HUSSH TECHNOLOGIES CORPORATION

🤫 One Puppy

Partner Master Pack

Everything the product team hands to partners: the offer, the partnership, the platforms, and the company and technology behind Agentic Personal Sovereign Supercomputing. One document, current as of the date below.

Rio Grande — the original Puppy One.

Built by hushh Technologies on Apple hardware, extensible to the broader ecosystem, with Google Cloud as the governed surge tier.

CONFIDENTIAL — partner distribution. Final-copy candidates; external use after McDermott sign-off. Revision 2026-06-12.

Reserve any machine for $0.69 at hushh.ai/one · partners@hushh.ai

WHAT'S INSIDE

Contents

How to use this pack: Sections 01–03 are customer- and partner-facing and may be shown or excerpted directly. Section 04 is for technically deep partners and recruiters. The full eight-vendor kit set, unit economics, and the GTM/partner operations kits ship as separate files alongside this master pack.

01

The Offer

Brochure, lineup, and the field leave-behind — what a partner shows a customer.

🤫 One Puppy

Own your intelligence.

Built on the Apple hardware you own — a fully on-device, always-on personal data agent that works with Siri and your other AIs, fetching your information back from every institution that banks it against your phone number. You are the owner of record, 24/7/365.

Your data, your business. · 24/7/365, four seasons style.

POCKET

🤫 One Puppy Max

Built on iPhone 17 Pro Max — the phone becomes the agent that holds your data for you.

$3,669.69

Device · Day 0 setup · 12 months of service

DESK

🤫 One Puppy Max Pro

Built on 16″ MacBook Pro M5 Max — extreme personal computing you own outright.

$8,469.69

Device · Day 0 setup · 12 months of service

Reserve any machine for $0.69 at hushh.ai · Your data, your business.

Start at $69.69/mo on the device you own · Your agents. Yours to own.

Day 0, not Day 30.

Every software pilot dies in setup. hushh One Puppy skips it — configured before the box opens.

Fetch

With your consent, the Puppy collects your picture from every information banker — accounts, policies, statements, inboxes.

Organize

On-device models assemble one living view. Nothing leaves the hardware you own.

Guard

Every access is consent-gated and written to a transparency log you can read.

Supercomputer. Singular owner.

18-core
CPU · M5 Max
Fusion Architecture

128GB
Unified memory for
on-device models

40-core
GPU · Neural Accel.
in every core

Day 0
Provisioned before
unboxing · Wi-Fi 7

hushh One Puppy Max Pro reference build. Hardware manufactured and sold by Apple Inc. Specs per Apple, March 2026. hussh configures and services the agent layer.

60 seconds

From inbox to full financial picture.

Connect one account. Watch the whole picture assemble — balances, policies, renewals, the dollars hiding in plain sight — with your consent, on your device.

On-device, on principle

Your data lives on hardware you own. Nothing leaves without explicit consent.

You are the owner of record

Every enterprise banking your data against your phone number answers to you.

Loyal to exactly one party

No ads, no data resale, no lock-in. Only you hold the bowl.

Reservations: $0.69 deposit, fully credited toward purchase and refundable anytime; a place in line, not a contract of sale. Specifications, pricing and availability confirmed at purchase.
Apple, iPhone, iPad, Mac, MacBook Pro; NVIDIA, DGX; Google, Pixel, Android; Samsung, Galaxy; Microsoft, Windows are trademarks of their respective owners, used nominatively.
Hushh Technologies Corporation (brand: “hussh”) is independent and not affiliated with, endorsed by, sponsored by, or partnered with any company named here. © 2026 Hushh Technologies Corporation.

🤫 One · The Puppy Family

One agent. Every door in.

Agentic Personal Sovereign Supercomputing on the Apple hardware you own.

ENTRY

🤫 One Puppy — $69.69/mo

Software only · your existing iPhone or Mac

Self-serve setup · on-device agent · consent fabric · the 60-second first fetch

POCKET

🤫 One Puppy Max — $3,669.69*

iPhone 17 Pro Max · arrives configured

White-glove Day 0 · 12 months managed service · dedicated device for client data

DESK

🤫 One Puppy Max Pro — $8,469.69*

16″ MacBook Pro M5 Max · 128GB-class

Extreme on-device AI · local models & memory · the personal supercomputer you own

PRACTICE

🤫 One Puppy Practice — $17,969.69*

3 seats: 1× Max Pro + 2× Max

Shared consent fabric · onsite onboarding · built for how an RIA practice actually buys

SERVICE

Self-serve $69.69/mo · Managed seat $169.69/mo (included in year 1 of every hardware bundle)

* Bundle pricing includes year-one managed service. Ultra and roadmap platforms: reservation only; specifications confirmed at purchase.

Reservations: $0.69 deposit, fully credited toward purchase and refundable anytime; a place in line, not a contract of sale. Specifications, pricing and availability confirmed at purchase.
Apple, iPhone, iPad, Mac, MacBook Pro; NVIDIA, DGX; Google, Pixel, Android; Samsung, Galaxy; Microsoft, Windows are trademarks of their respective owners, used nominatively.

FOR YOUR PRACTICE

Your client's whole picture.
In 60 seconds.

gent on the Apple hardware you own — works with Siri, fetches with consent, nothing in anyone's cloud. Agentic. Pe

1

Ask for the demo

We connect one account, live, in your office. Sixty seconds later you see the full financial picture — with a dollar figure attached.

2

Start where you are

hushh One Puppy runs on the iPhone or Mac you already own — $69.69/mo, self-serve, cancel anytime.

3

Step up when ready

Dedicated, pre-configured devices for the practice — arriving on Day 0 already working, client data off personal machines.

Three promises. No fine print.

On-device, on principle — your data lives on hardware you own, every access logged where you can read it.
You are the owner of record — every institution banking your data against your phone number answers to you.
Loyal to exactly one party — no ads, no data resale, no lock-in.

Book the 60-second demo · hushh.ai · hello@hushh.ai

Reservations: $0.69 deposit, fully credited toward purchase and refundable anytime; a place in line, not a contract of sale. Specifications, pricing and availability confirmed at purchase.
Apple, iPhone, iPad, Mac, MacBook Pro; NVIDIA, DGX; Google, Pixel, Android; Samsung, Galaxy; Microsoft, Windows are trademarks of their respective owners, used nominatively.

02

The Partnership

How to sell the category and the partner economics.

ONE PUPPY · CHANNEL & PARTNER PROGRAM

Sell the category, not a gadget.

Personal supercomputing — owned, on-device, consent-first. A new line for your existing book.

WHO WE PARTNER WITH

- **Apple-ecosystem MSPs & consultants**
  You already manage their Macs and iPhones. hushh One Puppy adds a managed agent layer — and a recurring line — to devices you already touch.

- **RIA platform & custodian ecosystems**
  Conference circuit, practice-management consultants, compliance vendors: hushh One Puppy makes your client's practice stickier.

- **Insurance FMOs / IMOs**
  Your agents drown in carrier portals. The household coverage picture in 60 seconds is a renewal-season weapon.

PARTNER ECONOMICS (DRAFT)

16.9% referral on year-1 bundle revenue · 10% of service revenue, months 13–36

Example: one Practice sale = $3,036.88 year-1 referral + service trail. Registered deals protected 90 days.

Co-marketing kit included: dark-mode collateral, the 60-second demo script, and the field leave-behind.

WHAT YOUR CLIENTS GET

The hushh One Puppy family: BYO software at $69.69/mo on devices they own → hushh One Puppy Max ($3,669.69, iPhone 17 Pro Max, arrives configured) → hushh One Puppy Max Pro ($8,469.69, 16″ MacBook Pro M5 Max) → hushh One Puppy Practice ($17,969.69, 3 seats + onsite onboarding). Day 0 setup, fully on-device agent, consent-gated data flows with a transparency log their compliance officer can actually read.

Partner onboarding: one call · one demo · one registered deal. partners@hushh.ai

Reservations: $0.69 deposit, fully credited toward purchase and refundable anytime; a place in line, not a contract of sale. Specifications, pricing and availability confirmed at purchase. Apple, iPhone, iPad, Mac, MacBook Pro; NVIDIA, DGX; Google, Pixel, Android; Samsung, Galaxy; Microsoft, Windows are trademarks of their respective owners, used nominatively.

03

The Platforms

One vendor lineup shown as the pattern; the full set ships separately.

🤫 ONE PUPPY ULTRA

The supercomputer is personal now.

NVIDIA-class personal supercomputing, with a supercomputing-powered AI agent — owned, governed, and controlled by exactly one party: you. Managed end to end by the hussh One Platform.

ULTRA S

🤫 One Puppy Ultra S

NVIDIA DGX Spark · GB10 Grace Blackwell

1 petaFLOP of AI compute · 128GB unified memory · runs up to 200-billion-parameter models on your desk — the same software stack as a datacenter.

Hardware from $3,999 · reserve for $0.69

ULTRA MAX

🤫 One Puppy Ultra Max

NVIDIA DGX Station · GB300 Grace Blackwell Ultra

~20 petaFLOPs · 748GB of coherent memory · trillion-parameter models beside your chair. The most powerful personal computer ever offered to an individual.

$85K-class hardware · reserve for $0.69

BUILT TO ORDER

🤫 One Puppy Ultra Custom

1–4× RTX PRO 6000 Blackwell · 96GB GDDR7 each

Your exact local-training rig, specified to the workload — configured, delivered, and installed by a human.

Configured to quote · reserve for $0.69

THE hussh ONE PLATFORM PROMISE

Every Ultra ships with the White-Glove Human-First install: a person delivers it, the consent ceremony sets every permission to your choice, the first fetch runs live in 60 seconds, and the install isn't done until you can drive it alone. Agent runtime for this platform is on our roadmap — your machine runs under hussh-managed infrastructure from day one, honestly labeled, with the full Puppy experience delivered the moment the port ships.

Reservations: $0.69 deposit, fully credited toward purchase and refundable anytime; a place in line, not a contract of sale. Specifications, pricing and availability confirmed at purchase.
NVIDIA, DGX, Grace and Blackwell are trademarks of NVIDIA Corporation, used to describe third-party hardware on which hussh software and services run.
Hushh Technologies Corporation (brand: “hussh”) is independent and not affiliated with, endorsed by, sponsored by, or partnered with any company named here. © 2026 Hushh Technologies Corporation.

04

The Company & The Technology

The story and the systems design — for partners who go deep.

🤫

HUSSH TECHNOLOGIES CORPORATION

Personal Supercomputing Infrastructure

A technical brochure on why we are building a fully on-device, consent-first personal data agent — and the infrastructure that lets humans and machines use personal information efficiently, safely, and on the owner's terms.

Rio Grande — the original Puppy One. Est. August 2021.

Agentic. Personal. Sovereign. · Your data, your business.

01 · WHERE THIS STARTED

A break-in, and a puppy.

Hushh Technologies Corporation began in August 2021, in the same week as two arrivals. The first was the company — founded after our founder's family experienced identity theft, and after watching how casually the institutions entrusted with personal information actually treated it. The founding conviction fit in five words: your data, your business.

The second arrival was a Bernedoodle puppy named Rio Grande. It took four years to see that the second explained the first. A dog is a creature of remarkable capability that is loyal to exactly one family. It fetches what you ask. It guards the house. It is always on. And it never, once, sells your secrets to a stranger. No one reads a privacy policy to trust a dog — the incentives are written on its face.

That became the design brief for an entire class of software: an agent with real capability and absolute, legible loyalty to a single owner. We called it hushh One Puppy. The company has spent the years since turning that metaphor into architecture.

THE PROBLEM WE KEPT HITTING

Every person is, in effect, a distributed database scattered across institutions that bank their information against a phone number — banks, brokerages, insurers, carriers, clinics, retailers, inboxes. The data describes you, but you can neither see it whole nor direct it. The owner of record has the least access of anyone. Existing AI tools made this worse, not better: they offered help in exchange for custody — upload everything to our cloud, trust our policy. That is the opposite of ownership.

02 · THE THESIS

Personal supercomputing infrastructure.

The hardware finally caught up with the idea. The most capable computers ever manufactured now fit in a pocket, on a wrist, on a desk. A modern phone or laptop carries enough local compute and memory to run capable models entirely on-device. For the first time, the supercomputer can belong to the person — not be rented from a cloud that reads what crosses it.

So we stopped thinking about “an app” and started thinking about infrastructure: the layer that lets a person's own machines hold, organize, and act on their own information — and lets trusted software and AI agents use that information efficiently, with consent, without ever taking custody of it. Three principles fixed the architecture:

On-device by default

Personal data lives on hardware the person owns. Models run locally. Nothing leaves without an explicit, logged consent event. The cloud is a tier you rent and govern, not a place your life is stored.

Consent is a protocol, not a checkbox

Every access produces a signed consent receipt and a line in a transparency log the owner can read. We call this PCHP — the permission and consent layer. Revocation is a first-class operation, not a support ticket.

The owner is sovereign

No ads, no resale, no lock-in. The agent is loyal to exactly one party and is legible about it. Capability without legible loyalty is surveillance; we refused to ship that.

03 · THE hussh ONE PLATFORM

How it actually works.

hussh One is the platform that binds a person's devices and governed cloud tenant into one fabric: one identity anchor, one consent fabric, one transparency log, one owner. The agent — hushh One Puppy — runs on top of it.

AGENT

hushh One Puppy

Fetches, organizes, and acts on the owner's information. Interoperates with Siri and other assistants via platform intents. Gives no advice it isn't asked for; the human decides.

INFRASTRUCTURE

hussh One Platform

Identity anchoring (the phone number as key), PCHP consent receipts, the transparency log, device/fleet lifecycle, and the on-device → governed-cloud burst boundary.

HARDWARE

Owned device + governed cloud

Apple silicon today (on-device inference, Secure Enclave). Burst to a customer-owned cloud tenant (confidential computing, customer-managed keys) only when a workload exceeds the device.

AGENT-NATIVE BY DESIGN

The platform exposes a Model Context Protocol (MCP) surface, so any MCP-capable agent — the owner's, or a trusted partner's — can read the catalog, place consent-gated actions, and check status. Agent-to-agent (A2A) flows and human-authorized payment mandates (AP2-class) mean even autonomous transactions carry a real human's authorization and a consent receipt. The same discipline that protects the human protects the machine acting for them.

04 · THE SILICON, AND THE HONEST LIMITS

Built on what's real.

We are precise about hardware because engineers judge precision. On-device inference runs on Apple silicon (M-series unified memory; A-series with the Secure Enclave as the identity root). When a workload exceeds the device, it bursts to a governed cloud tier — where independent analysis puts Google's Ironwood TPU at materially lower total cost than comparable GPU servers, and Arm-based instances deliver the cheapest honest substrate for always-on work. Owned core, governed surge.

We are equally precise about what is not yet true — because that is how trust is earned:

The 60-second fetch is the proof, and it is still being built.
The demonstration that defines the product — one connected account to a full picture in about a minute — is in active development. We will show it before we sell on it.

Agent runtimes are per-platform ports.
Apple is first and native. Android, Windows, and Linux (for NVIDIA-class personal supercomputers) are on the roadmap; until each ships, those devices run under managed infrastructure, labeled honestly.

Small devices are consent surfaces, not compute.
A smartwatch has no NPU for local LLM inference; it approves, revokes, and senses. Saying otherwise would be a spec lie, and we don't ship those.

Certifications are stated as in-progress.
FedRAMP High authorization is being pursued, never claimed as held. The honest status is itself the selling point to security-conscious buyers.

05 · WHERE THIS GOES

A computer that is finally yours.

The personal computer promised that the machine would work for you. Somewhere in the cloud era, that inverted — you began working for the machine, feeding data to systems that monetized it. Personal supercomputing infrastructure is the correction: capability that is local, loyal, legible, and owned.

We build it the way Rio earns trust — by being useful every day and never, once, betraying the household. Every device certified to that standard. Every access written where you can read it. Every agent loyal to exactly one party.

That is the company. That is the platform. And that is what we are asking the best engineers in the world to come build with us.

Build with us · partner with us · reserve a Puppy

engineering@hushh.ai · partners@hushh.ai · reserve for $0.69 at hushh.ai/one

Built by hussh Technologies on Apple hardware, extensible to the broader ecosystem, with Google Cloud as the governed surge tier.

Apple, iPhone, Mac, MacBook Pro, Siri (Apple Inc.); NVIDIA, DGX (NVIDIA Corp.); Google, Pixel, Android, Tensor, Titan, Ironwood, Axion (Google LLC) are trademarks of their owners, used nominatively. ushh Technologies Corporation (“hussh”) is independent and not affiliated with, endorsed by, sponsored by, or partnered with any company named here. MCP, A2A and AP2 refer to open agent/commerce protocol
© 2026 Hushh Technologies Corporation. Forward-looking architecture; roadmap items labeled in-progress are not yet generally available.

🤫

HUSSH TECHNOLOGIES CORPORATION

Personal Supercomputing Infrastructure

Technical Specifications · Systems Design · Next Generation

A specification for the next generation of personal computing: a fully on-device, consent-first data agent and the infrastructure that lets a person's own machines — and the AI agents acting for them — hold, organize, and act on personal information at supercomputing scale, without ceding custody. This document pairs the company's story with verified hardware specifications and measured results from the implemented system.

Rio Grande — the original Puppy One.

Est. August 2021. The spec every Puppy is built to.

Agentic. Personal. Sovereign.

Revision 2026-06-12 · specifications are reference values, re-verified at quote · TARGET figures are budgeted design goals

01

Origin: a break-in, and a puppy.

Hushh Technologies Corporation was founded in August 2021. Two arrivals marked that week. The first was the company — incorporated after the founder's family suffered identity theft, and after watching how casually the institutions entrusted with personal information actually handled it. The founding conviction compressed to five words: your data, your business.

The second arrival was a Bernedoodle puppy named Rio Grande. It took four years to recognize that the second explained the first. A dog is a system of remarkable capability bound to a single, legible loyalty: it fetches what you ask, guards the house, runs always-on, and never once sells the household's secrets. Nobody audits a privacy policy to trust a dog — its incentives are observable. That property — high capability with absolute, legible, single-party loyalty — is exactly what personal-data software lacks. It became the design brief for an entire class of system we call hushh One Puppy.

> **Capability without legible loyalty is surveillance. We refused to ship that.**

The problem, stated as a systems problem

Model the individual as a database. Each person's records are sharded across dozens of institutions — banks, brokerages, insurers, carriers, clinics, retailers, mail providers — each keying its shard on one natural key: the phone number. The data describes the person, but the person has the weakest read path of any party in the system: they cannot query across it, cannot see it whole, cannot revoke a downstream copy. The cloud-AI pattern of the early 2020s made this worse by offering intelligence in exchange for custody. We wanted the inverse primitive: intelligence that runs where the data already lives, under the owner's key, creating no new custodian.

02

The thesis: move compute to the data.

The classical supercomputing insight applies directly to personal data: when data is large, sensitive, or gravitationally heavy, you move computation to the data, not the data to computation. For seventy years personal information moved to the mainframe, then the server, then the cloud — each migration widening the gap between the owner and control. The hardware finally makes the inverse possible. The most capable computers ever manufactured now fit in a pocket, on a desk, beside a chair, with enough local compute and unified memory to run capable models entirely on-device.

Next-generation personal supercomputing infrastructure is the layer that exploits this: it lets a person's own machines hold and process their own information, and lets trusted software and AI agents use that information efficiently — with consent, with a complete audit trail, and without taking custody. Three invariants fix the architecture.

- **Locality.** Personal data is processed on hardware the owner controls; models run on-device for the common case. Remote compute is the exception, provisioned in a tenant the owner governs — never a new custodian.
- **Legible consent.** Every access emits a cryptographically signed, content-addressed receipt and an entry in an append-only, hash-chained transparency log the owner can read in full. Consent is a protocol, not a checkbox.
- **Single-party sovereignty.** One principal. No second reader — not the vendor, advertiser, or model provider. Revocation is a first-class, low-latency operation.

03

System architecture

The system separates a data plane (where personal data is read, transformed, stored — always on owned hardware) from a control plane (consent grants, receipts, the transparency log, device lifecycle — small, auditable, the only component requiring coordination). Personal data never traverses the control plane; the control plane carries capabilities and proofs, not content. This is what allows a complete audit trail to exist without the audit system ever seeing the data.

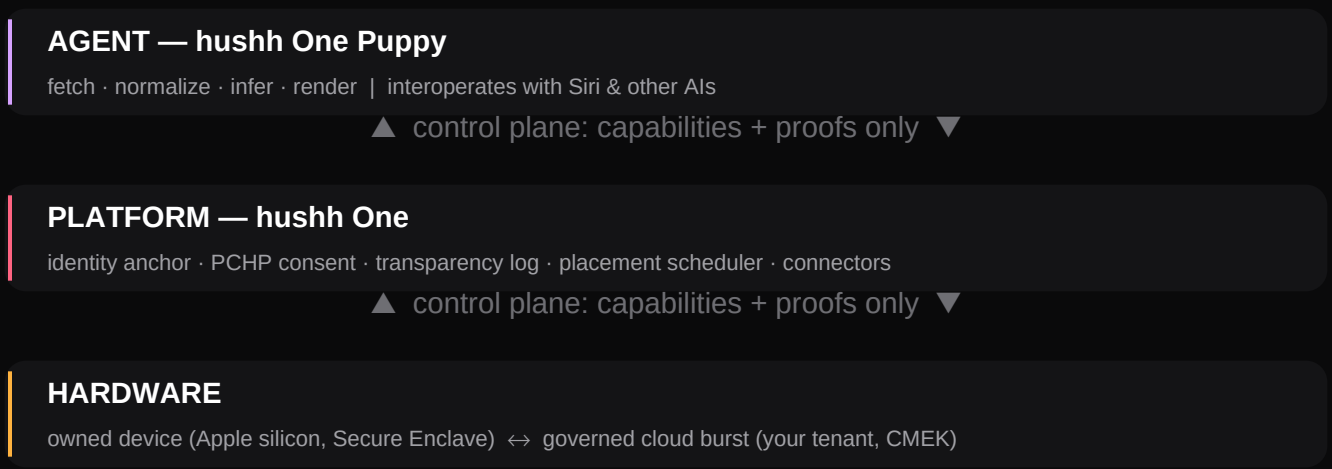

Each layer is independently substitutable. The agent is loyal to one principal and interoperates with platform assistants via intents. The platform is the durable IP: identity anchoring on the phone-number key, the consent engine, the transparency log, the connector adapters, and the placement scheduler that decides on-device versus burst. The hardware layer spans owned silicon and a governed cloud tenant.

04

Hardware specifications

The infrastructure is hardware-plural by design: it runs on the owner's existing device and scales to personal supercomputers when the workload demands. All figures are reference values from public vendor materials (June 2026); re-verified at quote.

On-device tier — Apple silicon (shipping; agent-native)

Device	Compute / memory	Role
iPhone 17 Pro Max	A19 Pro · 12GB · Secure Enclave	Identity anchor · key root
MacBook Pro 16" M5 Max	40-core GPU w/ per-core Neural Accel · ≤128GB	Extreme on-device inference
iPad Pro M5	M5 · ≤2TB · Pencil Pro	Consent-ceremony / advisor-desk

Unified memory bandwidth removes the discrete-GPU memory wall for local models.

Personal supercomputer tier — NVIDIA (reservation; Linux runtime on roadmap)

System	Peak / memory	Class
DGX Spark (GB10)	1 PFLOP · 128GB unified · ≤200B-param models	Desk-side Ultra entry
DGX Station (GB300 Ultra)	~20 PFLOPS · 748GB coherent · 1T-param models	Flagship personal
RTX PRO 6000 Blackwell	96GB GDDR7 / card · 1–4× in tower	Local fine-tuning ceiling

Governed-cloud burst tier — Google Cloud (your tenant; CMEK)

Resource	Spec	Burst use
Ironwood TPU v7	4.6 PFLOPS FP8 · 192GB HBM3e · 7.37 TB/s · ≤9,216/pod	Model-scale train/serve · ~44%
A4 / A4X (Blackwell)	B200 192GB / GB200 NVL72 >1 EFLOP	CUDA-native bursts ·
Axion C4A (Arm)	≤72 vCPU · 576GB DDR5 · ~65% better $/perf	Always-on orchestration

Cloud is rented capacity in a tenant the owner governs — owned core, governed surge. Pricing regional; quote live.

05

PCHP: consent as a protocol

PCHP (the permission & consent layer) treats every data access as a transaction that must present a valid, unexpired, sufficiently-scoped grant and must emit a receipt. Defaults are closed: a freshly provisioned device consents to nothing until the owner grants, item by item, during the in-person setup ceremony. Receipts are content-addressed and signed by an ed25519 device key held in the Secure Enclave — a receipt proves a specific action over specific fields occurred without disclosing the field values.

CONTROL-PLANE TYPES (implemented)

```
Grant   { grant_id, principal, source, scope[], expiry, revoked, sig }
Receipt { grant_id, action, fields[], purpose, content_hash, ts, sig }
LogEntry{ prev_hash, receipt, sig, hash = sha256(prev || canon(body)) }

# revocation writes a tombstone; the scheduler refuses any future action
# whose grant chain is tombstoned — and the refusal is itself logged.
```

The transparency log is append-only and hash-chained: tampering with any past entry invalidates every subsequent hash, giving the owner a verifiable O(n) audit of all access. “What does this agent know about me, and who authorized it” becomes a query the owner can actually run.

MEASURED — implemented & tested

ed25519 receipt sign/verify: tamper to any field breaks the signature (test: pass).
Hash-chained log: mutating one past entry fails verify_chain() (test: pass).
Field-scope enforcement: withheld fields never reach the normalized picture (test: pass).
Revoked / expired grant: access raises ConsentError before any fetch (tests: pass).

06

The fetch pipeline & tail-tolerant executor

The operation that defines the product is first-fetch: from one connected account to a complete, normalized picture with a quantified insight, on the owner's device, in about a minute. The fetch stage dominates the budget and talks to sources we do not control, so it runs through a deadline-aware, fault-tolerant executor applying the tail-tolerance primitives of large-scale systems.

- **Deadline propagation.** A shared end-time; every attempt receives only the remaining budget, so the stage cannot overrun.
- **Bounded concurrency.** A semaphore caps in-flight fetches; connectors run in parallel with progressive assembly (no blocking on the slowest source).
- **Retry + circuit breaking.** Full-jitter exponential backoff for transient failures; a per-source breaker opens after N consecutive faults and half-opens after cooldown.
- **Graceful partial results.** On deadline, return completed sources and mark the rest degraded; the insight reports coverage. Bounded latency with degradation beats unbounded waiting for completeness.

MEASURED — graceful degradation under a hung source

Setup: 3 sources; one deliberately hung for 5.000s; stage deadline 1.000s.
Result: pipeline returned in 1.004s with a correct picture from 2 healthy sources.
Hung source cleanly marked degraded (reason: timeout); insight: 'partial: 2/3 sources'.
Transparency log held exactly 2 entries — no receipt for data never accessed.
Suite: 16 tests (consent, log, pipeline, executor) — all passing.

SLO

Latency budget for first-fetch (TARGET)

The defining operation is budgeted stage-by-stage with headroom and instrumented so the target is measured, not asserted. p50 is materially below the tail target.

Stage	Budget	Engineering note
OAuth consent + token	≤ 8 s	human-paced; one tap, scoped read grant
Source enumeration	≤ 6 s	identify financial senders/domains, 12 mo
Fetch + parse	≤ 25 s	parallel connectors; incremental parse (dominant)
Normalize → picture	≤ 8 s	accounts, balances, recurring, renewals
Local inference	≤ 10 s	on-device model; dollar-denominated insight
Render + receipts	≤ 3 s	picture view + transparency-log writes
TOTAL	≤ 60 s	tail-latency target; bursts a step only if over budget

07

Placement, security & threat model

Placement: owned core, governed surge

Placement is a scheduling decision, not dogma. Default site of computation is the owned device, because that is where the data already is and moving data is the expensive, risky operation. Compute bursts to the governed cloud tenant only when a workload provably exceeds the device — a large fine-tune, a batch beyond local memory, a model larger than unified memory can hold — carrying confidential-computing isolation, customer-managed keys, and the same receipt/log discipline. Personal data crossing that boundary is itself a logged consent event.

Threat model

- **Central breach — mitigated by construction.** There is no central corpus; compromising the platform yields capabilities and hashes, not anyone's personal data.
- **Silent access — detectable.** Every read is grant-gated and produces a hash-chained log entry the owner can audit.
- **Custody creep — prevented.** No second reader; partner/model agents act only through scoped, revocable grants and never receive a durable copy.
- **Device loss — contained.** Keys in the Secure Enclave; data encrypted at rest; lost device revoked from the control plane and its grants tombstoned.

Not yet claimed

FedRAMP High authorization is in progress, never stated as held. An attacker with persistent code execution on an unlocked device is outside the model; we reduce blast radius (Enclave keys, least-scope grants) but claim no immunity. Honesty about the boundary is part of the specification.

08

Implementation status

A design is judged by the honesty of its status section.

Component	Status
PCHP consent: ed25519 receipts + hash-chained log	Implemented · tested
Fetch pipeline (fetch/normalize/infer/render) + budget instrument	Implemented · tested
Tail-tolerant fetch executor (deadline/concurrency/breaker/partial)	Implemented · tested
Reservation backend + Order MCP server (agent-native commerce)	Implemented · tested
60-second on-device first-fetch vs REAL sources (live OAuth)	In development — critical path
On-device model selection + inference budget	Prototyping
Governed-cloud burst boundary	Designed; not yet wired
Agent runtimes: Android / Windows / Linux	Roadmap, demand-gated
FedRAMP High authorization	In progress

The critical path is the live first-fetch. Everything else exists to make that minute trustworthy.

INVITATION

A computer that is finally yours

The personal computer promised the machine would work for the person; the cloud era quietly inverted that. Next-generation personal supercomputing infrastructure is the correction — capability that is local, loyal, legible, and owned, built to the standard a good dog sets: useful every day, and never once betraying the household.

We are recruiting systems engineers, on-device ML engineers, and security engineers who find the above under-specified in the right places. The hard, open, interesting problems are the on-device inference budget, the connector reliability surface, and the placement scheduler.

Build with us · partner with us · reserve a Puppy

engineering@hushh.ai · partners@hushh.ai · hushh.ai/one

Apple, iPhone, iPad, Mac, MacBook Pro, Siri, Secure Enclave (Apple Inc.); NVIDIA, DGX, Grace, Blackwell (NVIDIA Corp.); Google, Pixel, Tensor, Titan, Ironwood, Axion, Google Cloud (Google LLC) are trademarks of their owners, used nominatively.

Hushh Technologies Corporation (“hushh”) is independent and not affiliated with, endorsed by, sponsored by, or partnered with any company named here. MCP/A2A/AP2 denote open agent & commerce protocols.

© 2026 Hushh Technologies Corporation. Hardware specifications are public reference values, re-verified at quote. TARGET figures are budgeted design goals, not measured production results.