

CHAPTER 1 / HSIMC BLUEPRINT

The Ten-Billion-Phone Opportunity

From a mobilizing ambition to verified, useful capability

PROPOSITION

Ten-billion-phone activation is a strategic invitation and a research hypothesis, not an audited count of idle, safe, accessible or deployable handsets. The result that matters is a useful Work accepted by a real person—not a device merely counted or connected.

1.1 Consumer obsolescence versus remaining productive capability

A phone may be retired from daily personal use while retaining a camera, screen, microphone, sensors, local storage, networking or accessible user interface. Consumer replacement and technical uselessness are different events. A device need not host a large language model to contribute: it may capture evidence, present a consent decision, act as a supervised terminal or connect a qualified local capability to an authorized Work.

Residual functionality is not sufficient evidence of readiness. Battery condition, electrical safety, operating-system support, network isolation, retained personal data, permission controls and accessibility must be tested for the proposed role. A device that cannot be operated safely or securely may belong on a repair or responsible recycling path instead. Reuse is an option to qualify, not an environmental or economic benefit to assume.

1.2 Dormant phones as distributed starting points—not a centrally owned inventory

Possible devices sit with households, companies, schools, repairers, charities and other independent holders. Ownership, possession and permission remain with those parties; HSIMC cannot treat them as a centrally controlled stockpile. A dormant phone is not automatically abandoned, available, legally transferable or safe to enroll. Participation must begin with an identifiable, willing holder and a narrowly defined purpose.

An honest activation record therefore starts with the holder's permission, a minimally necessary device description, an inspection result, an intended role and an explicit stop or revocation route. Entry may be through an ordinary URL or QR code; qualified NFC and local adapters are optional routes, not prerequisites. No universal factory reset, rooting, sideloading or administrator access is presumed.

1.3 Existing stock, future retirement and alternative AI adoption scenarios

Scale indicators establish relevance, not an actionable inventory. ITU estimates that 6.0 billion people used the Internet in 2025, while 2.2 billion remained offline [E1]. GSMA reports 4.8 billion mobile-Internet users at end-2025 [E2]. Those figures describe people under different definitions—not physical phones available for a new Work. ITU and UNITAR report 62 million tonnes of electronic waste generated in 2022 across many equipment categories, with 22.3% formally documented as collected and recycled [E3]. That mass cannot be converted into a stock of reusable handsets.

The initiative's ten-billion label should consequently be tested as a long-horizon participation thesis, never displayed as ten billion presently eligible nodes. Stock and flow must be reported separately: phones already retained, future devices leaving primary use, devices lost or recycled, and repeat registrations. Low, reference and high scenarios will require observed cohorts, deduplication rules, transparent eligibility assumptions and dated revision—not extrapolation from subscription counts.

1.4 From estimated devices to eligible nodes, active contributors and accepted work

The measurement chain is explicit: **observed candidate** → **holder-authorized device** → **safely operable device** → **reachable interface** → **qualified capability** → **completed action** → **human-accepted Work**. Each transition has its own numerator, denominator, evidence and rejection reason. A phone with Wi-Fi is not therefore an MCP server; a successful tool call is not therefore a delivered service; a completed action is not necessarily an accepted outcome.

HSIMC should measure time to first useful evidence, independent-operator success, supervision effort, repeat acceptance, total cost and failure or withdrawal rates. Record unique devices separately from operators, work requests and accepted outcomes. A small cohort with reproducible value is stronger evidence than a large map of unverified icons. Results should allow an independent reviewer to follow the work from Intent and Authority through Reality, Action, Evidence, Acceptance and Continue.

For a comparable pilot, the ledger should distinguish a device record from an operator, a requested Work and a signed acceptance decision. Keep personal information and device identifiers confidential by default, retaining only what a legitimate, bounded Work requires. Publish aggregated results only where permitted. A verified Work means its evidence can be inspected or independently judged; it does not mean that the phone is perpetually online or available.

1.5 The phone-first invitation and the wider existing-asset opportunity

The first invitation can be concrete: bring an existing phone, inspect it without unnecessary data access, choose one bounded task, disclose required permissions, and produce a result another person can verify. For example, an authorized operator might use an eligible phone to record a product's visible condition and submit a dated evidence item for human review. Acceptance, rejection, deletion and withdrawal must be possible; this example is a proposed Work, not a claim of field deployment.

The phone is the entrance, not the boundary. Qualified charging and data cables, NFC-enabled physical identifiers, fixtures, sensors, local compute and human services may supply conditions that the handset cannot meet alone. Some cases will need no additional product; some should not proceed. The infrastructure opportunity is to discover and complete the *smallest missing condition* for a valuable, authorized Work, and to verify the outcome locally.

The first template also needs a genuine rejection gate: unsafe power, unsupported software, unclear evidence or a holder who withdraws consent should stop the attempt and preserve only permitted failure evidence. Correct refusal is part of a trustworthy service. This method should be replicable in different languages, premises and regulatory settings without transferring device ownership, personal data or corporate control to HSIMC.

WORK / DECISION TEST

Chapter 1 advances only when a real pilot reports distinct counts for candidates, authorized devices, safety-qualified devices, usable capabilities and independently accepted Works. If cost, risk or missing permissions make the pathway inferior to alternatives, the correct outcome is to revise or reject it—not to enlarge the phone count.

PUBLIC EVIDENCE · SOURCE LINKS

[E1] [International Telecommunication Union \(ITU\) · Measuring digital development: Facts and Figures 2025](#) (2025)

[E2] [GSMA · The State of Mobile Internet Connectivity 2026](#) (2026)

[E3] [ITU / UNITAR · The Global E-waste Monitor 2024 \(UNITAR release\)](#) (2024)

STATUS / BOUNDARY. Author-origin infrastructure proposal informed by public secondary indicators; no evidence of ten billion eligible handsets, no third-party endorsement, no deployment authorization.

NEXT GATE. After English technical/editorial acceptance, complete Part I (Chapters 1–3), freeze source revision, and run six-language Part-level QA; do not translate an unfrozen draft into separate authoritative versions.

CHAPTER 2 / HSIMC BLUEPRINT

From App-Centric Use to Intent-Driven Physical Tools

Preserve the human interface; make a bounded real-world outcome computable

PROPOSITION

A useful GenAI Physical Tool is not defined by controlling a screen. It is a declared capability that can help realize a specific human intent in a known physical Reality, with accountable authorization, evidence and acceptance.

2.1 The shift from operating an app to requesting a real-world outcome

The app-centered model asks a person to find an application, navigate menus and interpret results. An intent-driven Work begins instead with a human outcome: document a sample, verify a condition, ask for accessible guidance, or obtain a locally usable component. The intent is not itself an instruction to act. It becomes a Work only when the intended result, permitted scope, responsible parties, environment and acceptance criteria are explicit.

For example, “identify this connector” can be decomposed into a holder-approved image, a qualified identification tool, evidence of uncertainty, and a decision a technician can confirm. A mistaken identification must be rejectable. The infrastructure preserves the relation between requested outcome and supporting evidence; a polished answer alone does not complete the Work.

2.2 Keeping human interfaces while making bounded functions available to AI

People still need displays, accessible controls, multilingual explanations, consent and a visible stop button. Intent-driven design does not eliminate apps or replace the human user. A web page or local application may remain the host that asks permission, explains consequences and presents the completed result; a model-assisted client selects only the qualified tools exposed by that host or an independent service.

HSIMC separates interface, capability and authority. Reading a camera preview is different from changing system settings, deleting records or starting a purchase. Sensitive actions require explicit, appropriately scoped authorization and potentially a second human check. The availability of an automation API or accessibility service is not permission to automate an entire phone.

2.3 A phone as a sensor, interface, computing resource and physical gateway

An existing phone may be an observation point, a work screen, a connectivity bridge, limited local compute or an input to a nearby qualified device. These roles are alternatives, not promises about every handset. Each role has distinct conditions: operating-system version and patches, battery and power health, transport availability, camera/sensor performance, permissions, storage protection and operator accessibility.

Where the phone lacks one condition, a cable, fixture, NFC label, edge device or supervised work site may supply it. Such parts are not automatically intelligent tools; a passive cable can be a necessary material dependency without claiming to be an MCP server. A Work graph should identify what each component contributes and keep procurement, physical inspection and service delivery as separate accountable steps.

2.4 Why useful GenAI participation need not mean running a large model on the phone

A local device can contribute structured observations or invoke a permitted remote model without hosting a large model itself. This reduces the need to discard hardware solely because its processor cannot execute the newest AI model. It does not remove costs: online inference introduces latency, transmission, account terms, security, energy and availability dependencies. Local-only, hybrid and remote-assisted routes must be evaluated against the same Work acceptance criteria.

If connectivity fails, the candidate Work may safely stop, queue non-sensitive evidence with consent, or use a tested offline path. It must not silently transmit private images or retry a consequential action indefinitely. The best route is determined by the purpose, data sensitivity, local capabilities and the human time window—not by placing “AI” on the device label.

2.5 Beyond App Control: from operating applications to authorizing capabilities

Screen-driving shortcuts, debug bridges and privileged operating-system APIs can be useful in authorized laboratories, but they are poor universal service contracts. Their permissions and behaviors change across operating systems and devices, and they can confer excessive agency. OWASP identifies excessive functionality, permission and autonomy as distinct paths to harmful outcomes [E4]. Prefer narrow, typed capabilities with predictable inputs, bounded side effects and visible records.

A capability contract should state who may request it, what information is exposed, when human confirmation is required, and how the result is measured or revoked. Execution still belongs to a specific supervised Work and its actual Reality. This is the design transition: from “the AI can touch an app” to “an authorized person can obtain and inspect a useful outcome through a declared capability.”

WORK / DECISION TEST

A candidate must demonstrate one end-to-end Work in which a human states an outcome, grants bounded access, can refuse an unsafe invocation, and accepts or rejects evidence. Show the host UI, invoked capability, actual device role and revocation path; do not count screen taps as value.

PUBLIC EVIDENCE · SOURCE LINKS

[E4] OWASP GenAI Security Project · LLM06:2025 Excessive Agency (2025)

STATUS / BOUNDARY. HSIMC architecture proposal. No universal device permissions, third-party application endorsement, mass deployment or accepted Work is implied.

NEXT GATE. Freeze the documented Work contract and operator-visible consent test with Part I, then evaluate one local object-to-evidence workflow under fresh authorization.

CHAPTER 3 / HSIMC BLUEPRINT

MCP as an Enabler—not the Complete Production System

Protocol interoperability is distinct from operational readiness and accepted Work

PROPOSITION

MCP can make tools and contextual resources discoverable and callable across implementations. It does not by itself establish the right to act, the physical conditions needed for safe execution, or proof that a person accepted the resulting Work.

3.1 The July 28, 2026 protocol shift and its implications for distributed capabilities

The published MCP 2026-07-28 revision retires the protocol-level initialize handshake and the `Mcp-Session-Id` session header. Each request carries version and capability metadata; server discovery is available when needed, and standard HTTP routing becomes simpler. This is a concrete protocol change, not a claim that all deployed MCP clients, legacy servers or phones already implement the new revision [E5].

A physical capability can be offered by an authorized server without treating a persistent transport session as its identity. Yet the device may still need calibration, a human operator, electricity, legal permissions and durable evidence. HSIMC uses MCP as one possible interoperability layer inside a larger Reality-and-Work system; compatibility and deployment tests must name the actual protocol revision.

3.2 Discovery, invocation and interoperability across independent implementations

MCP distinguishes callable tools from resources that provide context. A host can present appropriately available functions to a model-assisted client, which requests a narrowly typed tool call. An independent provider may advertise equivalent outcomes through a different implementation. A declared schema helps interoperability but does not verify a supplier, guarantee an input is truthful or authorize a physical act [E6].

A useful capability record should identify the tool and version, provider, expected inputs, output and error states, limits, permissions, environment and evidence format. Discovery, successful invocation and safe operation are distinct gates. Clients should validate returned data and treat untrusted tool descriptions, annotations and retrieved text as data—not a source of new authority. Human approval remains available for consequential operations.

3.3 Stateless protocol interactions, persistent work identities and physical state

Stateless communication means the protocol does not require one hidden server-side conversation session across requests. It does not mean the application has no memory or that a physical task has no state. A Work can persist through an explicit Work ID, signed or otherwise verifiable evidence, owner-controlled durable storage and a declared continuation handle, rather than assuming that a network connection equals authorization.

Physical state is independently real: a battery charges, a fixture moves, a package is present or a person withdraws permission. The next invocation must reconcile that Reality before acting. A lost response should not automatically produce a second irreversible action. Idempotency, retries, stop conditions, independent readback and operator override belong in the Work contract, whether or not the transport is stateful.

3.4 Protocol conformance, device readiness and accepted delivery as separate milestones

Three tests are required. **Protocol conformance** asks whether the client and server actually exchange valid messages at a named version. **Reality readiness** asks whether this specific device, place, input, adapter, safety condition and authorization can support the proposed task now. **Accepted delivery** asks whether evidence of the resulting outcome meets a named human or organizational acceptance rule.

Passing one gate is never shorthand for all three. A Wi-Fi connection or a successful tools/list call does not prove camera permission, trusted hardware power or valid commercial rights. Conversely, a manual demonstration may create real evidence without being a compliant MCP product. Failed checks and safe refusals should be recorded as useful findings instead of hidden to inflate node counts.

3.5 Date-based protocol conformance and the Sony pilot's older baseline

The HSIMC Sony reference journey is a historical research case, not automatic proof of conformance with a later protocol. Existing experiments must retain their own date, hardware, transport, software revision, test scope and original receipt. If an older prototype used a prior MCP handshake or an app-specific adapter, that evidence must remain correctly labeled. The 2026-07-28 release includes breaking changes, so migration is a testable new Work—not a retroactive relabeling [E5].

A new pilot should report the negotiated or declared protocol revision, authenticated identity, minimal tool set, consent decision, observable action, evidence digest, failure behavior and independent acceptance. Only those tested facts can support “MCP-ready” for that configuration. The future work is to make different implementations cooperate under one inspectable contract, not to announce universal compatibility by terminology alone.

WORK / DECISION TEST

On one explicitly authorized test configuration, separately demonstrate a named MCP revision exchange, bounded and safe physical readiness, and a completed Work reviewed by a second person. Reject partial success as proof of production deployment.

PUBLIC EVIDENCE · SOURCE LINKS

[E5] [Model Context Protocol Core Maintainers · The 2026-07-28 Specification](#) (2026)

[E6] [Model Context Protocol specification · 2026-07-28 Tools](#) (2026)

STATUS / BOUNDARY. Protocol claims are dated to the published 2026-07-28 MCP revision; local Sony or other devices are not asserted to have passed that revision. No partner endorsement or permission escalation is implied.

NEXT GATE. Use the Chapter 3 three-gate distinction when writing Chapters 11, 16–18 and 30; prepare a test-specific Work acceptance record before live execution.