototype. A hotel may pay for a service test. A manufacturer may sponsor an adapter for an existing machine. A foundation may sponsor accessible participation. A technology company may sponsor a public reference test. A government or university programme may sponsor a trial under its own procurement and governance requirements.

Sponsorship should not automatically buy control over every contributor or every later use.

The Work should state what the funding covers, what result becomes public or private, and what reuse rights remain with contributors.

5. Reproduce and review

Independent reproduction is one of the most valuable forms of participation.

A reviewer can ask: - Can I obtain the same versioned package? - Can I reproduce the stated result? - Which assumptions were missing? - Which evidence was insufficient? - Did the result depend on one specific person? - Which local change was required? - Did the economic or ESG claim survive the second site?

A network grows stronger when a second team can improve the first team’s method instead of merely believing it.

18.2 Begin with one Work, then earn the right to scale

Large visions are easy to state.

Global execution networks, intelligent manufacturing, autonomous logistics, accessible mobility, local-to-global production and GenAI-native commerce all sound important.

The discipline is to begin with a Work small enough to finish and rich enough to teach the network something.

A useful first Work has: - a named requester; - a named provider; - one or more Resources; - a clear Authority boundary; - a versioned input; - a result condition; - a realistic time window; - direct cost recording; - a simple evidence method; - an acceptor; - a known recovery path.

The first 30 days: make the interface real

The first phase should answer: - Is the problem real enough that someone cares about the result? - Can the capability be described clearly enough for another person to request it? - Is Authority explicit? - Can a useful versioned sample or service result be completed? - Can cost and time be recorded? - Can the evidence be preserved and retrieved?

The first phase does not need full automation.

A human operator can remain in the loop.

Automation should enter where repeated evidence shows that it improves the outcome.

Days 31-60: reproduce and compare

The second phase should attempt at least one form of independence: - another engineer uses the same contract; - another device runs the same function; - another maker reproduces the sample; - another location implements the same service pattern; - another client retrieves the same Work and evidence.

Differences should be documented, not hidden.

This is where a reference result begins to become a Profile.

Days 61-90: decide whether to scale, standardize or stop

The third phase should compare: - actual cash result; - actual released capacity; - user/service outcome; - environmental/material effect; - governance and evidence quality; - repeatability; - support burden; - next-site adaptation cost.

The decision can then be honest.

The Work may be ready to scale.

It may need another engineering iteration.

It may be valuable only in a narrow local niche.

It may not be worth continuing.

A network that can stop weak ideas is more credible than a network that converts every pilot into a success story.

The first visible global test

One of the strongest tests for HSIMC is simple:

Can a second team, in another network or location, receive a versioned package and complete the agreed function without depending on the original conversation?

This is where Stateless MCP, versioned files, explicit Resource identities, Authority and Evidence become more than architecture words.

They become a test of whether knowledge can travel while responsibility remains clear.

The same test can be applied to the card-phone / T736 path:

cloud simulation -> remote reference board -> frozen version -> shipment -> local binding -> continued local execution

The shipped hardware is only one stage in the Work.

The more important asset is the combination of Profile, software, evidence, recovery method and local Authority that lets the hardware continue to be useful after shipment.

18.3 Make AI real - without leaving people behind

The physical era of GenAI can develop in more than one direction.

One direction concentrates capability into increasingly large organizations, increasingly specialized infrastructure and increasingly automated systems. That can produce extraordinary efficiency.

Another direction uses the same intelligence to make smaller existing capabilities easier to discover, combine and improve.

The two directions do not have to be enemies.

Global scale and local participation can complement each other.

A powerful model can help a small workshop understand a drawing. A cloud platform can help preserve the version. A creative tool can help prepare the interface. A local machine can make the sample. A person can inspect it. A sensor can return evidence. A global engineering team can improve the design. A second maker can reproduce it. A customer can accept the result.

The opportunity is not that every household becomes a semiconductor fab.

The opportunity is that **more people can contribute a useful, bounded function to a sophisticated product or service**.

That is a different kind of industrial participation.

It is closer to a network of capabilities than to a hierarchy of factories.

No One Left Behind is a design requirement

The phrase should create practical questions.

Can a person join without speaking English?

Can a participant contribute without owning expensive compute?

Can a small workshop join without deploying an enterprise API?

Can an older worker use a card, QR or simple web interface instead of learning a developer console?

Can a person with a disability express Authority and receive a service in an accessible way?

Can a local maker see what will be paid before accepting Work?

Can a contributor withdraw or limit reuse?

Can a buyer see the actual evidence rather than only a sustainability claim?

Can a worker’s saved time create a benefit for that worker, not only for the organization?

Can a failed Work be reported without punishing the person who surfaced the failure?

If the network cannot answer those questions, it has not yet achieved the principle.

Local capability should remain local enough to matter

Global participation should not turn every contributor into an anonymous subcontractor competing only on price.

A local capability has context.

It may know the building, the machine, the customer, the language, the material, the neighborhood, the hotel, the route or the cultural expectation.

That context can be part of the value.

The purpose of a common Profile is not to erase local knowledge.

It is to make the transferable part explicit enough that local knowledge can be combined with global engineering.

Human Authority should remain visible as automation increases

As tools become more capable, the system may automate more planning, scheduling, inspection and execution.

That increases the importance of knowing: - which person or organization authorized the automation; - which Resource is being used; - which constraints apply; - which side effects are permitted; - how the action can be stopped; - who accepts the result; - who is accountable when the result is uncertain.

A faster autonomous tool should not make Authority less visible.

It should make Authority easier to state and audit.

Evidence should make trust portable

Trust in physical work is often local and personal.

A customer trusts a particular operator. A factory trusts a particular technician. A hotel trusts a particular contractor. A maker trusts a familiar supplier.

The purpose of Evidence is not to replace those relationships.

It is to make part of that trust portable.

A second team can see: - the version; - the test; - the result; - the conditions; - the responsible party; - the known limits.

Portable trust makes global participation more practical without requiring blind centralization.

The credit-card-sized starting point returns

The white paper began with a familiar physical footprint.

That small rectangle was never the final product definition.

It is a symbol of a larger engineering idea:

make the entry small enough to understand, then let capability extend beyond the object.

A card can carry identity, a QR or NFC entry, artwork, a functional surface, a sensor or a connection to a larger compute node. It can point to a T736-based system, a machine, a hotel service, a material Profile, a local maker or a global engineering package.

The intelligence does not have to fit entirely inside the card.

The participation path can begin there.

The final invitation

HSIMC proposes a practical response to the GenAI physical era:

Do not begin by asking everyone to buy a new machine.
Begin by asking what capability already exists.

Do not begin by asking every small participant to build a platform.
Begin by defining one function that another party can request and verify.

Do not treat a model plan as a completed result.
Connect it to people, materials, machines, time, Authority and Evidence.

Do not call every reference method a global standard.
Earn interoperability through independent implementations, testing and consensus.

Do not measure value only in activity or publicity.
Measure the useful outcome, the cost, the distribution of benefit and the evidence.

And most importantly:

Bring your existing capability.
Deliver one reproducible function.
Measure one useful outcome.
Make GenAI plans become real value - now, here, globally.

That invitation can begin with stone, thread, a printer, a machine, a card, a piece of code, a warehouse, a trolley, a hotel, a reference board or one person's practical skill.

The starting point can be small.

The collaboration does not have to be.

Reference notes for Part VI

The following sources support the external context used in this Part. They do not imply endorsement of HSIMC.

1. **OpenAI - Accelerating engineering for the physical world (3 November 2026).** Describes AI-assisted mechanical/electrical engineering, simulation, cross-disciplinary coordination and collaboration around real engineering workflows.
   https://openai.com/index/accelerating-engineering-for-the-physical-world/
2. **Model Context Protocol - The 2026-07-28 Specification (28 July 2026).** Describes the stateless protocol core, explicit per-request context, routing headers, cacheable lists, authorization hardening, extensions and explicit state handles where applications need continuity.
   https://blog.modelcontextprotocol.io/posts/2026-07-28/
3. **United Nations - SDG 17: Partnerships for the Goals.** Includes multi-stakeholder partnerships that mobilize and share knowledge, expertise, technology and financial resources.
   https://sdgs.un.org/goals/goal17
4. **United Nations - Multi-stakeholder partnerships.** Describes Goal 17's role for public, public-private and civil-society partnerships.
   https://sdgs.un.org/topics/multi-stakeholder-partnerships
5. **ISO - Developing standards.** Explains that ISO standards respond to stakeholder/market need, are developed by experts through a multi-stakeholder process and are consensus-based.
   https://www.iso.org/developing-standards.html
6. **ISO/IEC Directives, Part 1.** Defines consensus as an essential procedural principle and recognizes that globally relevant standards may require justified options for legitimate differences.
   https://www.iso.org/sites/directives/current/consolidated/
7. **ISO Glossary - conformity assessment.** Defines conformity assessment as determining whether something meets stated requirements and distinguishes first-, second- and third-party assessment.
   https://www.iso.org/glossary.html

End of Part VI - The Invitation: A Global Network of Existing Capabilities

---

Appendix A - The 17 SDGs: Relevance, Contribution, Evidence and Limits

The 17 Sustainable Development Goals are used in this white paper as a relevance and evidence map. This is not a claim that HSIMC has achieved any Goal. Each public case should identify the specific target or outcome pathway, the baseline, the observed result, who benefited, who carried burden, and what the case does not prove.

A.1 - 1 - No Poverty

Initial relevance: Application-dependent

Possible contribution pathway: Verified net income, resilience or reduced essential-service burden for low-income participants.

Minimum evidence: Participant baseline, actual net income/cost change, distribution of benefits and burdens.

Limit: Do not infer poverty impact from general job creation or revenue alone.

A.2 - 2 - Zero Hunger

Initial relevance: Application-dependent

Possible contribution pathway: Agriculture, food handling, storage or logistics Work that measurably reduces loss or improves access.

Minimum evidence: Quantity/quality of food, loss rate, time, cost, beneficiary and supply-chain records.

Limit: Not relevant to non-food Work merely because a farm is a partner.

A.3 - 3 - Good Health and Well-Being

Initial relevance: Application-dependent / safety-critical

Possible contribution pathway: Assisted mobility, reduced unsafe manual handling, or access to services.

Minimum evidence: Defined population, safety case, incident data, service outcomes and human oversight.

Limit: Cargo or logistics evidence never substitutes for person-carrying safety evidence.

A.4 - 4 - Quality Education

Initial relevance: Direct initial relevance

Possible contribution pathway: Capability learning, reproducible engineering tasks, second-team reproduction and skills transfer.

Minimum evidence: Training content, pre/post capability, independent completion, time to reproduce and retained skill.

Limit: Attendance alone is not learning evidence.

A.5 - 5 - Gender Equality

Initial relevance: Cross-cutting

Possible contribution pathway: Participation, pay, decision rights and burden should be visible by relevant groups.

Minimum evidence: Disaggregated participation, compensation, Authority, safety and advancement data where lawful and appropriate.

Limit: Do not assume neutral technology produces equal outcomes.

A.6 - 6 - Clean Water and Sanitation

Initial relevance: Application-dependent

Possible contribution pathway: Relevant when materials/processes use water, create wastewater, or serve water/sanitation Work.

Minimum evidence: Process-water inputs, discharge route, contaminants, treatment and baseline comparison.

Limit: No water claim without a defined water boundary.

A.7 - 7 - Affordable and Clean Energy

Initial relevance: Application-dependent

Possible contribution pathway: Energy efficiency, shared equipment utilization, renewable-powered local operation or lower energy per accepted result.

Minimum evidence: Metered energy, output quantity, baseline, grid/fuel context and additional equipment energy.

Limit: A small device is not automatically energy efficient at system level.

A.8 - 8 - Decent Work and Economic Growth

Initial relevance: Direct initial relevance

Possible contribution pathway: Paid contribution, safer work, higher capability, SME participation and productive time.

Minimum evidence: Actual pay, working time, overtime, safety, skill progression and incremental revenue/cost.

Limit: Released time is not automatically a wage saving or decent-work outcome.

A.9 - 9 - Industry, Innovation and Infrastructure

Initial relevance: Direct initial relevance

Possible contribution pathway: SMEs and local makers connect to versioned interfaces, digital tools, hardware profiles and reproducible engineering.

Minimum evidence: Accepted capability profiles, conformance results, second-site reproduction, utilization and business continuity.

Limit: A prototype alone does not prove resilient infrastructure.

A.10 - 10 - Reduced Inequalities

Initial relevance: Direct initial relevance

Possible contribution pathway: Lower barriers created by language, device, geography, company size, disability or lack of API access.

Minimum evidence: Who participates, who is excluded, alternative access path, Authority and benefit distribution.

Limit: Access to a webpage is not the same as meaningful participation.

A.11 - 11 - Sustainable Cities and Communities

Initial relevance: Application-dependent

Possible contribution pathway: Local repair, shared mobility, building/airport/hotel logistics, reuse and accessibility Work.

Minimum evidence: Local baseline, space/time effect, safety, service availability, material and transport impact.

Limit: City-scale claims require city-scale evidence.

A.12 - 12 - Responsible Consumption and Production

Initial relevance: Direct initial relevance

Possible contribution pathway: Reuse, repair, material traceability, measured yield, waste, life extension and responsible sourcing.

Minimum evidence: Material identity, mass flow, yield, scrap, lifetime, repair, disposal and comparative baseline.

Limit: Reuse is not automatically environmentally preferable.

A.13 - 13 - Climate Action

Initial relevance: Application-dependent

Possible contribution pathway: Comparative GHG reduction or climate-related capability where a defined baseline and boundary exist.

Minimum evidence: Recognized accounting method, activity data, emission factors, life-cycle boundary and uncertainty.

Limit: Avoided-emissions claims must remain separate from organizational inventory accounting.

A.14 - 14 - Life Below Water

Initial relevance: Guardrail / application-dependent

Possible contribution pathway: Relevant where material leakage, wastewater, packaging or coastal/marine applications create a plausible pathway.

Minimum evidence: Material/waste pathway, discharge, leakage risk and mitigation.

Limit: Do not claim marine benefit merely from recycling.

A.15 - 15 - Life on Land

Initial relevance: Guardrail / application-dependent

Possible contribution pathway: Relevant to sourcing, land use, agricultural Work, biodiversity effects and disposal.

Minimum evidence: Source, land/ecosystem context, material choice, operational effect and baseline.

Limit: Local reuse can still shift burden elsewhere.

A.16 - 16 - Peace, Justice and Strong Institutions

Initial relevance: Cross-cutting governance relevance

Possible contribution pathway: Transparent Authority, auditability, correction, redress, data responsibility and accountable decisions.

Minimum evidence: Permission trail, disputes, incident/recovery records, audit access and correction process.

Limit: Project governance is not evidence of institutional reform at national scale.

A.17 - 17 - Partnerships for the Goals

Initial relevance: Direct initial relevance

Possible contribution pathway: Executed multi-stakeholder interfaces that share knowledge, expertise, technology, resources and measurable Work.

Minimum evidence: Named roles, accepted commitments, delivered result, value/burden, reproducibility and data quality.

Limit: Partnership logos or MoUs alone are not executed partnership evidence.

A.18 Evidence rule for SDG statements

A useful SDG contribution statement should contain **Work, participants, baseline, observed result, distribution, method, evidence stage, relevant Goal/target, and limits.** If one of these elements is absent, present the statement as a hypothesis or relevance assessment rather than verified impact.

A.19 Strongest initial cluster

The initial cluster with the most direct relationship to the operating model is SDG 4, 8, 9, 10, 12 and 17. These correspond to learning and capability, decent work and small enterprise, inclusive industrial participation, reduced barriers, responsible production and multi-stakeholder partnership. Even for these Goals, results remain Work-specific.

Appendix B - ESG Measurement Framework

ESG in this white paper is an operating measurement framework, not a substitute for the SDGs and not a single composite score.

B.1 Five ledgers

Cash ledger

Records realized revenue, verified avoided cash expense, new operating expense and initial investment. Typical unit: currency. Do not confuse forecast value, released time or an internal payment with newly created cash value.

Released-capability ledger

Records waiting, setup, rework, throughput, staff time made available and second-team reproduction. Typical units: minutes, hours, units and jobs. Released time becomes cash saving only when actual cash changes.

User/service-outcome ledger

Records timing, choice, access, error, disruption, completion, dignity or convenience where these are defined service outcomes. Typical units include percentages, minutes, counts and survey measures.

Environmental/material ledger

Records material input/output, reuse, scrap, energy, transport, repair, lifetime and disposal. Typical units include kg, kWh, km and cycles. Material reuse is not automatically a net environmental benefit.

Governance/evidence ledger

Records Authority, custody, acceptance, incident, retry, dispute, recovery and provenance. Typical units are events, timestamps and completeness measures. Good governance evidence is not a substitute for technical performance.

B.2 Cash and network accounting

For participant *j* during a defined period:

Incremental cash contribution(j) = new realized cash revenue + verified avoided cash expense - new incremental operating cash expense.

Initial investment should be shown separately. If initial investment is non-zero, simple period ROI may be reported as:

Simple period ROI = (cumulative incremental cash contribution - initial investment) / initial investment.

This is not automatically an annualized return.

At network level, internal payments between included participants must be eliminated. If a hotel pays a local service provider HKD 100, one participant records -100 and the other +100; the network does not record HKD 200 of newly created value.

B.3 Baseline discipline

Every comparative claim should state the functional unit, boundary, baseline or credible counterfactual, observation period, allocation rule and uncertainty.

B.4 Environmental reporting

Product-level environmental work should consider the relevant life-cycle stages: raw materials, manufacturing, transportation, storage, use, repair/reuse where applicable, and disposal. Project/intervention claims such as avoided emissions should remain distinct from an organization’s greenhouse-gas inventory.

A local reuse project therefore records not only the reused material but also cleaning, processing, new fasteners/electronics, transport, failure, maintenance and end-of-life route.

B.5 Social reporting

Social reporting should make distribution visible. Minimum questions include who was expected to benefit, who actually participated, who could not participate and why, who was paid, who contributed unpaid time/equipment, whether physical or schedule burden moved to another person, what new independent capability remained, and whether an alternative path existed for people without the preferred device, language, account or technical skill.

B.6 Governance reporting

Governance evidence should show who could authorize each action, who accepted the Work, what version was used, what happened when the result was uncertain, how cancellation or recovery was handled, and how a participant could dispute or correct a record.

B.7 Scaling gate

Do not scale because a demo looked impressive. Scale when the marginal useful outcome continues to exceed the marginal burden, and when a second team or second location can reproduce the result under declared conditions.

Appendix C - Common Capability Contract

A Common Capability Contract is the shared boundary that lets different teams work before final hardware is available. It describes the function, not an unlimited permission to use a Resource.

C.Field - contract_id

Stable identifier for the function contract or profile.

C.Field - function

Human-readable action/result, e.g. locate_marks.

C.Field - input_schema

Versioned input objects, units, allowed ranges and required files.

C.Field - output_schema

Versioned result, quality fields, units, evidence references and error states.

C.Field - resource_target

HSIMC Resource or target class; actual physical instance is bound separately.

C.Field - execution_mode

simulation / physical_reference / physical_delivered / operational.

C.Field - conditions

Environment, tooling, calibration, material and prerequisite conditions.

C.Field - authority_required

Minimum Authority required to discover, read, reserve, configure, execute, stop, accept or administer.

C.Field - effect_identity

Digest or identifier that locks the intended physical effect independently of short-lived credentials.

C.Field - operation_id

Stable identity used for retry/read-back of a physical-effect request.

C.Field - freshness_requirement

Maximum acceptable age for observations used to authorize or plan the action.

C.Field - resource_locking

Reservation, lease, busy/conflict behavior and release rule.

C.Field - safety_boundary

Conditions that must remain local/deterministic and cannot be delegated to a remote model.

C.Field - acceptance

Measurement, tolerances, evidence and human/organizational acceptance step.

C.Field - recovery

Stop, unknown-result, interrupted-network and restart behavior.

C.Field - version

Contract, adapter, algorithm, hardware/BOM, calibration and deployment version references.

C.Field - cost_time

Expected and actual time, material, energy and monetary cost fields.

C.Field - rights

Contribution, reuse, attribution, confidentiality and withdrawal constraints.

C.1 Example - locate_marks


```
{
  "contract_id": "hsimc.locate_marks.v0",
  "function": "locate_marks",
  "input": {"image_ref": "sha256:...", "marker_profile": "marker-profile-v3", "calibration_ref": null},
  "target": {"resource_id": null, "mode": "simulation"},
  "authority": {"subject": "bound-at-runtime", "scope": ["read_input", "run_observation"]},
  "output": {"coordinate_space": "pixels", "marks": [], "quality": {}, "evidence_refs": []},
  "limits": ["No millimetre conversion without calibration evidence", "Observation only; no machine motion or laser action"]
}
```


The same contract can run against fixed images, a cloud simulator, a remote reference camera, or a delivered device. Each target still requires its own adapter and evidence.

C.2 Retry and effect rule

A retry should reuse the original operation_id when the requested physical effect is the same. If the outcome is unknown, query the original operation before creating a new effect. New credentials or a refreshed token do not by themselves create a new physical Work.

Appendix D - Physical Interface and Material Profiles

D.1 Profile principle

A Profile describes a possible capability under declared conditions. A physical instance must still be observed and bound before the profile becomes evidence about that particular device or material.

D.2 Material Profile - minimum fields

Identity

Source, supplier/product when known, batch/lot, previous use and owner.

Physical condition

Dimensions, mass, thickness, surface, defects, contamination and moisture.

Composition

Declared composition, coating, binder, conductive material and unknowns.

Process history

Cutting, heat, lamination, printing, cleaning and chemical treatment.

Intended function

Fixture, card layer, packaging, antenna substrate, non-structural sample or another declared function.

Tests and limits

Dimension, resistance, adhesion, abrasion, RF response, durability and other relevant tests; explicitly record unverified load, fire, hygiene, electrical, environmental or regulatory conditions.

Evidence

Photos, measurements, test report, operator, timestamp and version.

D.3 Functional-surface Profile

For printed, embroidered, foil or conductive surfaces, record geometry, line/thread/material type, process settings, substrate, bonding/curing, measured resistance or other functional property, connection method, inspection image and acceptance limit.

D.4 Hardware/adapter Profile

For cameras, NFC readers, UHF RFID readers, UWB modules, machine interfaces and motion controllers, record exact model/BOM/revision, connector and pin assignment, electrical standard and voltage levels, driver/OS/runtime dependencies, communication protocol and timing, power budget and peak behavior, calibration/coordinate frame where relevant, failure/timeout/recovery behavior, tested concurrency/resource conflicts and evidence stage.

D.5 T736 reference design boundary

The current documentary baseline describes a HelperT736C core module approximately 60 x 40 x 3.0 mm and a substantially larger reference carrier. The credit-card-sized footprint remains a design target, not a proven integrated product. A proposed minimum common carrier retains power, recovery/USB0, UART0 service access, one selected network route, powered USB Host and intentionally allocated serial expansion. Exact connectors, power rails, USB supply, interface multiplexing, thermal design and device-tree configuration remain configuration-specific engineering work.

D.6 Interface names are not electrical guarantees

UART logic level is not RS-232 and is not RS-485. A connector that physically resembles miniPCIe does not prove PCIe signaling. A camera connector does not prove driver, format, timing or calibration compatibility. The Profile must record the actual electrical and software contract.

Appendix E - Reference Engineering and Acceptance Roadmap

E.1 - 0. Problem and Work definition

Question answered: Demand, Authority, acceptance, baseline and responsible parties are explicit.

Minimum evidence: Versioned Work brief.

E.2 - 1. Documented

Question answered: Source documents support a possible capability.

Minimum evidence: Document hash, profile, conflicts/unknowns.

E.3 - 2. Simulated

Question answered: Contract works in a controlled simulator with no physical effect.

Minimum evidence: Test vectors, receipts and deterministic errors.

E.4 - 3. Remote simulated

Question answered: A remote client reaches the simulator through the intended route.

Minimum evidence: Route, client, protocol/runtime and remote receipts.

E.5 - 4. Physical reference

Question answered: A bound reference device performs limited real observations/actions under controlled conditions.

Minimum evidence: Physical Resource ID, BOM/OS/driver, test results and stop/recovery evidence.

E.6 - 5. Configuration freeze

Question answered: Hardware, image/runtime, adapters, contracts and test pack are fixed for shipment.

Minimum evidence: Freeze manifest and recovery procedure.

E.7 - 6. Physical delivered

Question answered: Delivered instance is bound and reproduces acceptance at destination.

Minimum evidence: Delivered identity, installation, local test and handover.

E.8 - 7. Operational

Question answered: The system completes real Work over a defined period with incident and value records.

Minimum evidence: Operational log, outcomes, cost/value, maintenance and incidents.

E.9 - 8. Independent reproduction

Question answered: Second team/location reproduces the agreed function.

Minimum evidence: Reproduction time, deviations, clarification count and comparative evidence.

E.10 Configuration freeze

A freeze is not merely a filename. It should bind at least hardware revision/BOM, image or OS, runtime, driver set, adapter versions, capability contracts, application version, calibration, test pack, known limitations, recovery method and evidence hashes.

E.11 Shipping rule

Ship after the reference configuration has passed the intended pre-shipment acceptance. Shipping does not turn the reference identity into the delivered identity. The delivered unit receives its own Resource identity and local binding record.

E.12 Local operation

Where the use case requires local continuity, basic operation and safe stop/recovery should not depend on a permanently open cloud session. Cloud services can coordinate, update and collect evidence without becoming an invisible single point of control for every real-time physical behavior.

E.13 Second-site reproduction

The second site returns not only pass/fail but also time to reproduce, substitutions, environmental differences, hidden assumptions, questions asked, failed attempts and final evidence. These results are inputs to the next shared Profile.

Appendix F - Evidence Register and Source Notes

F.1 Internal evidence register

H0-001 / 2026-11-11 - Cross-Project Hand Off

Evidence class: Documented handoff

Boundary: Consolidates prior decisions/evidence; did not re-probe physical T736 or deploy firmware/device.

C0-005 / 2026-11-09 - T736 hardware standardization baseline

Evidence class: Documented

Boundary: Supplier-document review, pin/interface mapping and conflicts; no new bench measurement.

C0-006 / 2026-11-09 - Remote engineering plan

Evidence class: Planned architecture

Boundary: Cloud simulation -> remote reference board -> freeze -> ship -> local bind/accept.

C0-007 / 2026-11-10 - Source-verified simulator

Evidence class: Simulated

Boundary: Simulator source and receipts; physical targets return unbound; does not prove hardware.

C0-008 / 2026-11-10 - Recorded Azure remote simulation

Evidence class: Remote simulated

Boundary: 14/14 recorded remote simulation checks; real 736 was unreachable in that record.

W0-003 / P1 - Invitation Card + One-Page Brief

Evidence class: Published digital

Boundary: Digital publication/QR checks completed; physical print/NFC read-back deferred.

W0-004 / P2 - Executive Summary

Evidence class: Published digital

Boundary: English summary HTML/PDF published and verified.

W0-005 / P3-A - Part I

Evidence class: Published digital

Boundary: English Part I published and verified.

W0-006 / P3-B - Part II

Evidence class: Published digital / current known public baseline

Boundary: English Part II published and verified.

W0-007 / P3-C - Part III

Evidence class: Generated and archived

Boundary: Digital publication package prepared; remote deployment deferred.

W0-008 / P3-D - Part IV

Evidence class: Generated and archived

Boundary: Digital publication package prepared; remote deployment deferred.

W0-009 / P3-E - Part V

Evidence class: Generated and archived

Boundary: Digital publication package prepared; remote deployment deferred.

W0-010 / P3-F - Part VI

Evidence class: Generated and archived

Boundary: Digital publication package prepared; remote deployment deferred.

W0-011 / P4 - English Master Consolidation

Evidence class: Generated master

Boundary: Consolidates Parts I-VI plus front matter, Appendices A-H, glossary, source and version registers.

F.2 Evidence labels

Proposed - a design, architecture, pathway or claim to be tested.

Documented - supported by source documentation, not necessarily by the physical instance.

Simulated - tested in a simulator or controlled synthetic environment.

Remote simulated - tested through the intended remote route, still without physical target acceptance.

Physical reference - verified on a specifically identified reference physical instance.

Physical delivered - verified on the delivered instance at the destination.

Operational - observed during real Work over a defined period.

Independently reproduced - reproduced by a second team or location using the published boundary.

F.3 Observation versus receipt

An operation receipt is evidence about a specific operation. A fresh observation is evidence about a current or recent state. The evidence object should carry observation time, receipt time, source, method, version and freshness limit so that a historical receipt is not silently treated as current state.

F.4 Source hierarchy

When evidence conflicts, preserve the conflict rather than silently reconciling it. A practical hierarchy is current physical measurement for the bound instance; current manufacturer/BOM-specific data; versioned engineering documentation; simulator behavior; design proposal. Different sources may answer different questions and therefore should not be forced into a single value without explanation.

Appendix G - Participation Guide

A participant can begin with one real task or one real capability. A complete platform is not required.

G.1 Five ways to join

Bring a real task - a result that someone genuinely needs and can accept.

Bring a local capability - material, machine, space, person, storage, transport or service.

Bring an engineering or model function - software, model, design, simulation, adapter, translation or analysis.

Buy or sponsor a useful result - fund a Work with a defined acceptance condition.

Reproduce and review - independently reproduce, test, challenge or improve a reference result.

G.2 Minimum Capability Note

A Capability Note can be written without an API:

Capability: What can you do?
Location: Where can it happen?
Availability: When or how much is available?
Responsible party: Who can accept the Work?
Conditions: What inputs, tools or limits apply?
Evidence: What will prove completion?
Cost/time: What does one accepted result require?
Rights: What use, reuse, confidentiality or attribution conditions apply?

G.3 First Work

The first Work should be deliberately small. It should make one interface real and produce evidence strong enough for another team to understand the result. Avoid selecting the most dramatic autonomous use case merely because it makes the best demonstration.

G.4 30 / 60 / 90 day pattern

Days 1-30: define one Work, baseline, Authority, contract and first accepted result.
Days 31-60: reproduce with a second team, device, maker or location; record differences and hidden assumptions.
Days 61-90: decide whether to scale, standardize, change the interface, keep the capability local, or stop.

A network that can stop weak ideas is more credible than one that converts every pilot into a success story.

G.5 Contribution and value

Contribution can be paid service, sponsored work, in-kind equipment, research contribution or another explicitly agreed arrangement. Unpaid contribution should be visible rather than disguised as free capacity. IP, reuse, attribution and confidentiality follow the specific Work and contribution agreement.

Appendix H - Multilingual Editorial Rules

The English W0-011 master is the translation source for the next language layer.

H.1 Stable structure - Keep Part, Chapter, section and appendix numbering aligned with the English master. - Keep stable IDs, Runtime identities, Card IDs, Resource IDs, file hashes and version numbers unchanged. - Translate descriptive labels; do not translate identifiers that are used for technical addressing.

H.2 Evidence words are controlled vocabulary

The evidence labels **Proposed, Documented, Simulated, Remote simulated, Physical reference, Physical delivered, Operational, Independently reproduced** should use a maintained glossary in each language. Translation must not upgrade or weaken evidence status.

H.3 Proper names and standards

Preserve official organization names in English on first occurrence when local-language usage could create ambiguity. ISO, MCP, OpenAI, SDG numbers, product/model names and legal entity names should be verified against authoritative local usage before publication.

H.4 Numbers and units

Numbers, dates, units, tolerances, formulas and version identifiers must remain equivalent to the English master. Localize punctuation and typography only where it does not change technical meaning.

H.5 Chinese

The UN official language set includes Chinese. Simplified Chinese can serve as the first Chinese translation for the mainland engineering and manufacturing audience; Traditional Chinese can be maintained as a regional written variant for Hong Kong and other markets. Both should map to the same English master revision.

H.6 Arabic and RTL

Arabic requires right-to-left layout testing. URLs, code, formulas, model numbers and identifiers remain left-to-right where appropriate. Do a short RTL proof before translating the entire book.

H.7 Source citations

A translated edition should cite the same underlying source, not a machine-translated copy of the source, unless an official translated source exists. The translation may add an official local-language source link while preserving the original reference identity.

H.8 Change control

When the English master changes, create a translation delta list: changed section IDs, factual changes, evidence-status changes, new/removed sources, and terminology changes. Do not silently patch six languages independently.

Glossary

Authority

Business-level permission and responsibility for what a person or organization may discover, read, contribute, reserve, configure, execute, stop, accept, release or administer.

Capability

A bounded function that can produce a defined result under declared conditions.

Capability Note

A lightweight human-readable description of what can be provided, where, when, by whom, under what conditions and with what evidence.

Capability Profile

A versioned, more formal description of a Capability including inputs, outputs, interfaces, conditions and evidence.

Common Capability Contract

The shared technical/business boundary used by different implementations to deliver the same function.

Conformance

Evidence that an implementation satisfies a defined Profile or specification under the stated test conditions.

Custody

Responsibility for possession/handling of a physical object during a defined interval or handover.

Evidence

Versioned record of what was observed, tested, delivered, accepted, rejected or recovered.

Evidence boundary

The explicit limit on what a piece of evidence proves and what it does not prove.

Fresh observation

A state observation whose age and source satisfy the Work’s freshness requirement.

HSIMC Resource

A real-world or business object such as person, place, machine, material, inventory position, service or capability.

MCP resource

In Model Context Protocol terminology, data/context exposed by a server for a client to read; distinct from an HSIMC Resource object.

MCP tool

A callable operation exposed by an MCP server.

M-Person

HSIMC participation concept for a person represented with explicit identity/relationships and Work-specific Authority; not a claim that a digital representation replaces the person.

M-Team

A Work-specific combination of people, models, software tools and local capabilities with explicit responsibility.

Operation

A request or action with a stable identity, especially important when physical side effects are possible.

Operation receipt

Evidence about a specific operation; not automatically evidence of current state.

Physical delivered

Evidence class for a delivered physical instance that has been bound and accepted at destination.

Physical reference

Evidence class for a specifically identified physical reference instance tested under controlled conditions.

Profile

A versioned description of capability/conditions; it does not by itself prove that every instance currently satisfies the profile.

Reference result

A result and evidence pack used as a reproducible basis for comparison; not a standard by itself.

Resource

Generic word used contextually; the document distinguishes HSIMC Resource from MCP resource.

Runtime

A deployed execution identity/configuration used by HSIMC projects; a Runtime identity is not by itself Authority to control a device.

Stateless protocol

Protocol behavior in which each request carries the context required for processing rather than relying on hidden transport-session state.

Work

A specific accepted unit connecting demand, responsible parties, Resource, conditions, Authority, execution, acceptance and value evidence.

Work contract

The versioned agreement/record that defines a Work’s inputs, outputs, conditions, roles, acceptance and recovery.

Bibliography and Source Register

The source register identifies external material used for factual context. HSIMC architectures, Work structures, capability contracts, participation models and proposed application journeys are HSIMC proposals unless a source explicitly describes an existing external implementation.

S01 - OpenAI: Accelerating engineering for the physical world

Date/status: 2026-11-03
URL: https://openai.com/index/accelerating-engineering-for-the-physical-world/
Role in this white paper: Engineering/physical-world context and collaboration invitation.

S02 - Model Context Protocol: The 2026-07-28 Specification

Date/status: 2026-07-28
URL: https://blog.modelcontextprotocol.io/posts/2026-07-28/
Role in this white paper: Stateless core, routing, MRTR, extensions, authorization hardening and deprecation context.

S03 - Model Context Protocol: 2026-07-28 Specification

Date/status: 2026-07-28
URL: https://modelcontextprotocol.io/specification/2026-07-28/
Role in this white paper: Normative specification family used for protocol terminology.

S04 - Model Context Protocol: 2026-07-28 Release Candidate

Date/status: 2026-05-21
URL: https://blog.modelcontextprotocol.io/posts/2026-07-28-release-candidate/
Role in this white paper: Migration/deployment context for the major protocol revision.

S05 - United Nations DESA: Transforming our world: the 2030 Agenda for Sustainable Development

Date/status: 2015
URL: https://sdgs.un.org/2030agenda
Role in this white paper: 17 SDGs, integrated/indivisible framing, LNOB pledge and Goal 17 language.

S06 - UN Sustainable Development Group: Leave No One Behind

Date/status: current guidance
URL: https://unsdg.un.org/2030-agenda/universal-values/leave-no-one-behind
Role in this white paper: Meaningful participation, evidence, disaggregation, transparency, privacy and accountability.

S07 - United Nations DESA: Goal 17

Date/status: current
URL: https://sdgs.un.org/goals/goal17
Role in this white paper: Partnership, science/technology cooperation, multi-stakeholder partnership and data targets.

S08 - United Nations DESA: Multi-stakeholder partnerships

Date/status: current
URL: https://sdgs.un.org/topics/multi-stakeholder-partnerships
Role in this white paper: Partnership framing around pooled knowledge, expertise, technology and resources.

S09 - ISO: ISO/IEC 7810 Identification cards - Physical characteristics

Date/status: standard reference
URL: https://www.iso.org/obp/ui/#iso:std:iso-iec:7810:ed-3:v1:en
Role in this white paper: ID-1 nominal 85.60 x 53.98 mm planar reference; full standard has wider physical requirements.

S10 - ISO: Developing standards

Date/status: current
URL: https://www.iso.org/developing-standards.html
Role in this white paper: Market need, expert/multi-stakeholder development and consensus process.

S11 - ISO/IEC: ISO/IEC Directives - consolidated

Date/status: current
URL: https://www.iso.org/sites/directives/current/consolidated/
Role in this white paper: Formal standard-development procedure and consensus framework.

S12 - European Commission: Construction and demolition waste

Date/status: current
URL: https://environment.ec.europa.eu/topics/waste-and-recycling/construction-and-demolition-waste_en
Role in this white paper: Composition, scale, resource value, sorting and hazardous-material context.

S13 - University of Manchester research: e-Textile Embroidered Wearable Near Field Communication RFID Antennas

Date/status: research paper
URL: https://research.manchester.ac.uk/files/76660584/e_Textile_Embroidered_Wearable_Near_Field_Communication_RFID_Antennas.pdf
Role in this white paper: Conductive-thread embroidery, fabrication techniques, resistance and NFC matching.

S14 - MDPI: Wearable/embroidered antenna research cited in Part II

Date/status: research paper
URL: https://www.mdpi.com/2224-2708/4/4/371
Role in this white paper: Example of decorative/embroidered textile RF functionality.

S15 - MDPI: Printed electronics manufacturing review cited in Part II

Date/status: research review
URL: https://www.mdpi.com/2504-4494/5/3/89
Role in this white paper: Ink, substrate, processing and repeatability context for printed electronics.

S16 - MDPI: Functional/printed materials research cited in Part II

Date/status: research paper
URL: https://www.mdpi.com/2673-3978/6/2/7
Role in this white paper: Additional materials/process context used in Part II.

S17 - GHG Protocol: Product Standard

Date/status: current
URL: https://ghgprotocol.org/product-standard
Role in this white paper: Product life-cycle emissions across raw materials, manufacturing, transport, storage, use and disposal.

S18 - GHG Protocol: Inventory and Project Accounting: A Comparative Review

Date/status: current
URL: https://ghgprotocol.org/blog/inventory-and-project-accounting
Role in this white paper: Distinction between organizational inventory and project/intervention accounting.

S19 - GHG Protocol: Estimating and Reporting Avoided Emissions

Date/status: current
URL: https://ghgprotocol.org/estimating-and-reporting-avoided-emissions
Role in this white paper: Cautious disclosure of comparative/avoided-emissions claims.

S20 - IATA: 2026 Global Passenger Survey

Date/status: 2026-10-28
URL: https://www.iata.org/en/pressroom/2026-releases/10-28-iata-2026-global-passenger-survey-travelers-biometrics-ai-entry-requirements/
Role in this white paper: 49% global reported use of AI for travel planning and survey scope.

S21 - Hong Kong International Airport: Baggage Pick-up and Delivery

Date/status: current
URL: https://www.hongkongairport.com/en/passenger-guide/airport-facilities-services/baggage-pick-up-and-delivery
Role in this white paper: Existing paid airport-to-hotel/home/downtown baggage delivery context.

S22 - Hong Kong International Airport: MyTAG

Date/status: current
URL: https://www.hongkongairport.com/iwov-resources/html/MyTag/en/mytagseries.html
Role in this white paper: RFID/QR-linked baggage notification context used in Part V.

S23 - Hong Kong International Airport: About Us / smart airport history

Date/status: current
URL: https://www.hongkongairport.com/en/about-us/?section=our-history
Role in this white paper: Airport technology/trolley-monitoring context referenced in Part V.

S24 - Model Context Protocol TypeScript SDK: Support/migration for 2026-07-28

Date/status: current
URL: https://ts.sdk.modelcontextprotocol.io/v2/migration/support-2026-07-28
Role in this white paper: SDK migration/support context cited in Part III.

S25 - ISO: Glossary

Date/status: current
URL: https://www.iso.org/glossary.html
Role in this white paper: Standards vocabulary context cited in Part VI.

Internal project source set

The English master also relies on HSIMC internal project records, especially H0-001 and the C0-005 through C0-008 evidence chain, plus publication versions W0-001 through W0-010. These are recorded in Appendix F and the Version and Evidence Register.

Version and Evidence Register

Publication versions

W0-001 - Editorial master / TOC

Date: 11 Nov 2026
Status: Planning
Scope: English architecture for six Parts, 18 Chapters and Appendices.

W0-002 - Publication production plan

Date: 11 Nov 2026
Status: Planning
Scope: Four reading depths, HTML/PDF, QR/NFC and language sequence.

W0-003 - P1 - Invitation Card + One-Page Brief

Date: 11 Nov 2026
Status: Published digital
Scope: Card and brief; physical printing/NFC read-back deferred.

W0-004 - P2 - Executive Summary

Date: 11 Nov 2026
Status: Published digital
Scope: English Executive Summary HTML/PDF.

W0-005 - P3-A - Part I

Date: 11 Nov 2026
Status: Published digital
Scope: Chapters 1-3.

W0-006 - P3-B - Part II

Date: 11 Nov 2026
Status: Published digital; current known public baseline
Scope: Chapters 4-6.

W0-007 - P3-C - Part III

Date: 11 Nov 2026
Status: Generated/archived; deployment deferred
Scope: Chapters 7-9.

W0-008 - P3-D - Part IV

Date: 11 Nov 2026
Status: Generated/archived; deployment deferred
Scope: Chapters 10-12.

W0-009 - P3-E - Part V

Date: 12 Nov 2026
Status: Generated/archived; deployment deferred
Scope: Chapters 13-15.

W0-010 - P3-F - Part VI

Date: 12 Nov 2026
Status: Generated/archived; deployment deferred
Scope: Chapters 16-18.

W0-011 - P4 - English Master Consolidation

Date: 12 Nov 2026
Status: Master generated; deployment deferred
Scope: Front matter, Parts I-VI, Appendices A-H, glossary, bibliography and registers.

Current engineering evidence boundary carried into W0-011

C0-005: documentary hardware baseline; no new bench measurement implied.

C0-007: source-verified simulator; physical targets remain unbound unless separately verified.

C0-008: recorded 14/14 Azure remote simulation checks from 10 November 2026; not a physical-board acceptance record.

H0-001: cross-project handoff; did not re-probe the physical T736 and did not convert a design target into a delivered product.

Credit-card-sized hardware: remains a design target; the standard ID-1 planar dimensions are used as a shared starting reference.

Deployment status

As of this P4 consolidation, the known public white-paper site baseline is W0-006 (Parts I-II). W0-007 through W0-010 were generated and archived with remote deployment intentionally deferred. W0-011 is prepared so that a later controlled deployment can move from the verified W0-006 public baseline to the consolidated English master after a unified review.

Translation baseline

Future United Nations six-language work should use W0-011 as the English source revision unless a later English master is explicitly frozen. Translation status should be maintained independently from deployment status.

Colophon

Title: From Stone and Thread to Shared Intelligence
Subtitle: A Credit-Card-Sized Starting Point for Local Making and Global Participation
Author: Benny (Yan) Liao - Founder, HSIMC
English master: W0-011
Date: 12 November 2026
Publisher/initiative: HSIMC
Working principles: Human Authority · Local Capability · Global Participation
Mission statement: No One Left Behind in the GenAI Physical Era

This master is designed for digital HTML, A4 PDF, later book layout and multilingual derivative editions. Physical printing and NFC writing/read-back are production steps separate from digital publication.