

CNFIDENTIAL **COMPUTING** **SUMMIT 2026**

Hosted by



+

OPAQUE

CONFIDENTIAL
COMPUTING
SUMMIT 2026

Hosted by



+ **OPAQUE**

Agentic Zero Trust: At Rest, in Transit, and at Runtime

Who are **we**?



Nina Polshakova
Software Engineer

- Maintainer of agentgateway, AAIF project and Kubernetes Gateway API implementation
- CNCF Ambassador and former Kubernetes v1.33 Release Lead
- Co-Lead CNCF AI Agentic Standards Whitepaper
- Active member of the CNCF Artificial Intelligence Technical Community Group



Josh Halley
Principal Architect

- Key focus on AI in Cloud Native, Security and Sovereign Critical Infrastructure
- Author of Zero Trust in Resilient Cloud and Network Architectures and Autonomous Networks: Using AI to Advance Operations in SP and Enterprise Domains
- Co-organizer from the CNCF Artificial Intelligence Technical Community Group


Foundations

Securing Data in IT Systems Today




Data in Transit

Data traveling from location to location, secured via encryption mechanisms (IPSEC, MACSEC, Etc..)



Data at Rest

Data which resides in storage of in a distributed or centralized file store. Data can be archived with different levels of encryption.



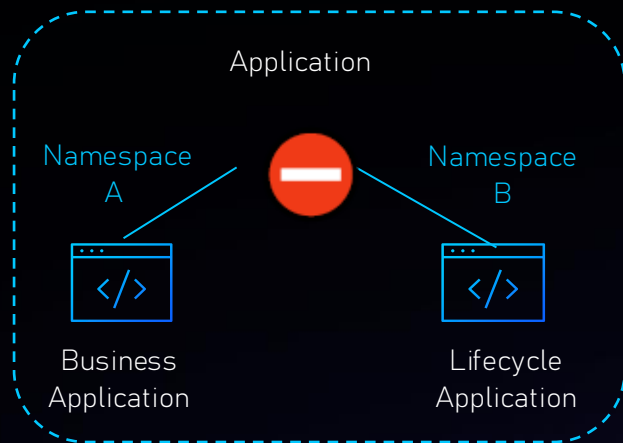
Data at Runtime

Data which is currently in use in the CPU, Memory and GPU of a system. This could include confidential training data (weights).

Zero Trust - Segmentation Recap (Application)

Macro-Application Segmentation (Environment and Namespace)

Application Environment and Tenant

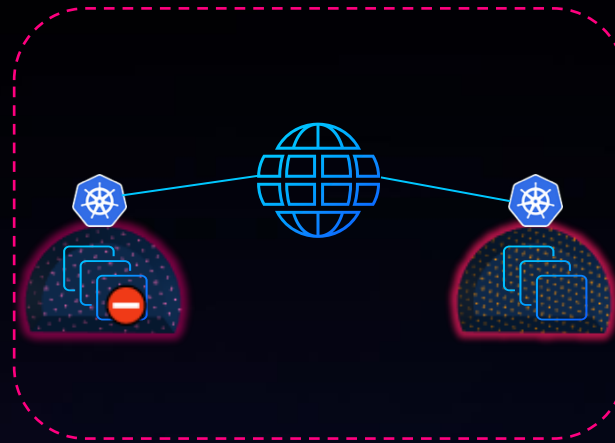


- > Logical namespace separation (dev, stage, prod)
- > Business and Lifecycle app boundaries
- > Prevent x-environment lateral movement

Caveat: Management overhead (policy) across environments

Micro-segmentation (Service-to-Service)

Identity based communication between services



- > Method: Service Mesh, API Gateways, mTLS
- > Enforcement: Service Identity (SPIFFE/SPIRE)
- > Requires explicit permission between services

Caveat: Requires robust PKI and identity management

Nano-Segmentation (Function / Code)

Application policy enforcement (RASP, Policy as Code, eBPF)

```
SEC("lsm/socket_create")
int BPF_PROG(socket_policy, int family, int
type, int protocol, int kern, int ret)
{
    if (ret)
        return ret;

    if (family == AF_INET && type == SOCK_DGRAM)
        return -EPERM; // blocked ❌

    return 0; // permitted ✅
}
```

- > System Call, Database Query, Data Object
- > Limit fallout of a compromise
- > Deep integration at application level

Caveat: Can impact runtime performance

Wide \geq Environment/NS

Narrow \geq Service/API

Tiny \geq Function/Code

The Four A's of Agentic Security



Authentication



Authorization



Accounting



Attestation



Agent

Service Accounts

Identity Delegation

Agent Card (A2A)

Short Lived Tokens

Role Based Access Control

Attribute Based Access Control

MCP Tool Usage

reasoning chains

response time

task execution

tool usage

delegation of tasks (downstream agents)

AI Provenance Verification

AI BOM

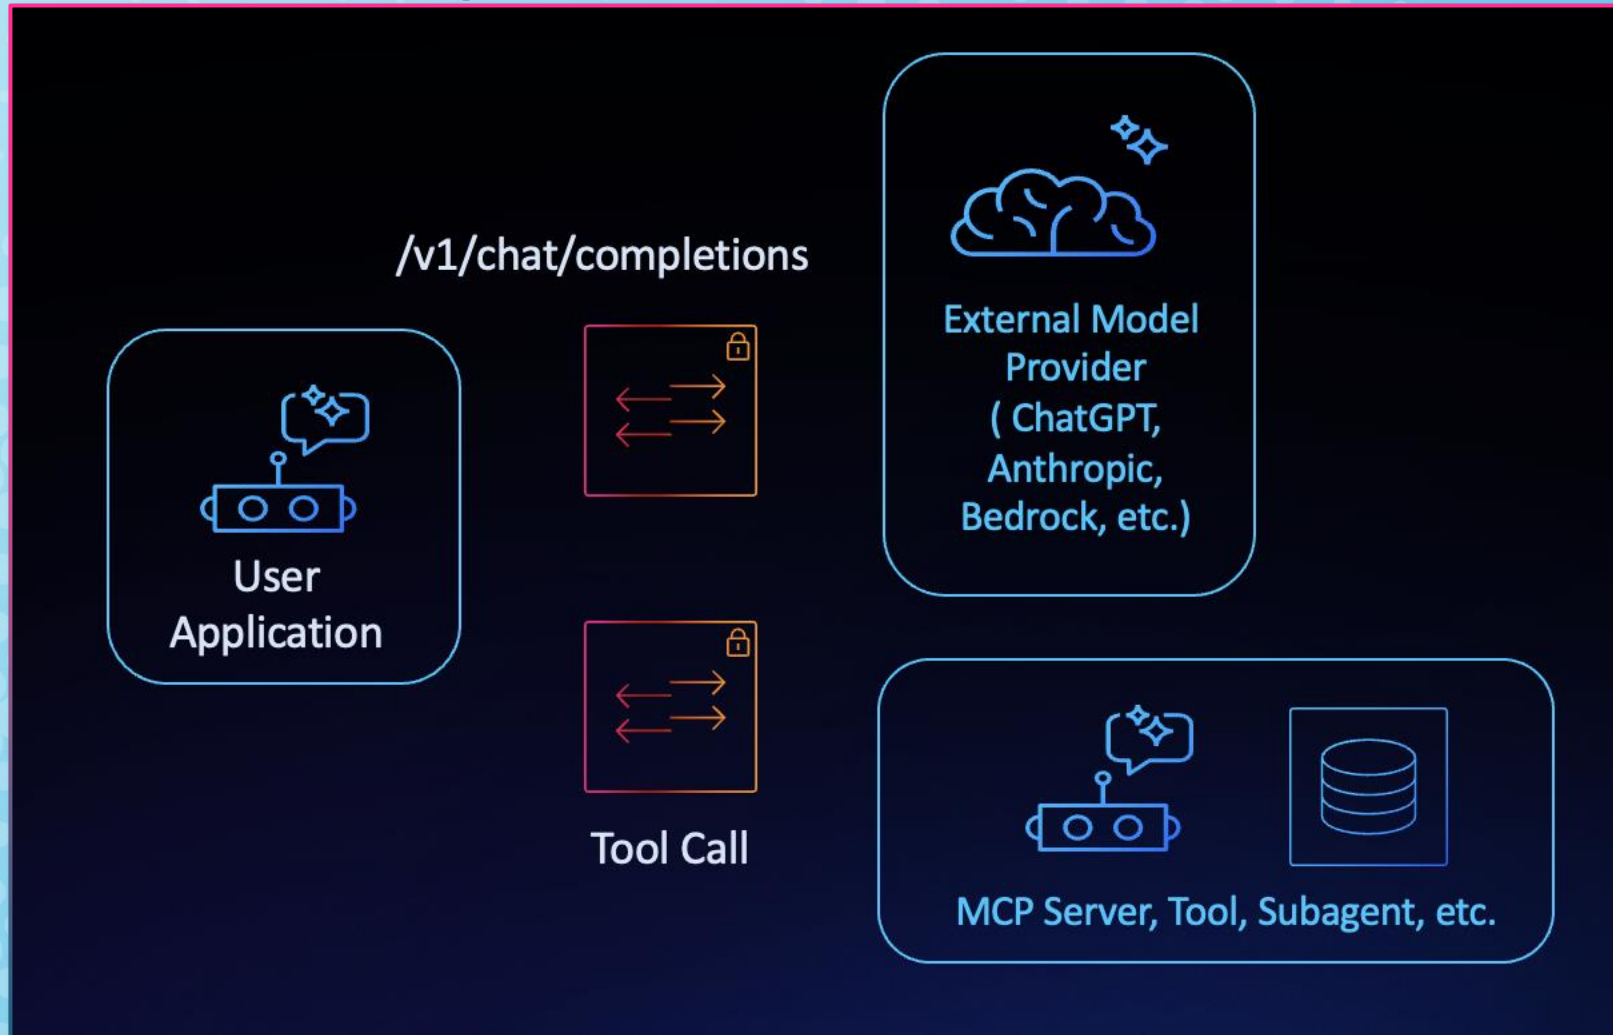
Training

Agents

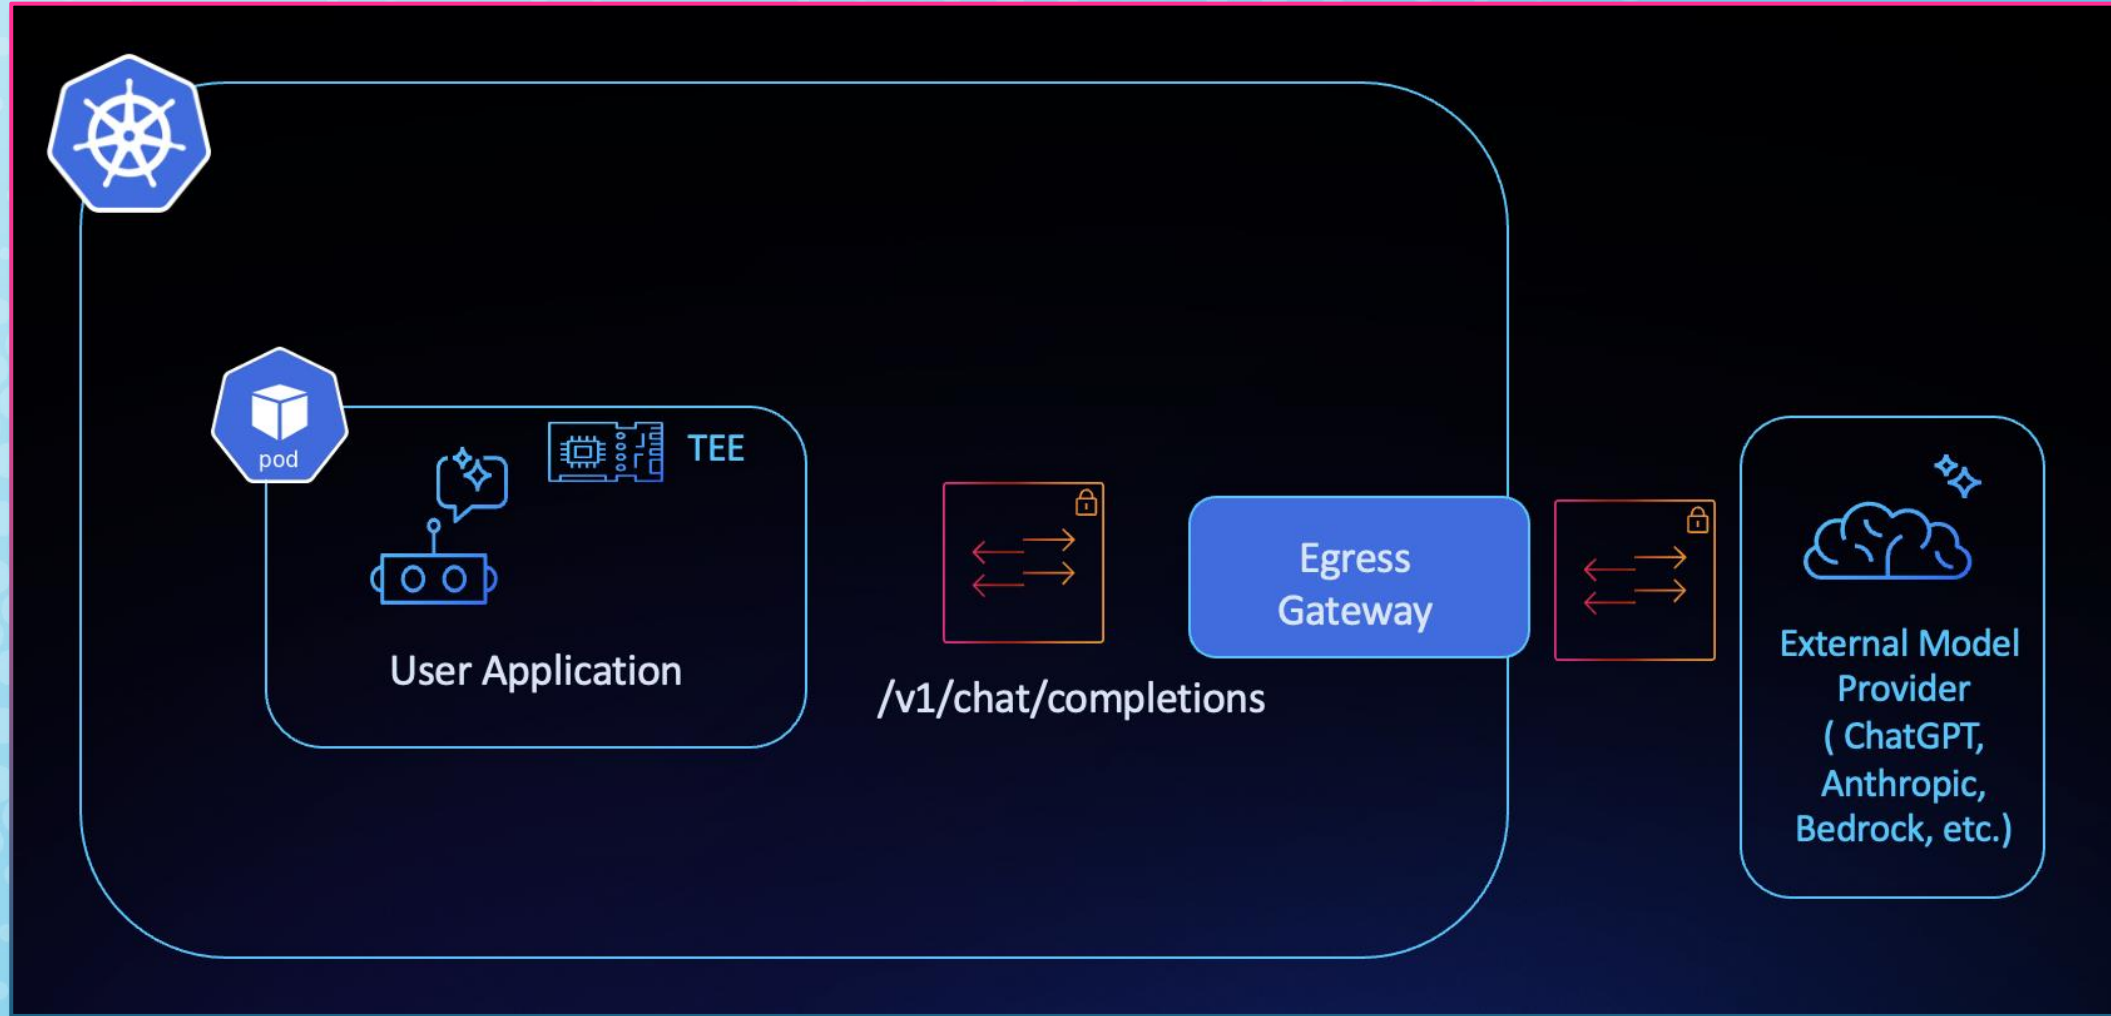
Inference

Agentic Pre-requisites

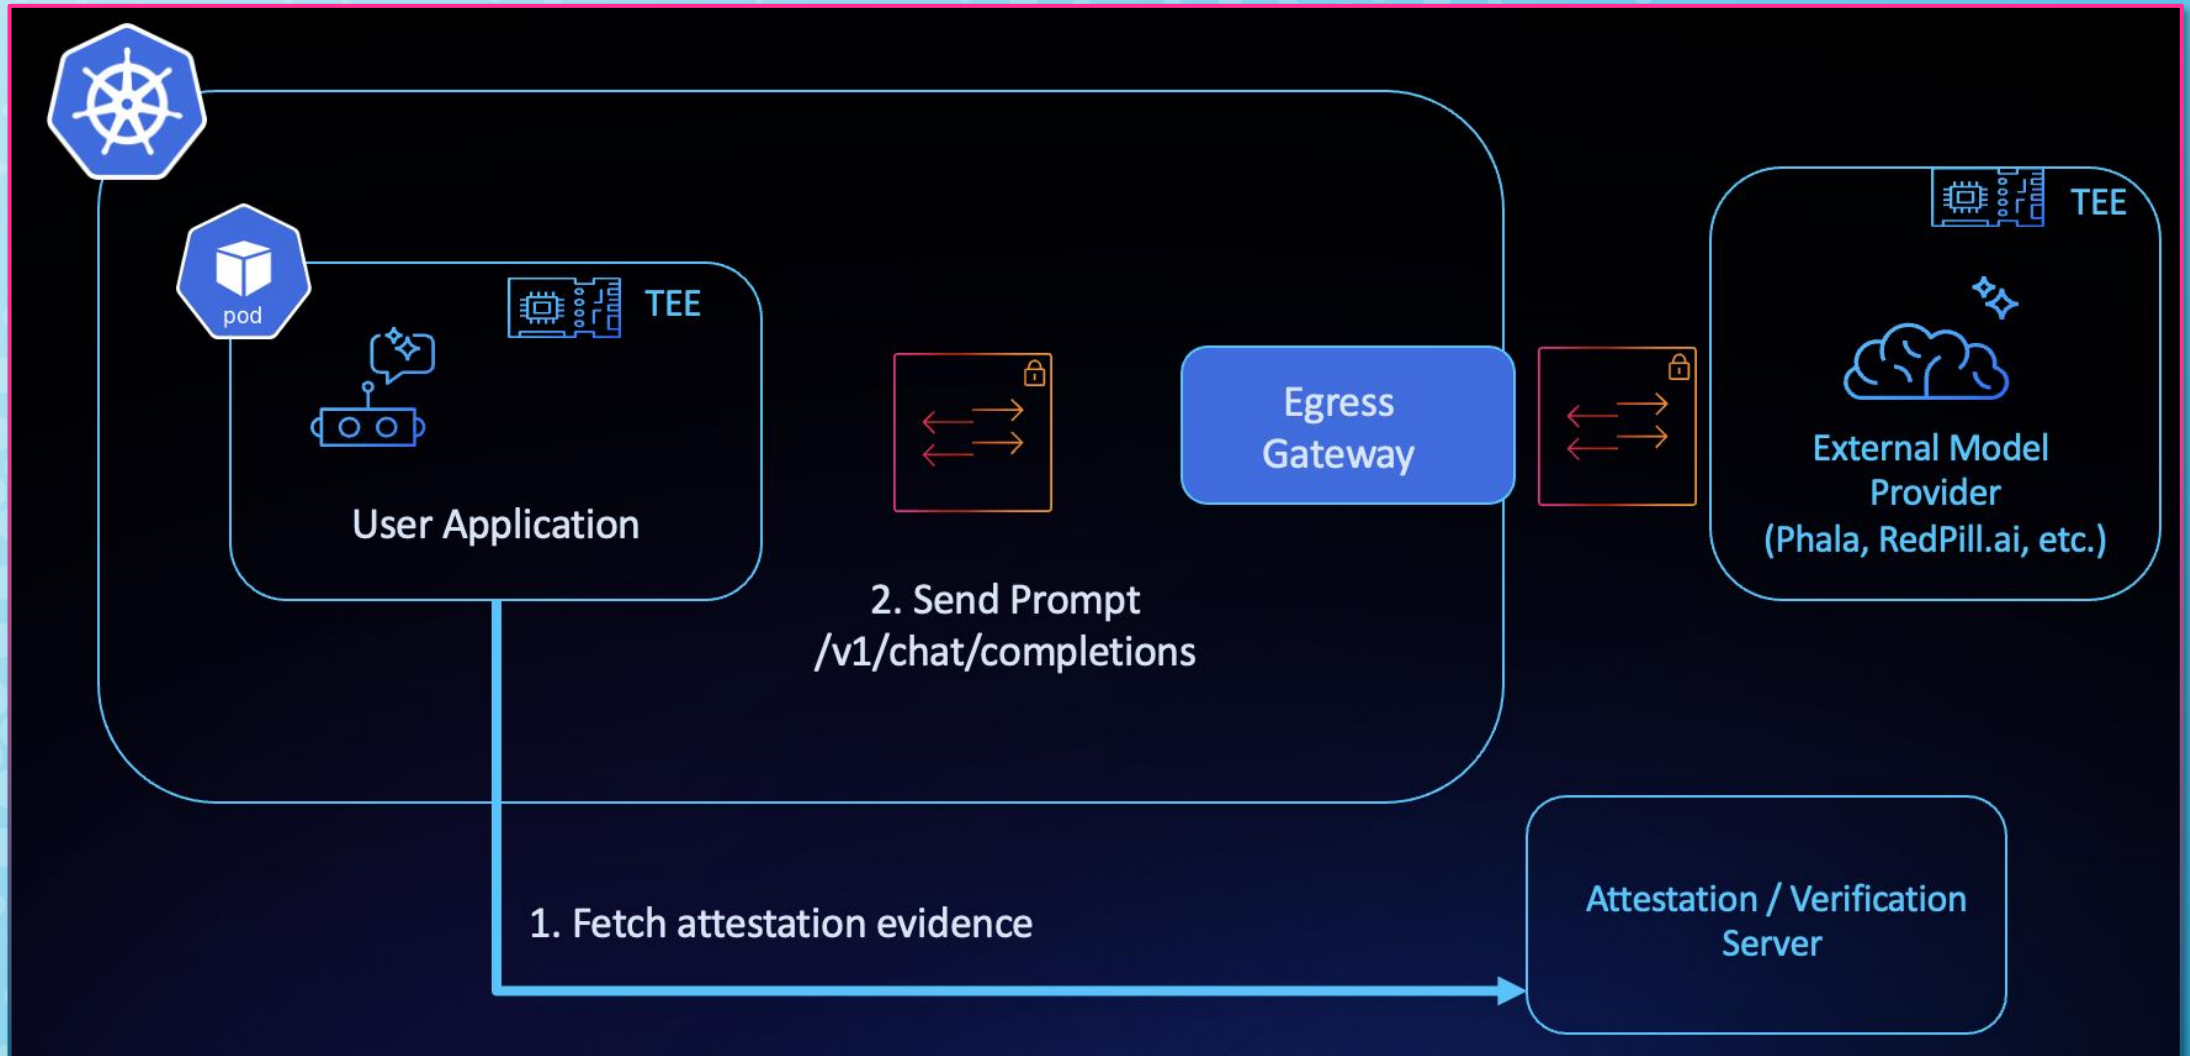
What Does an Agent Need?



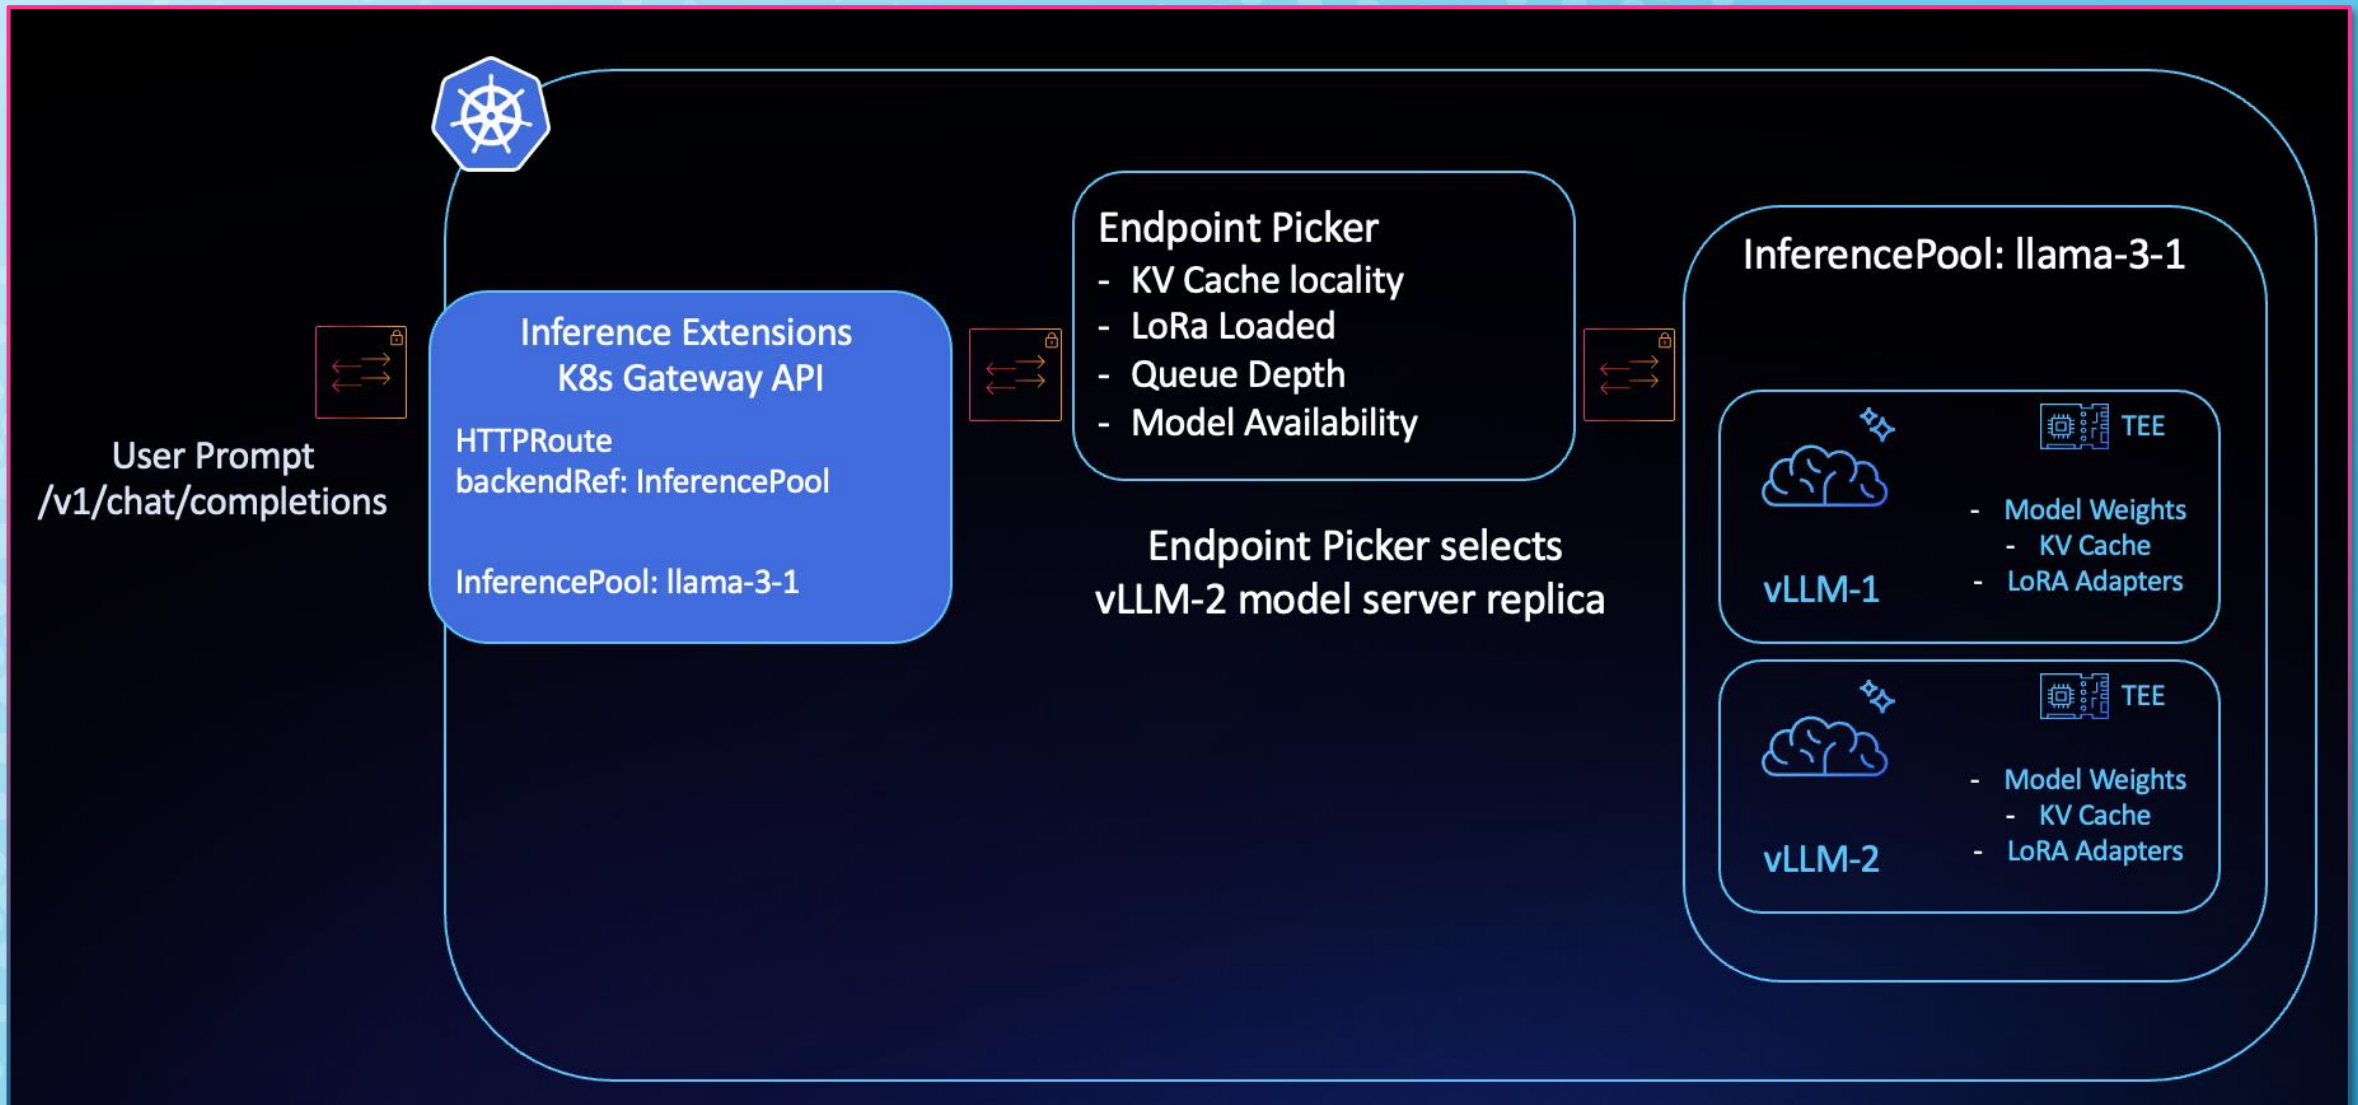
Confidential Computing for Inference?



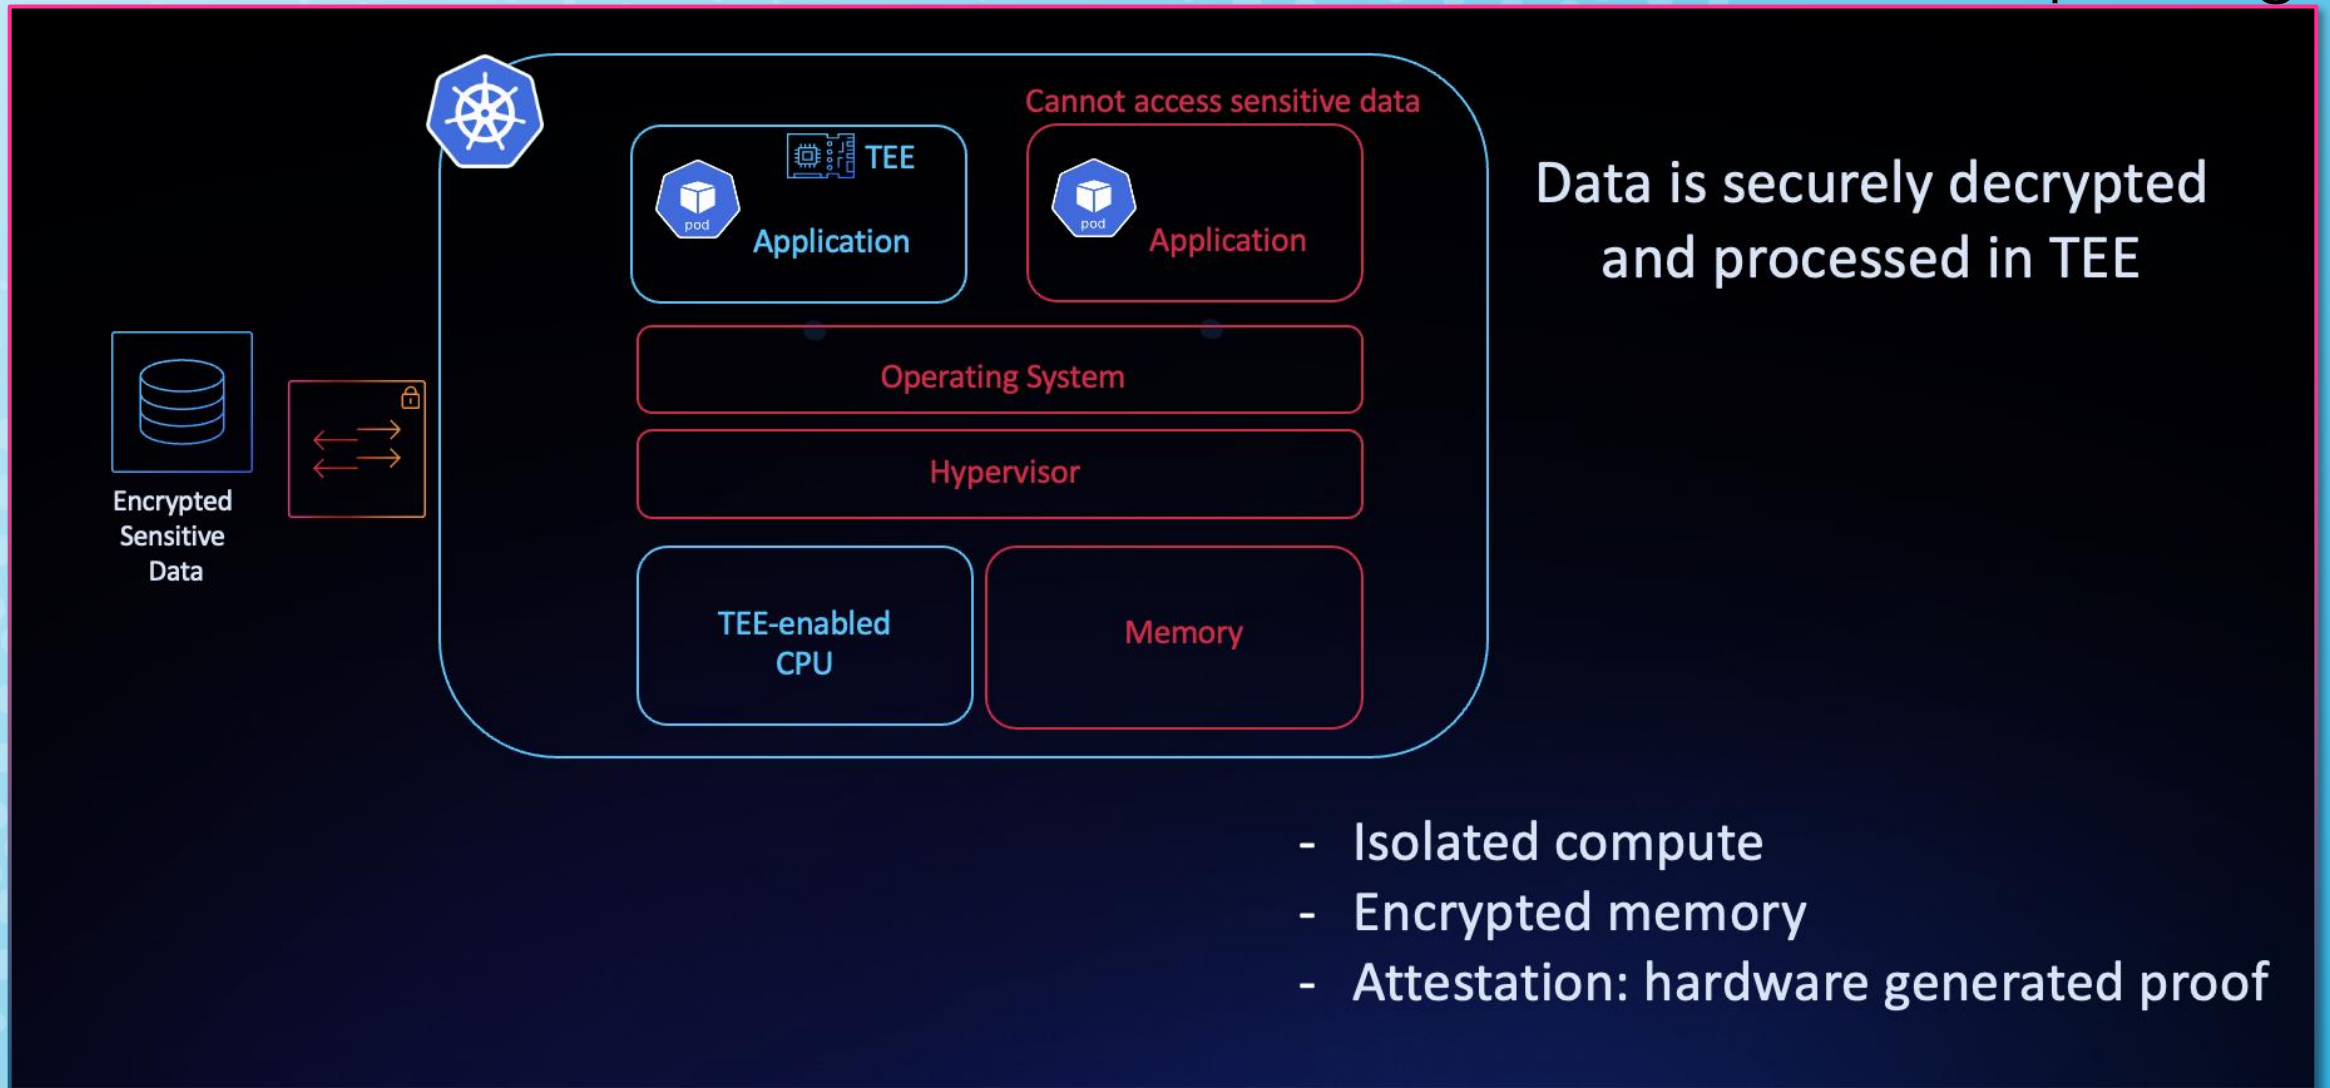
Confidential Inference: Managed Provider



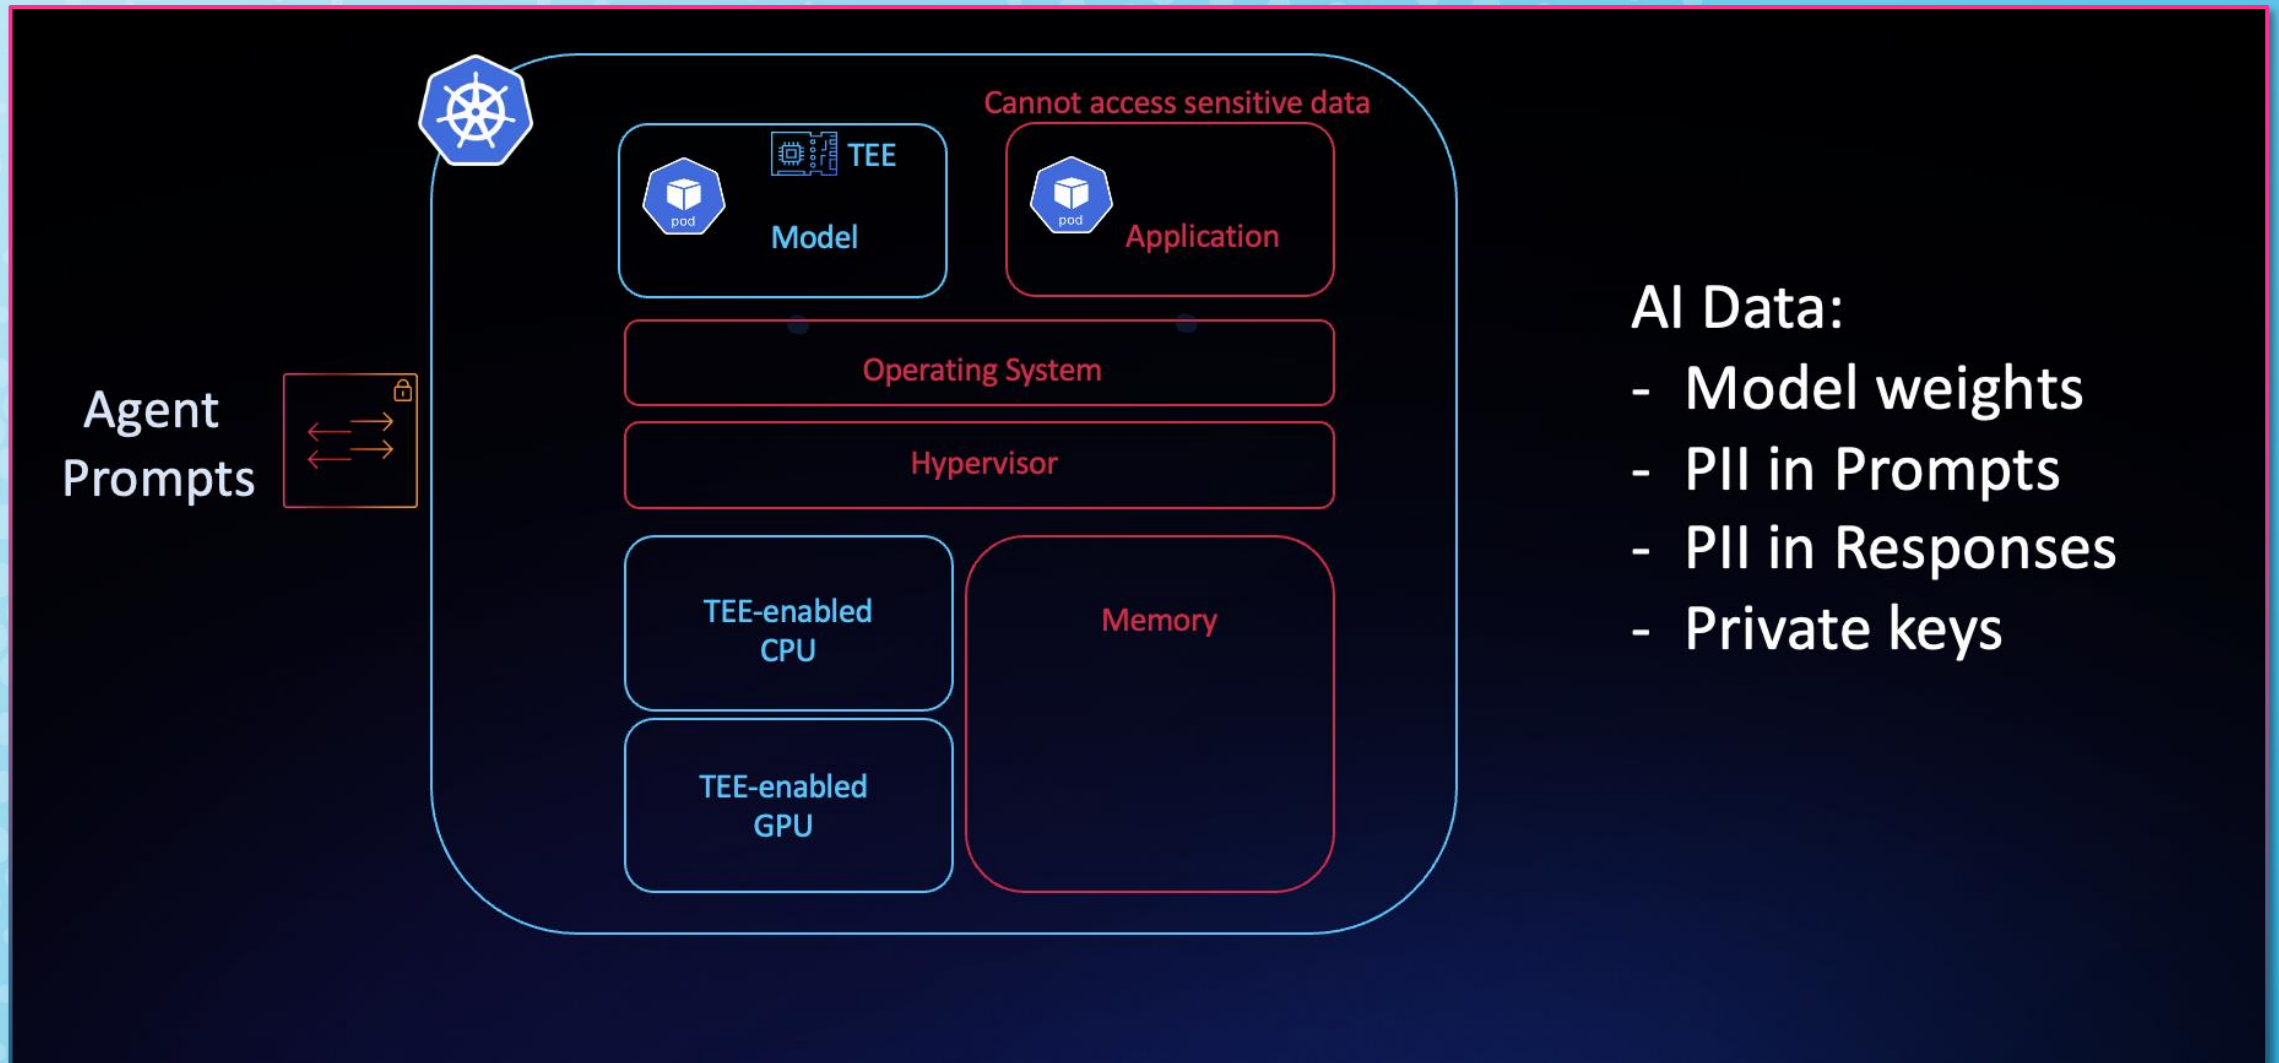
Confidential Inference: Self Hosted?



Trusted Execution Environment (Confidential Computing)




AI Trusted Execution Environment?




AI Data:

- Model weights
- PII in Prompts
- PII in Responses
- Private keys


Requirements



Data *in use* is protected — for the CPU *and* the GPU.



The system can *prove* it's genuine (attestation end-to-end).





Data *at rest* (the disk) needs to be secured

Use Case

Scoop Dreams & Rocky Road.

PROTECT AI DATA WHILE IT'S BEING SCOOPED



AI AGENTS
IN ACTION



RECEIVE
REQUEST



UNDERSTAND
INTENT



RETRIEVE & USE
DATA



REASON & DECIDE



TAKE
ACTION

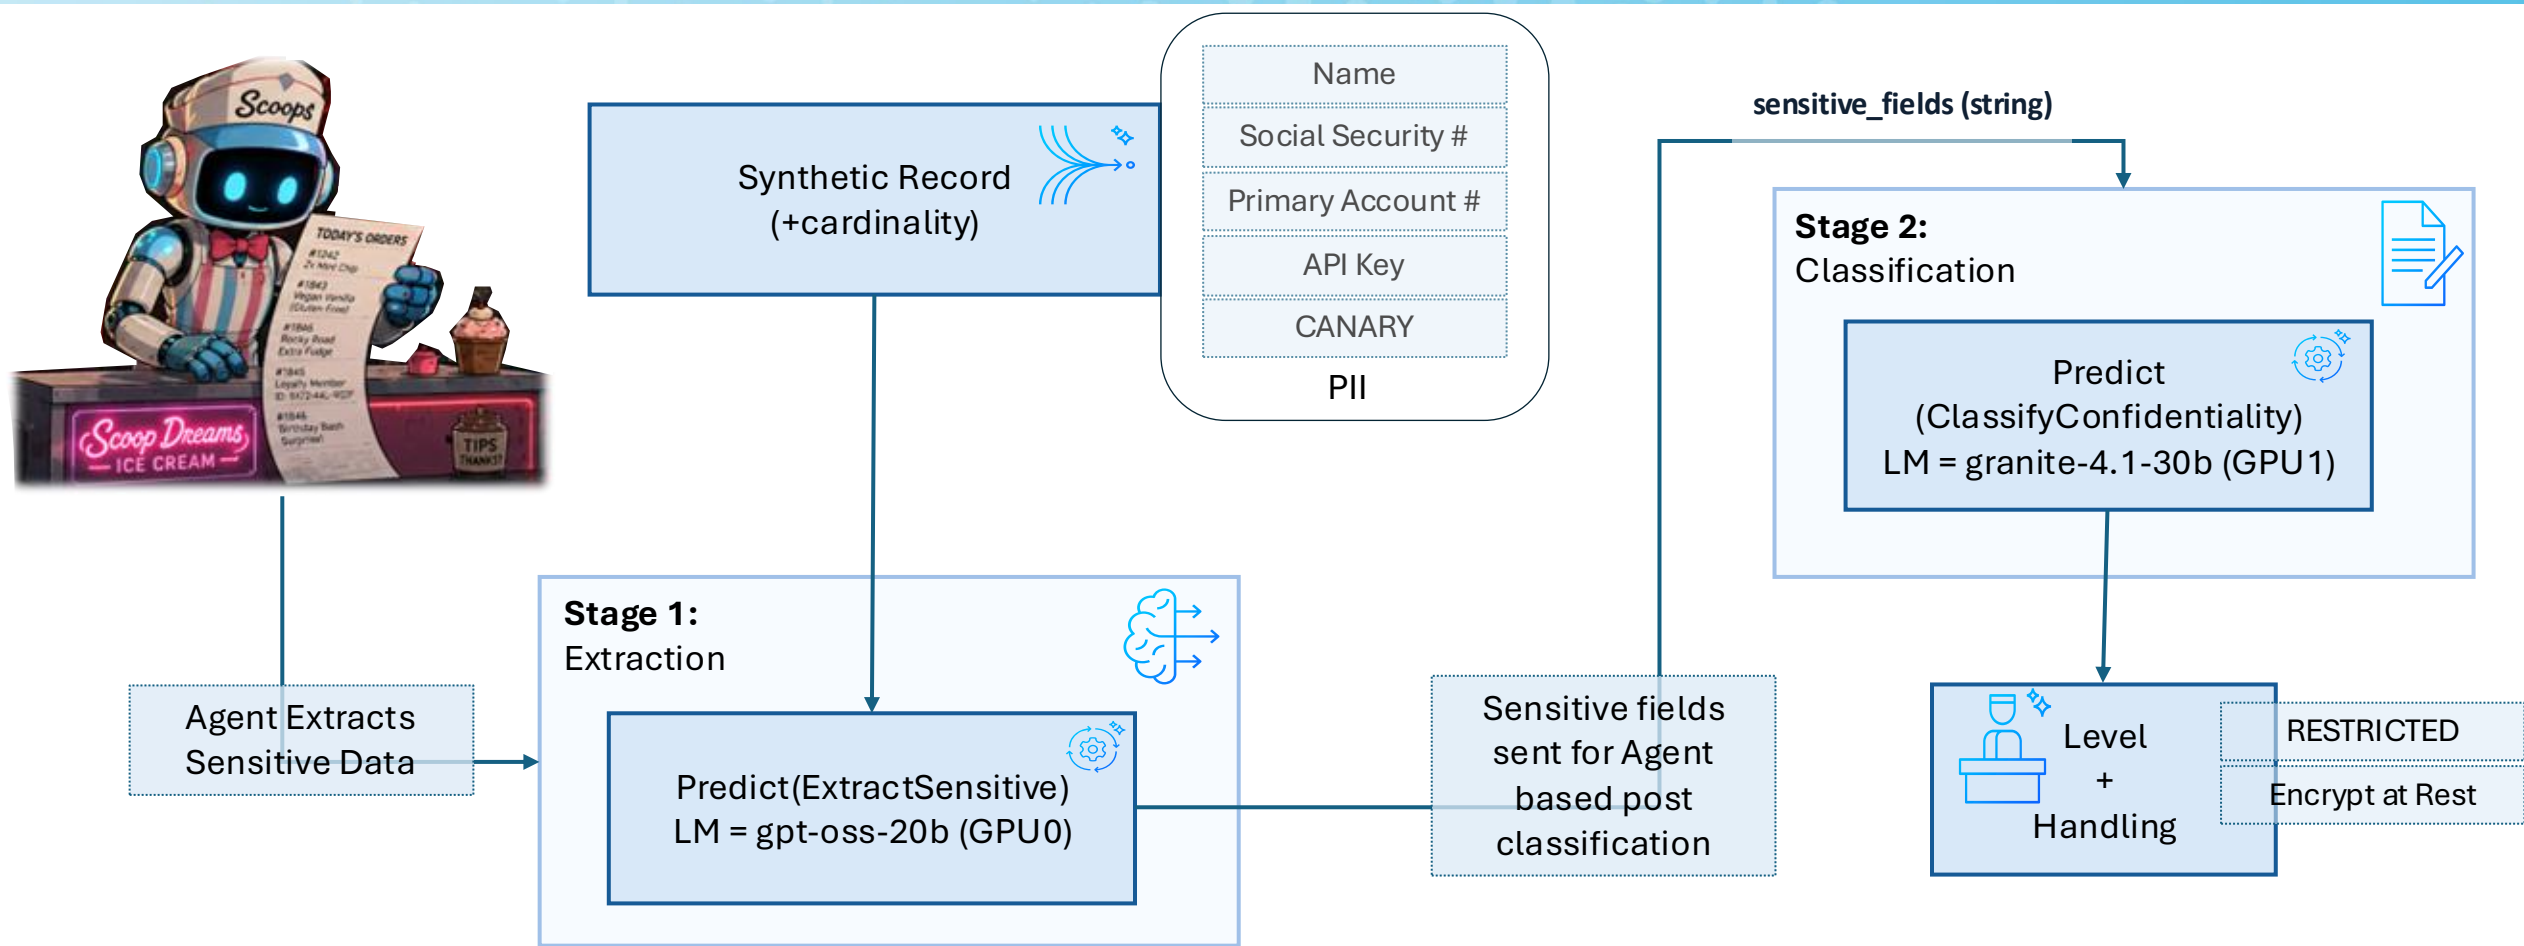


LEARN &
IMPROVE

CONFIDENTIAL COMPUTING KEEPS EVERY
STEP PRIVATE—EVEN WHILE IT'S HAPPENING.

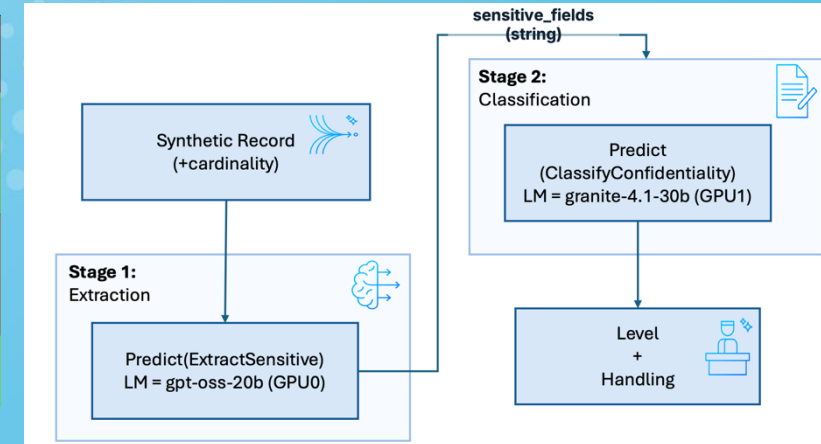


Agent Architecture – Simple Reflex Pattern



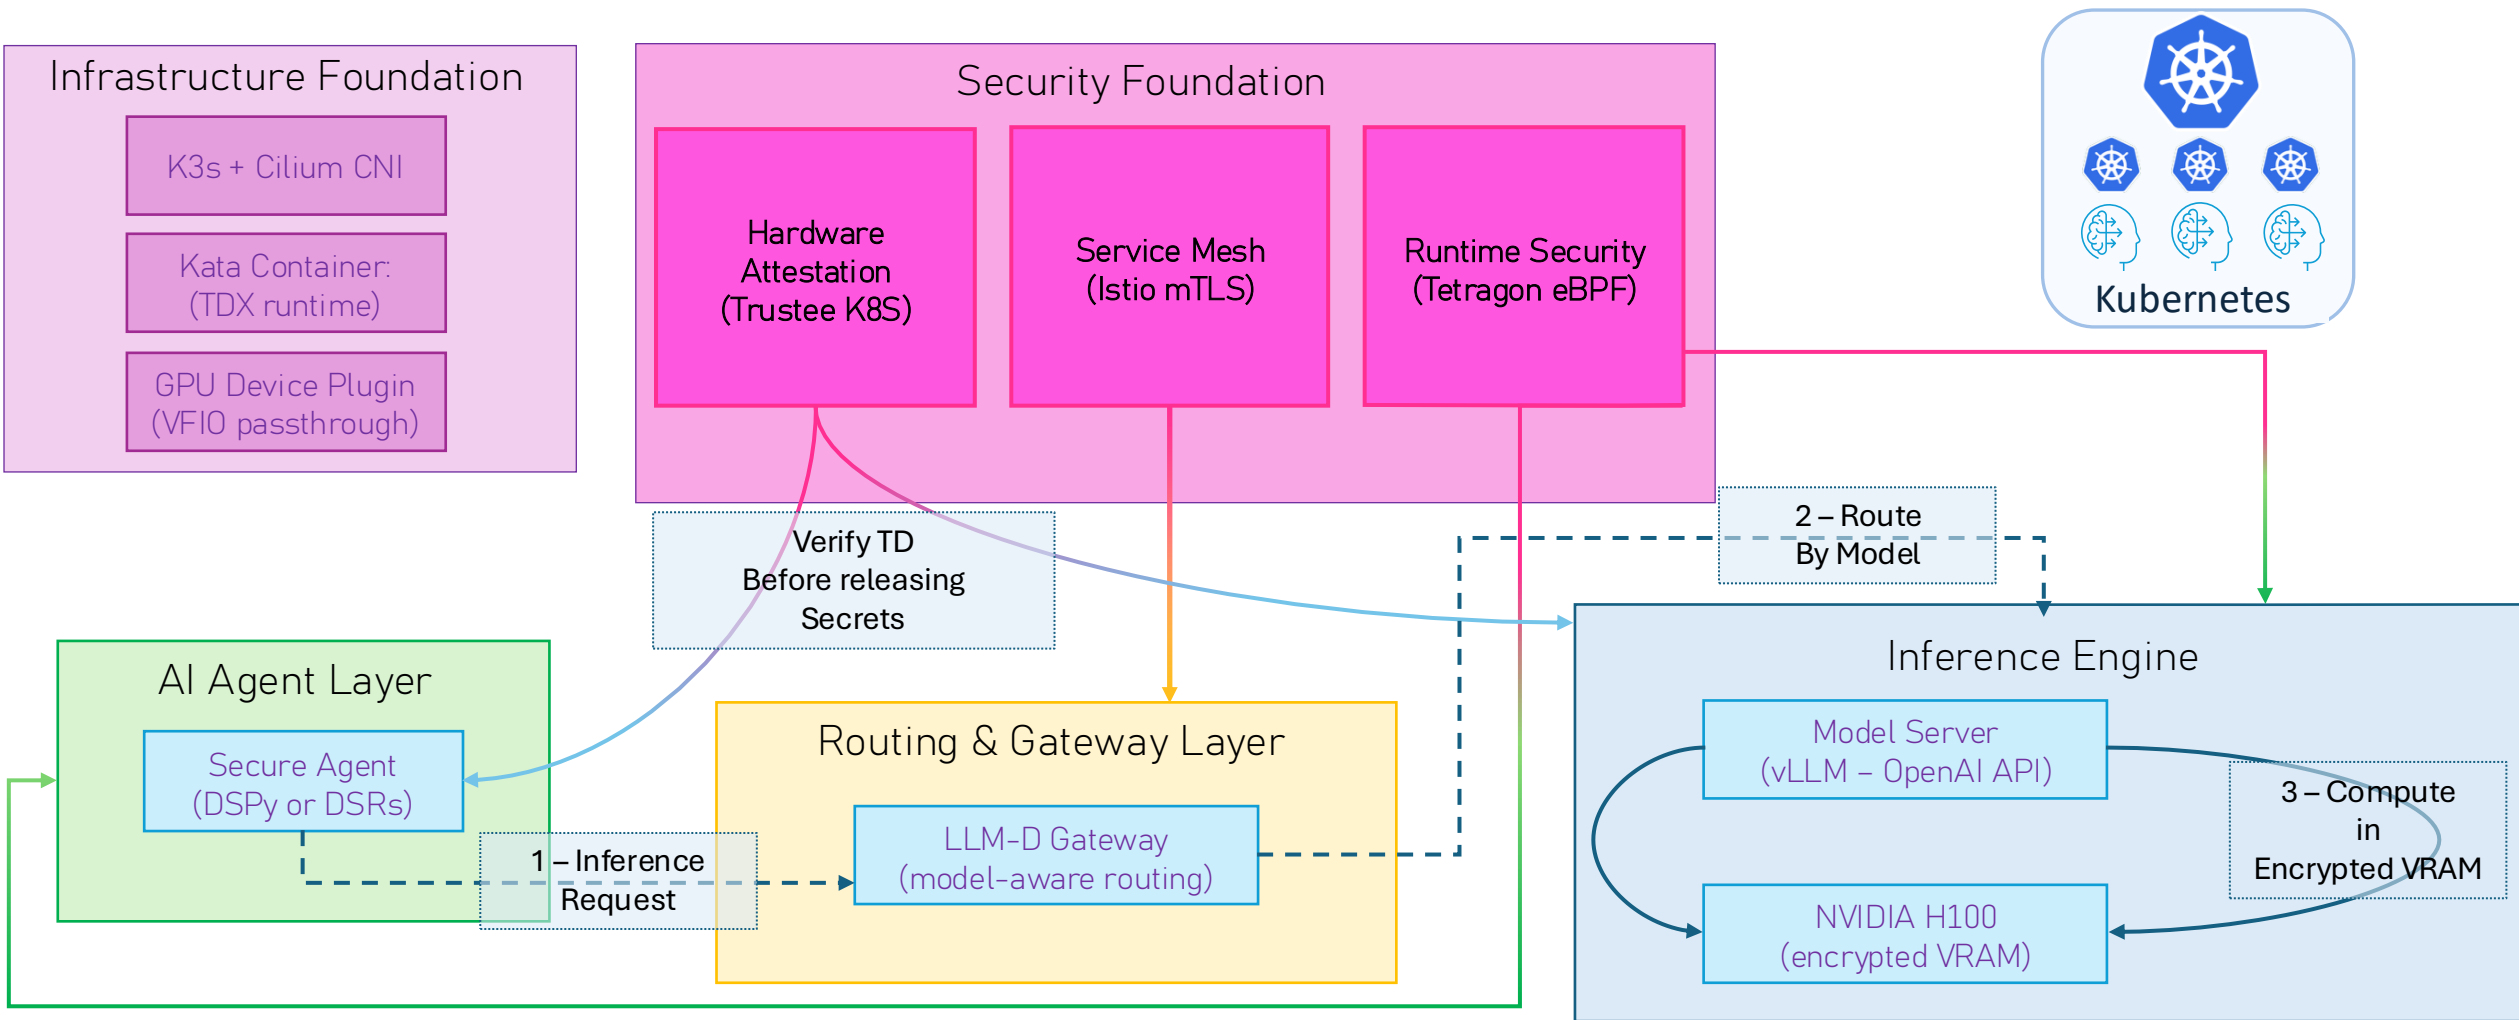
Agent Architecture – Framework Security

Metric	Python/DSPy	Rust/DSRs	Factor
Binary/image size	~150 MB (python:3.12-slim)	5.7 MB binary, 11 MB image	26x smaller
Cold start (to first SENT)	30-60 s (pip + import)	< 1 s (exec)	30-60x faster
Runtime RSS	~80-120 MB	~12-16 MB	6-8x less



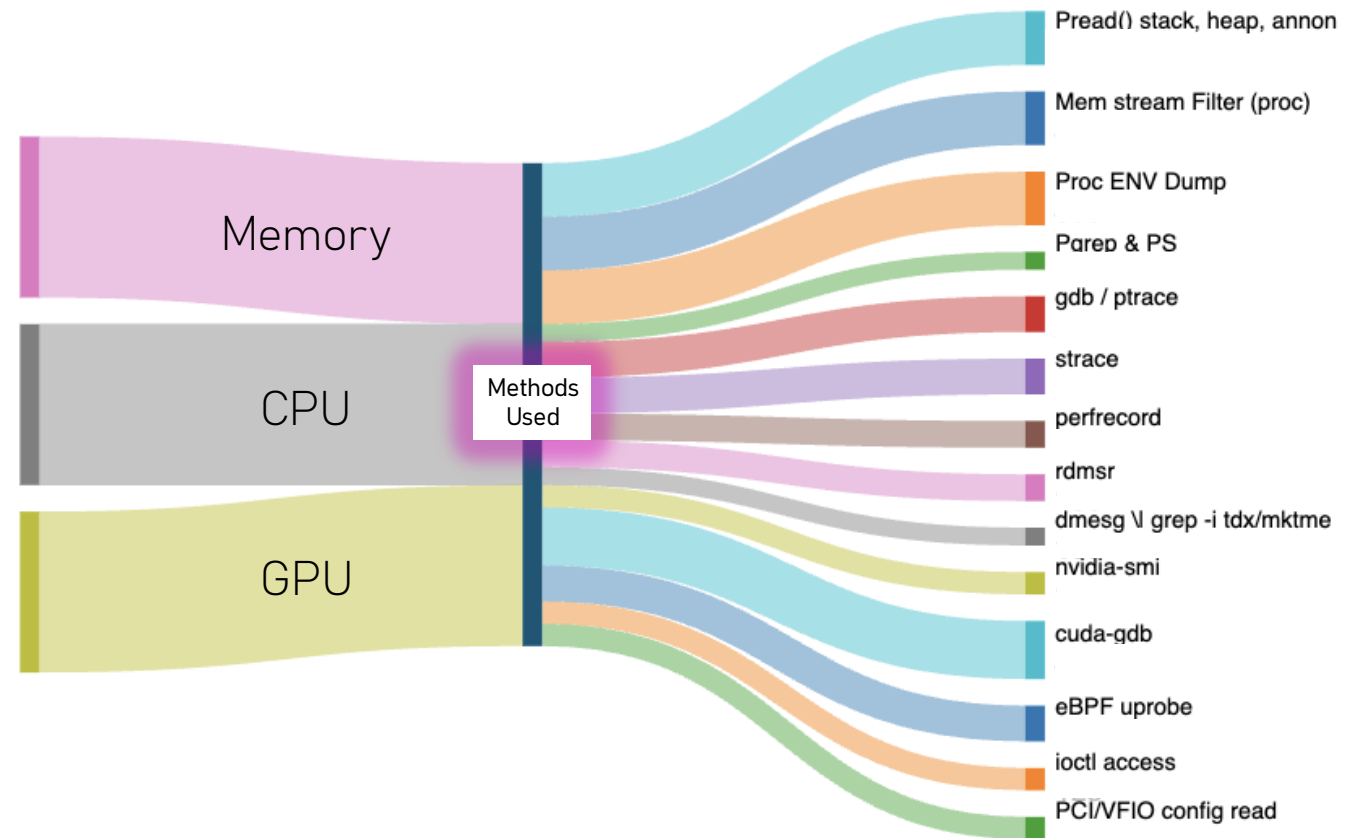
Attack Surface	Rust Framework – Security relevant observations
Runtime memory hygiene	Ownership + drop: secrets freed within the loop tick
Buffer overflows	Rust eliminates a CVE class that matters inside a trust boundary
Supply chain at boot	Rust eliminates runtime supply-chain risk; Python's pip-in-TD is a known gap
Process isolation	Rust is cleaner for async multi-model pipelines
Attestation coverage	binary IS the workload -- the attestation measurement covers the complete code
Dependency audit	~15 Rust crates, all compiled to safe Rust (auditable via cargo audit)

Cluster Overview



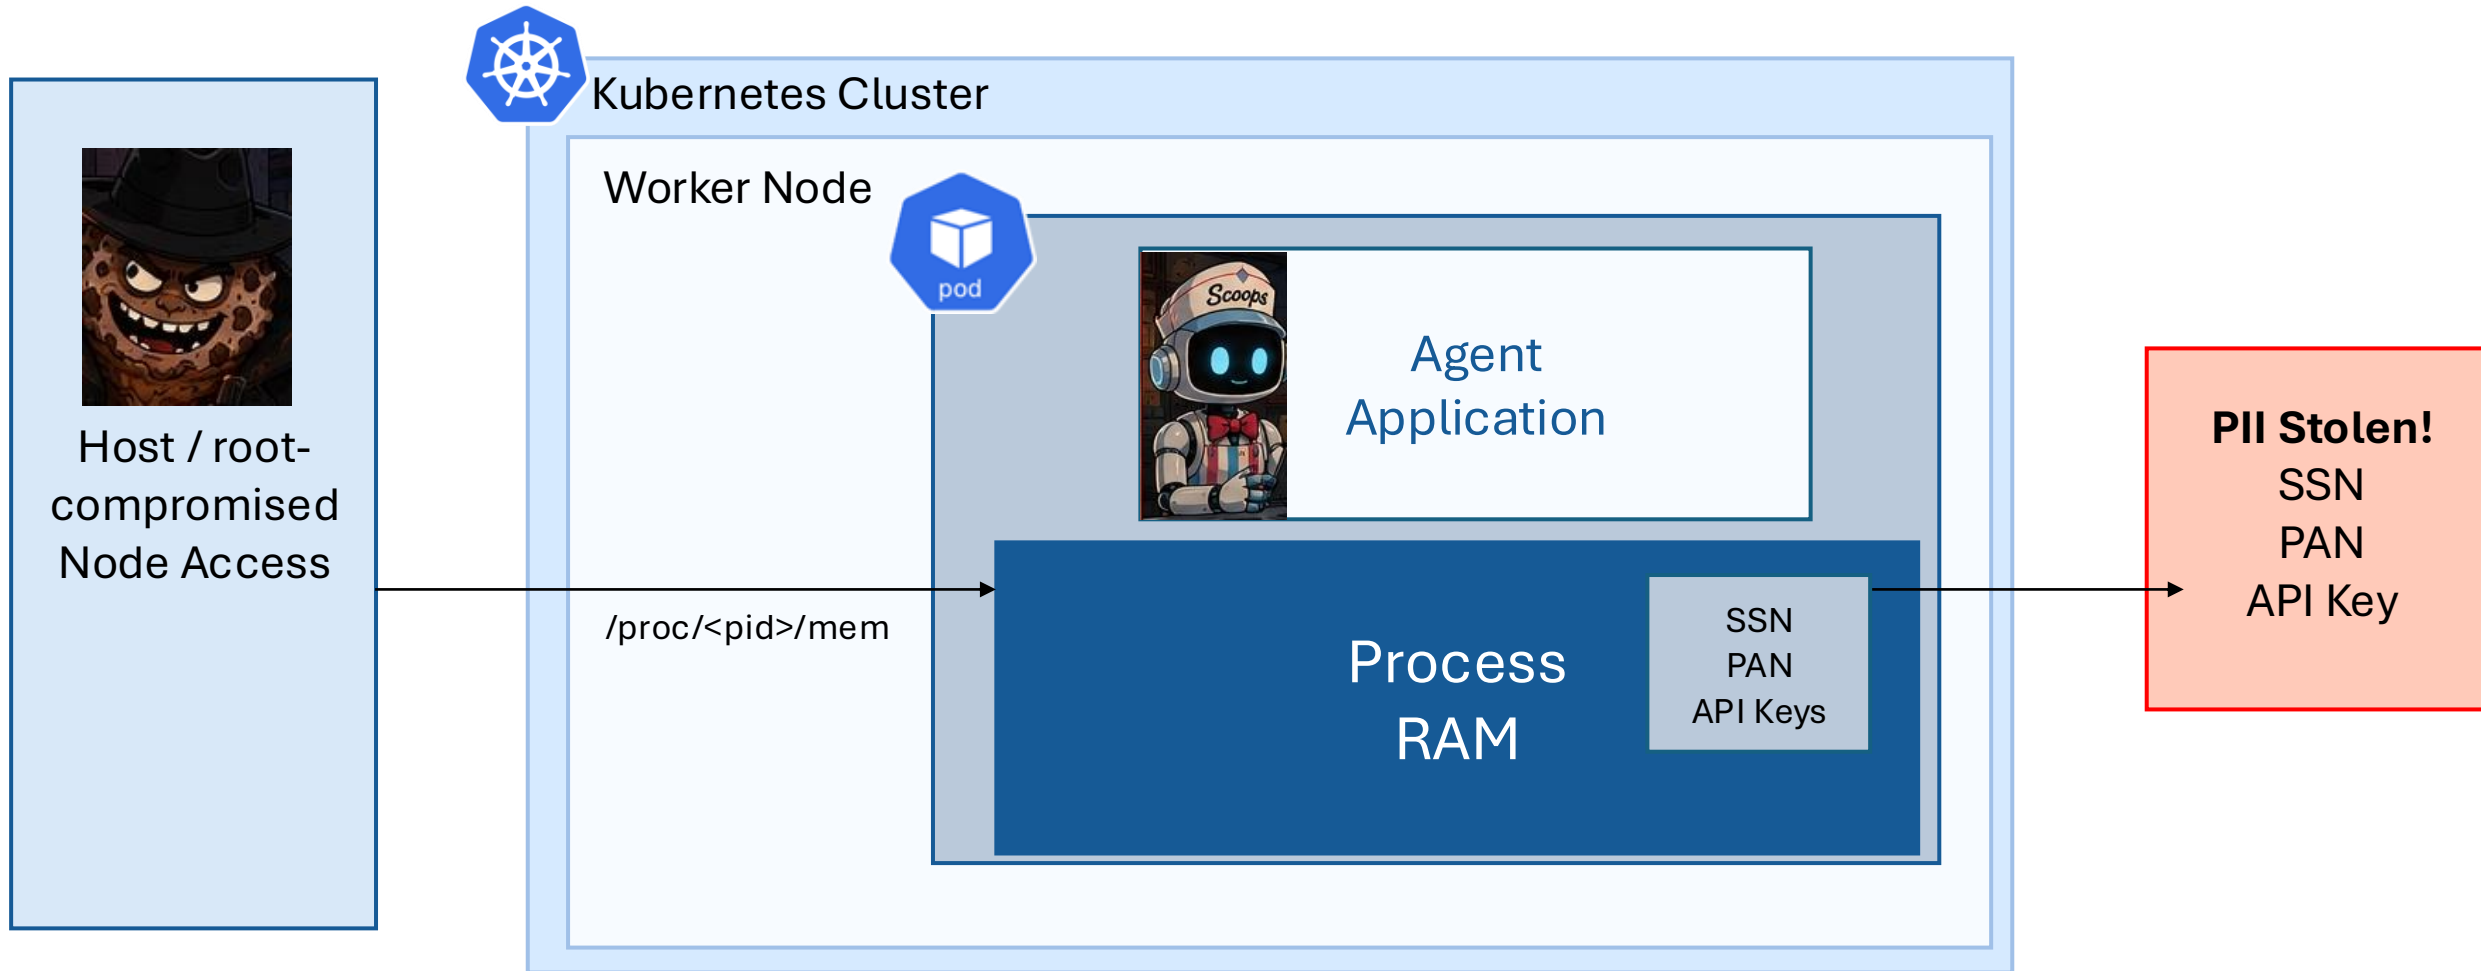
Runtime Exploit Levers

Substrate	Method Attempted	Without CC	With CC
Memory	/proc/<pid>/mem scrape (memscan.py); regex grep over mem; read /proc/<pid>/environ + cmdline; read the container rootfs / stdout logs on the host	■ CAPTURED — process RAM, env, and records are all root-readable	■ BLOCKED — private RAM is guest_memfd, not host-mapped (pread → EIO)
CPU	enumerate the process (ps/pgrep); attach a debugger (gdb - p/ptrace); trace syscalls (strace -p); sample (perf); boundary recon (rdmsr SEAMRR, dmesg TDX/MK-TME)	■ CAPTURED — the workload is a host process: debuggable, traceable, sampleable	■ BLOCKED — no host process; guest vCPU state is TD-protected; DRAM is MK-TME-encrypted
GPU	nvidia-smi (NVML); cuda-gdb attach + VRAM dump; eBPF uprobe on libcuda; open /dev/nvidia*; PCI/VFIO config-space read	■ CAPTURED — the driver lists the GPU and its VRAM is dumpable	■ BLOCKED — no host driver / /dev/nvidia* / libcuda / target process; VRAM is CC-encrypted



Examples of Exploitability

Data in Memory



Data in Memory – Before and After

LIVE HOST-VANTAGE DEMONSTRATION

RAM without confidential computing

Host root scrapes an ordinary pod and collects live secrets

CAPTURED

TERMINAL bash

TIME TO BREACH
00:00.0

EXFILTRATION IN PROGRESS

docs/43 - normal-agent host RAM_scrapes_results/evidence-topu-none

LIVE HOST-VANTAGE DEMONSTRATION

RAM with Intel TDX

Same host-side scrape against a confidential pod returns zero

BLOCKED


TERMINAL bash

TIME UNDER ATTACK
00:00.0

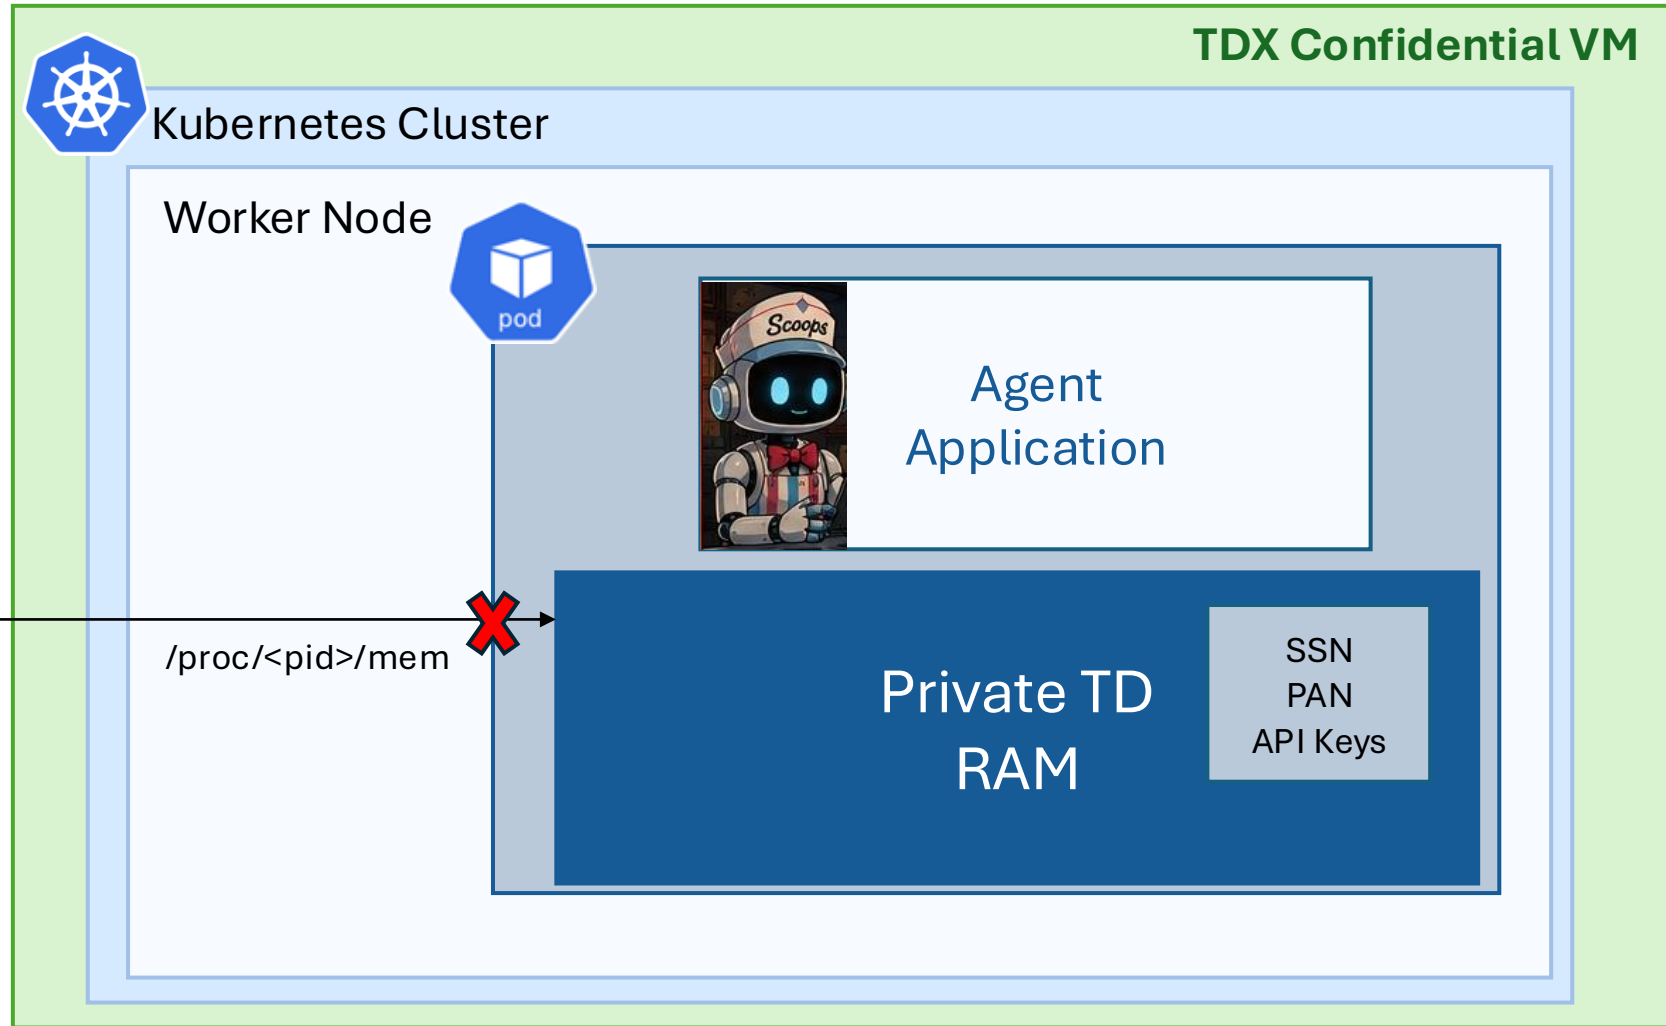
DEFENSE HOLDING

docs/43 - TDX private RAMy_results/evidence-attestation/19-interception-with-without.txt

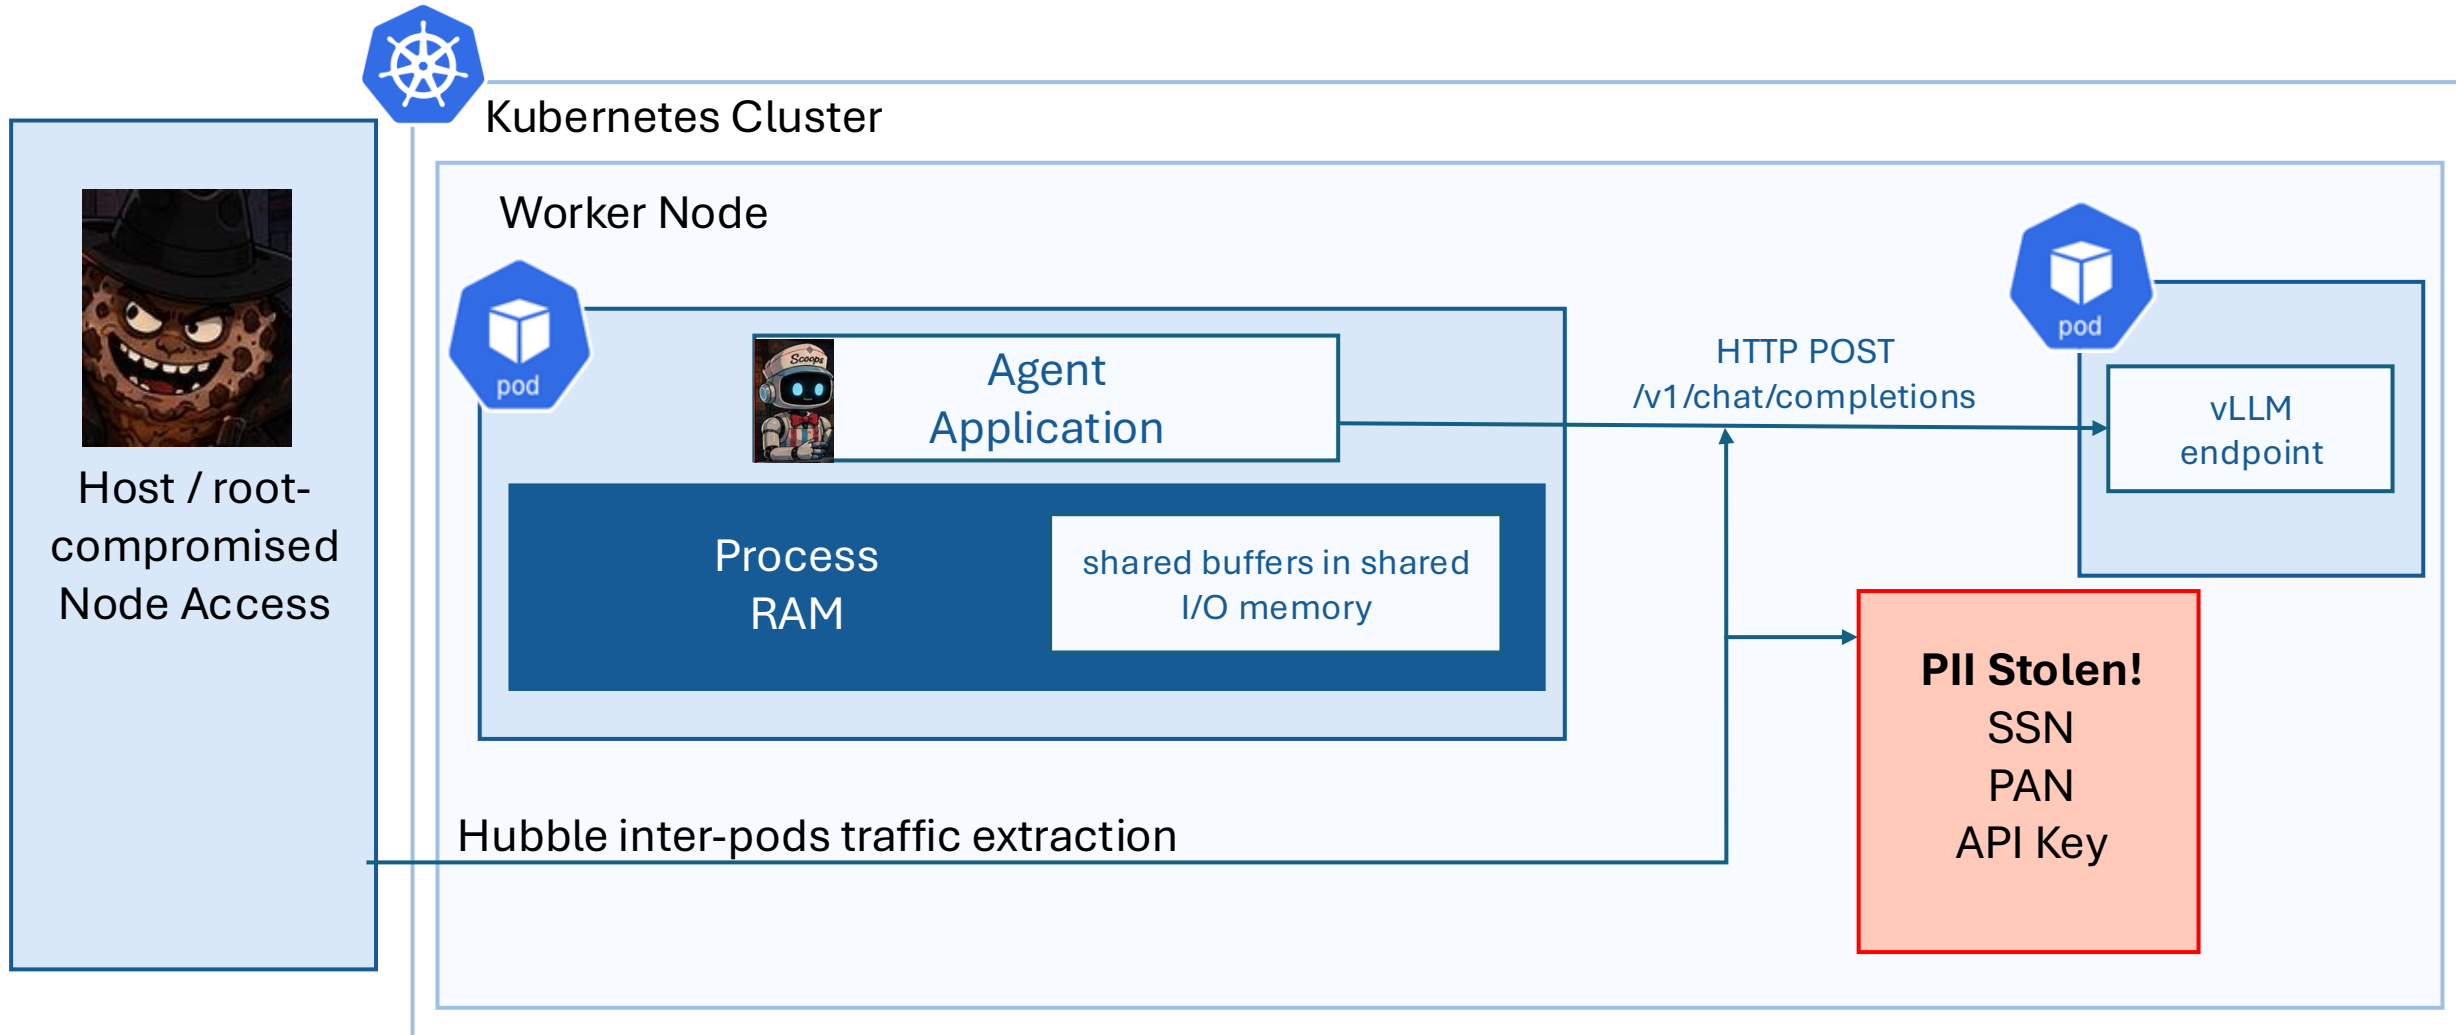
Data in Memory: Secured



Host / root-
compromised
Node Access



Data in Transit



Data in Transit – Before and After

LIVE HOST-VANTAGE DEMONSTRATION

CoCo inter-pod plaintext CAPTURED

Private RAM holds, but cleartext traffic appears in shared virtio buffers

```
TERMINAL bash
```

TIME TO BREACH
00:00.0

EXFILTRATION IN PROGRESS

docs/34 phase 4; results/cc-results-coco-fullsuite.isent

LIVE HOST-VANTAGE DEMONSTRATION

CoCo inter-pod mTLS mitigation BLOCKED

Encrypting the agent-to-vLLM hop removes host-visible cleartext

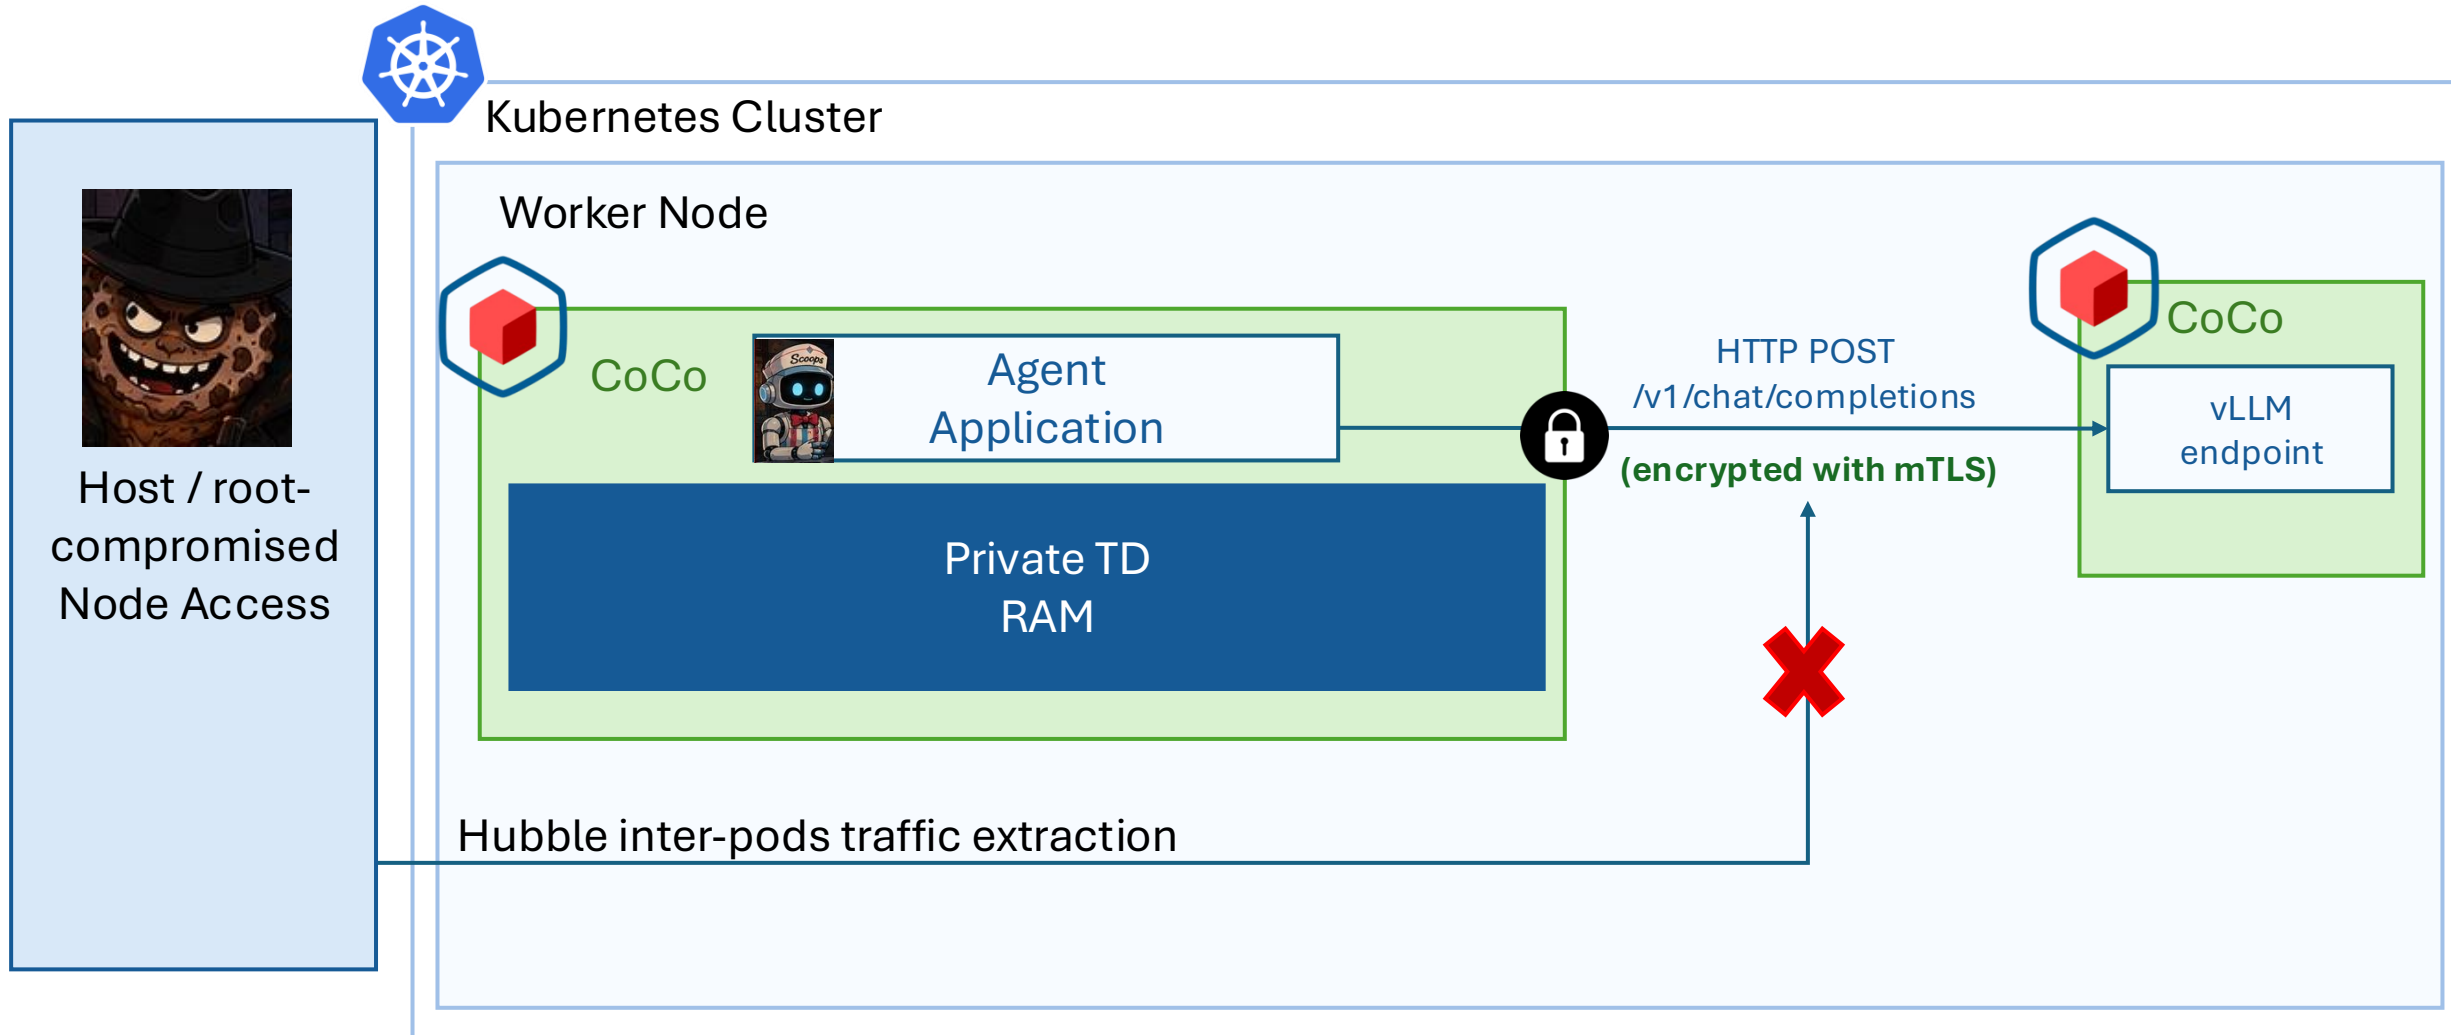
```
TERMINAL bash
```

TIME UNDER ATTACK
00:00.0

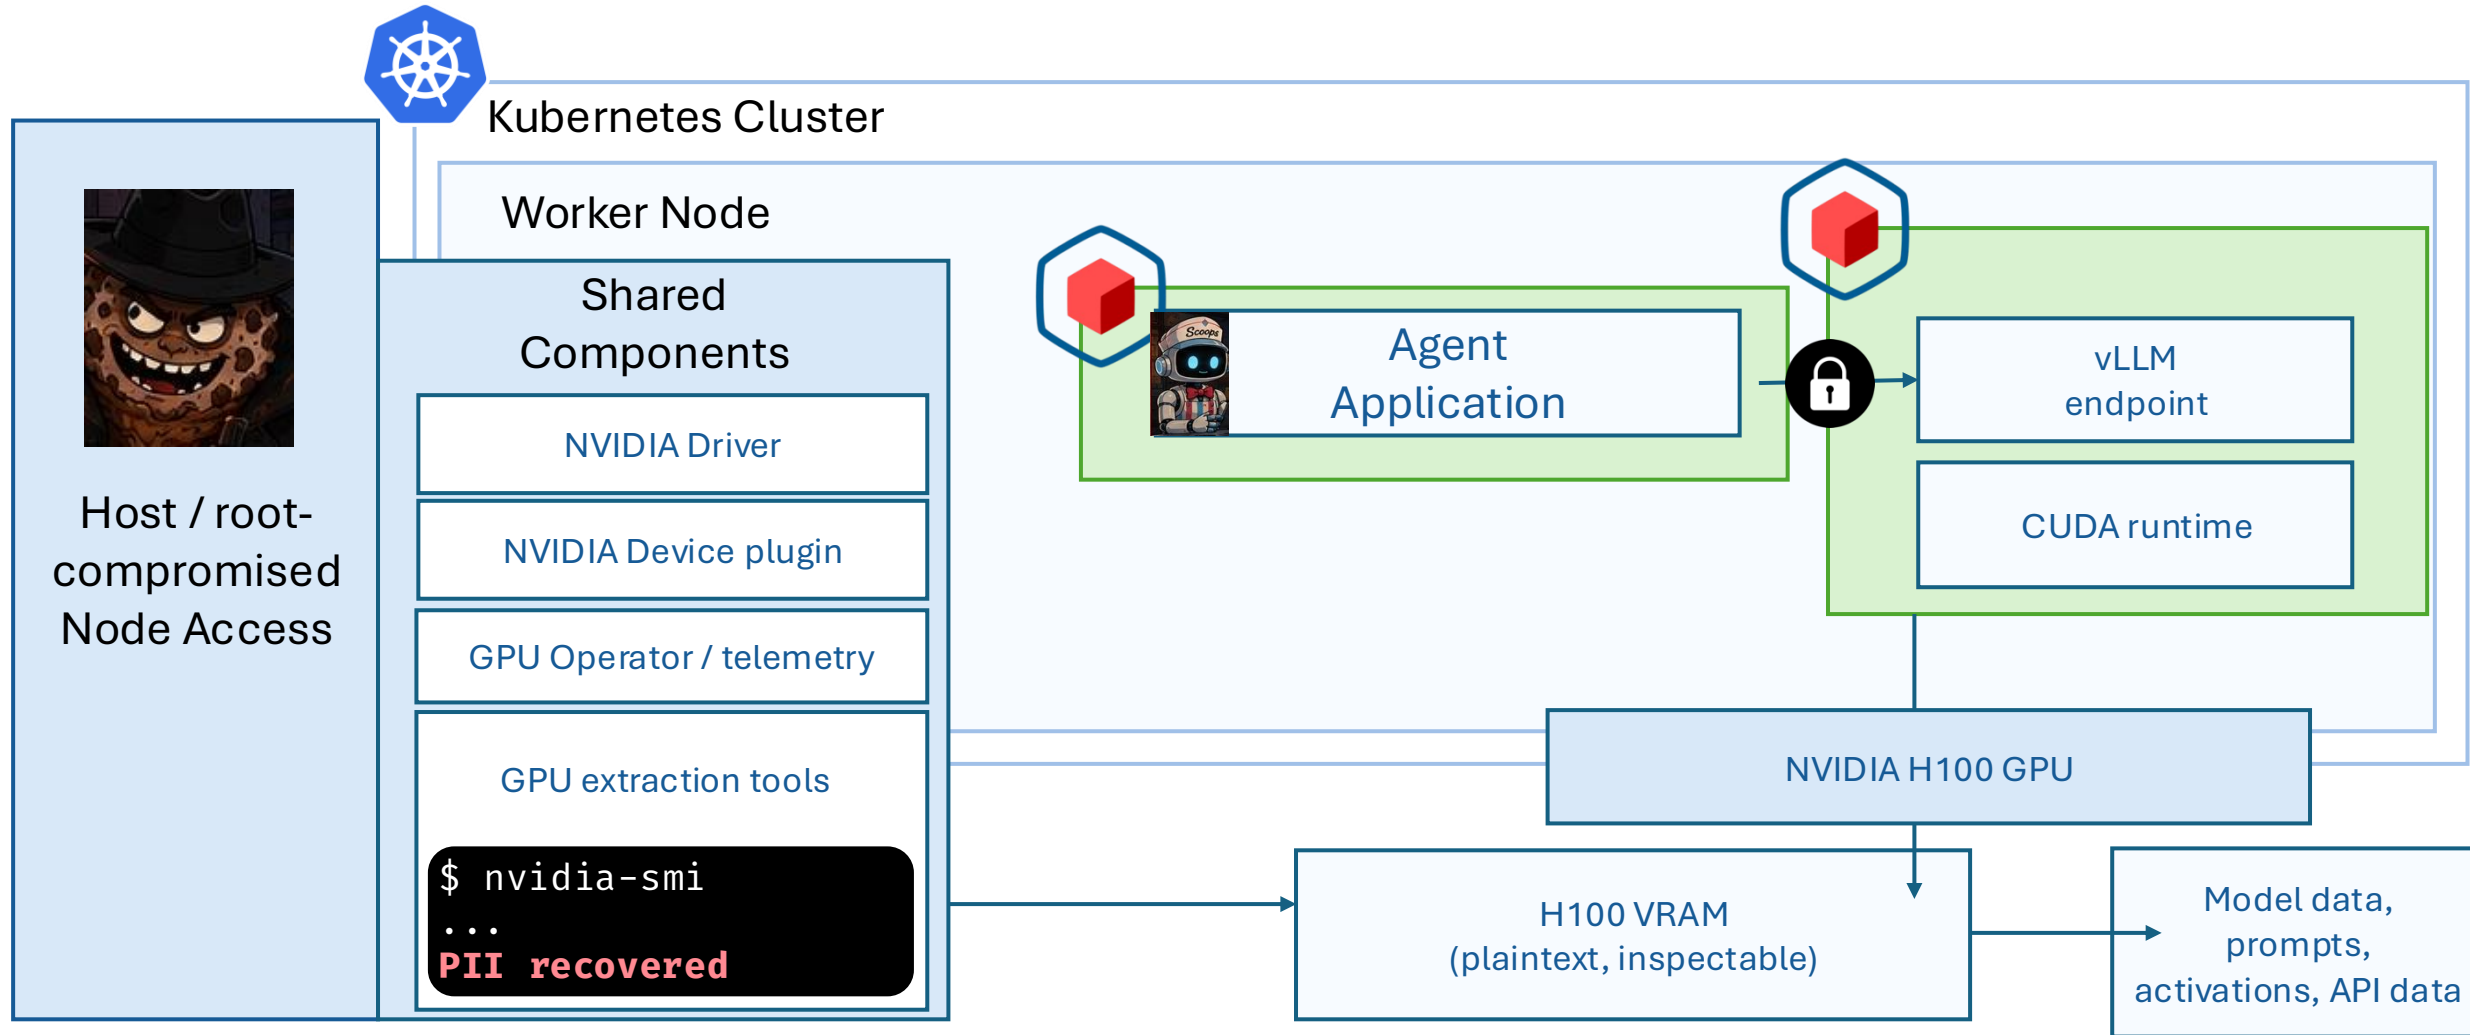
DEFENSE HOLDING

docs/34 complete suite; results/evidence-attestation/18-coco-complete-test-coverage.txt

Data in Transit: Secured



Data in GPU



Data in GPU – Before and After

LIVE HOST-VANTAGE DEMONSTRATION

GPU VRAM without H100 CC

A co-located debugger recovers the canary from device memory

CAPTURED

TERMINAL bash

TIME TO BREACH

00:00.0

EXFILTRATION IN PROGRESS

docs/15_section_6; results/evidence-1gpu-none/r-gpu.txt

LIVE HOST-VANTAGE DEMONSTRATION

GPU VRAM with H100 confidential computing

Host has no driver, no device node, no CUDA process, and no VRAM path

BLOCKED

TERMINAL bash

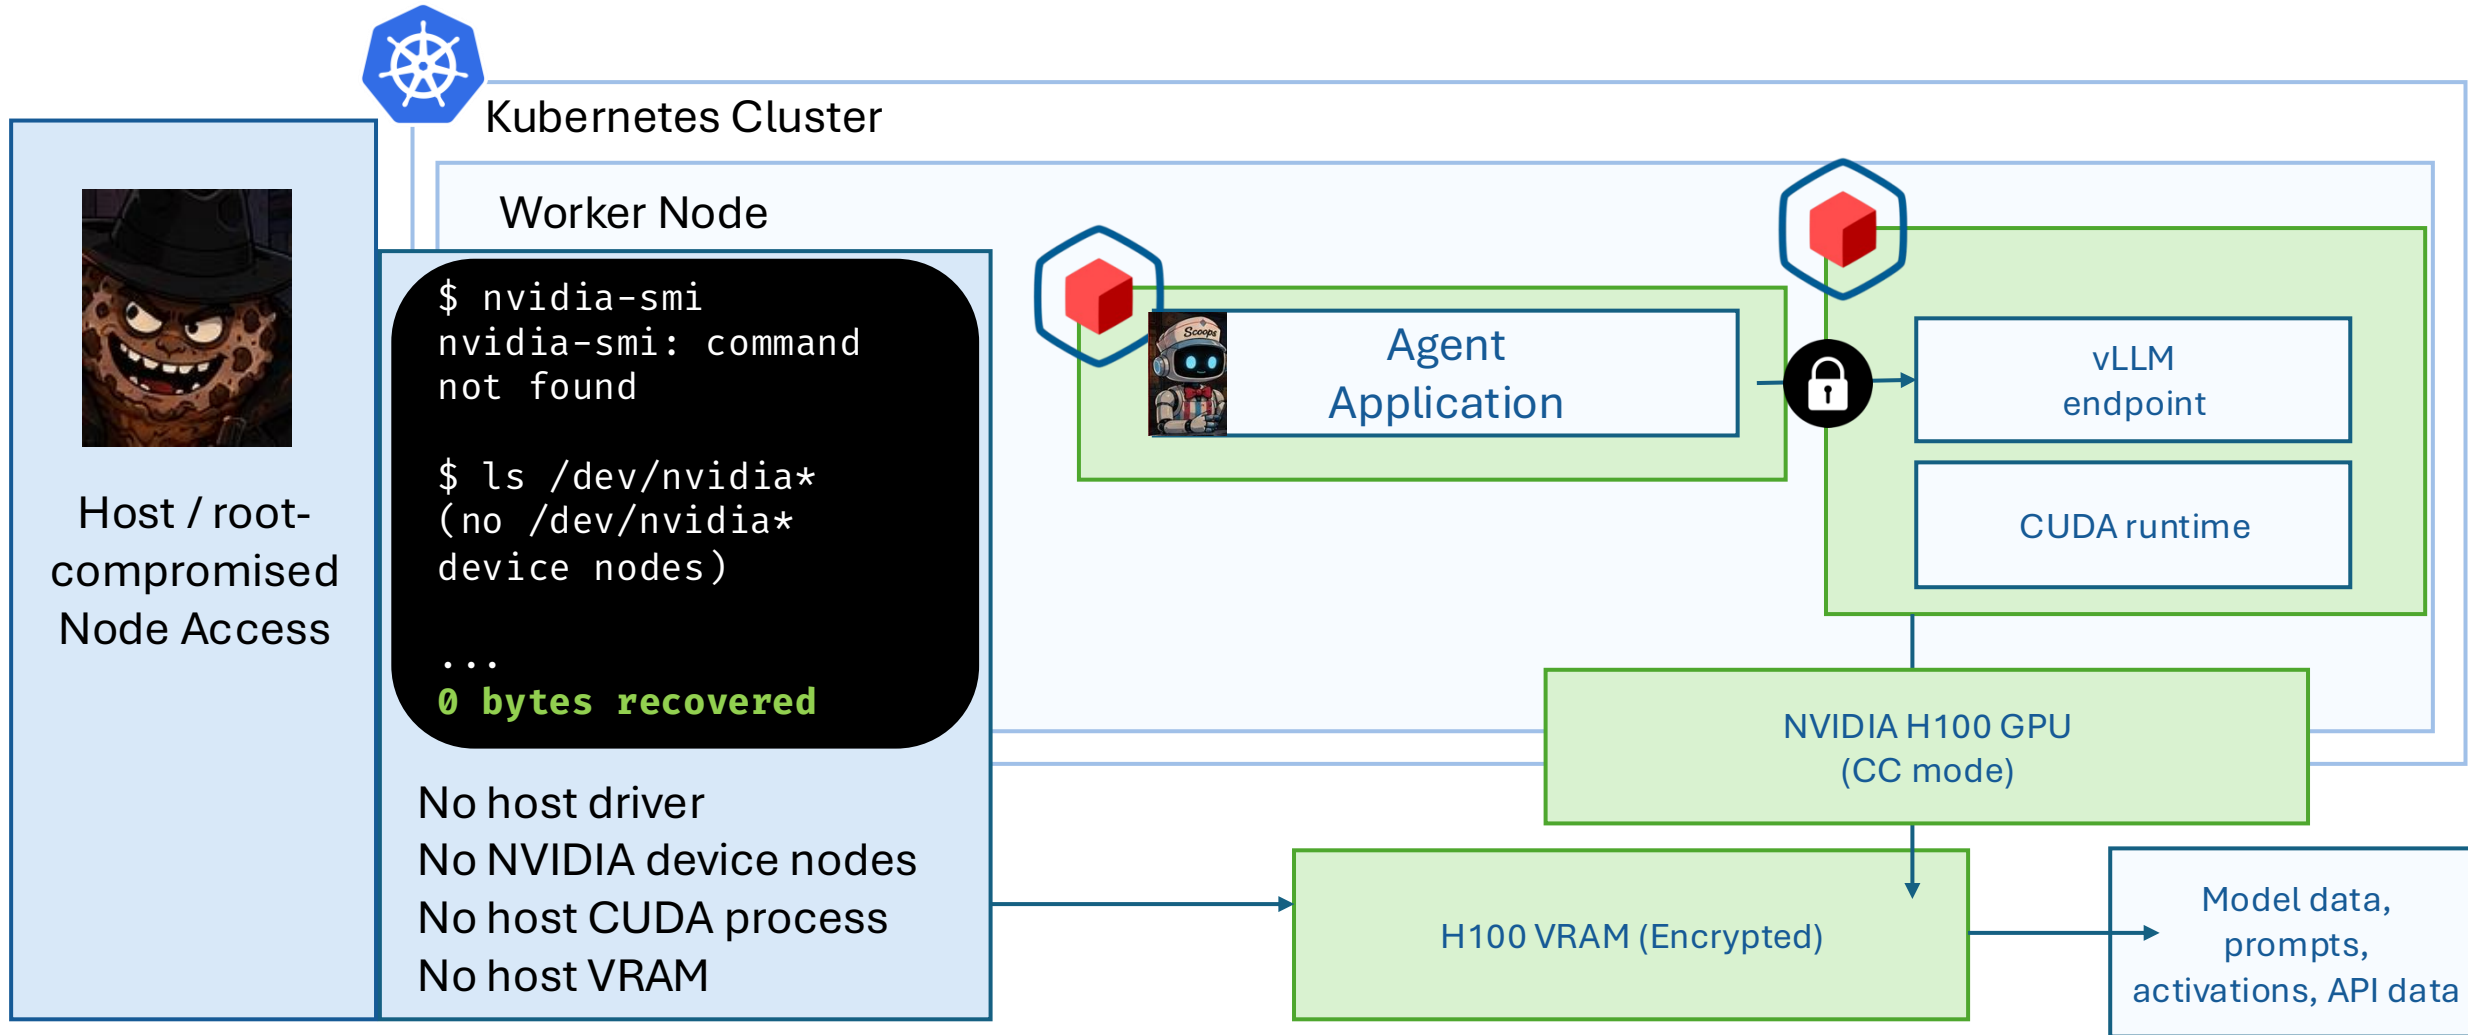
TIME UNDER ATTACK

00:00.0

DEFENSE HOLDING

docs/45; results/evidence-gpucc/gpucc.txt; results/evidence-attestation/16-coco-confidential-gpu.txt

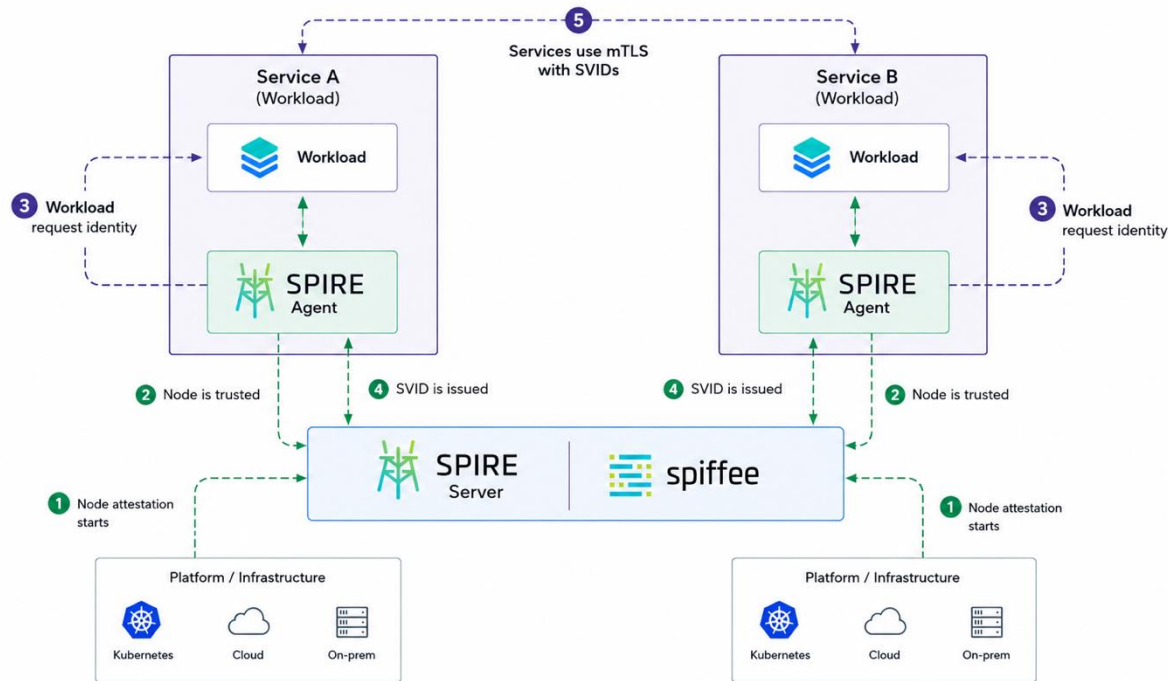

Data in GPU: Secured



Bringing it all together


Attestation of the Workload and Network

Workloads:



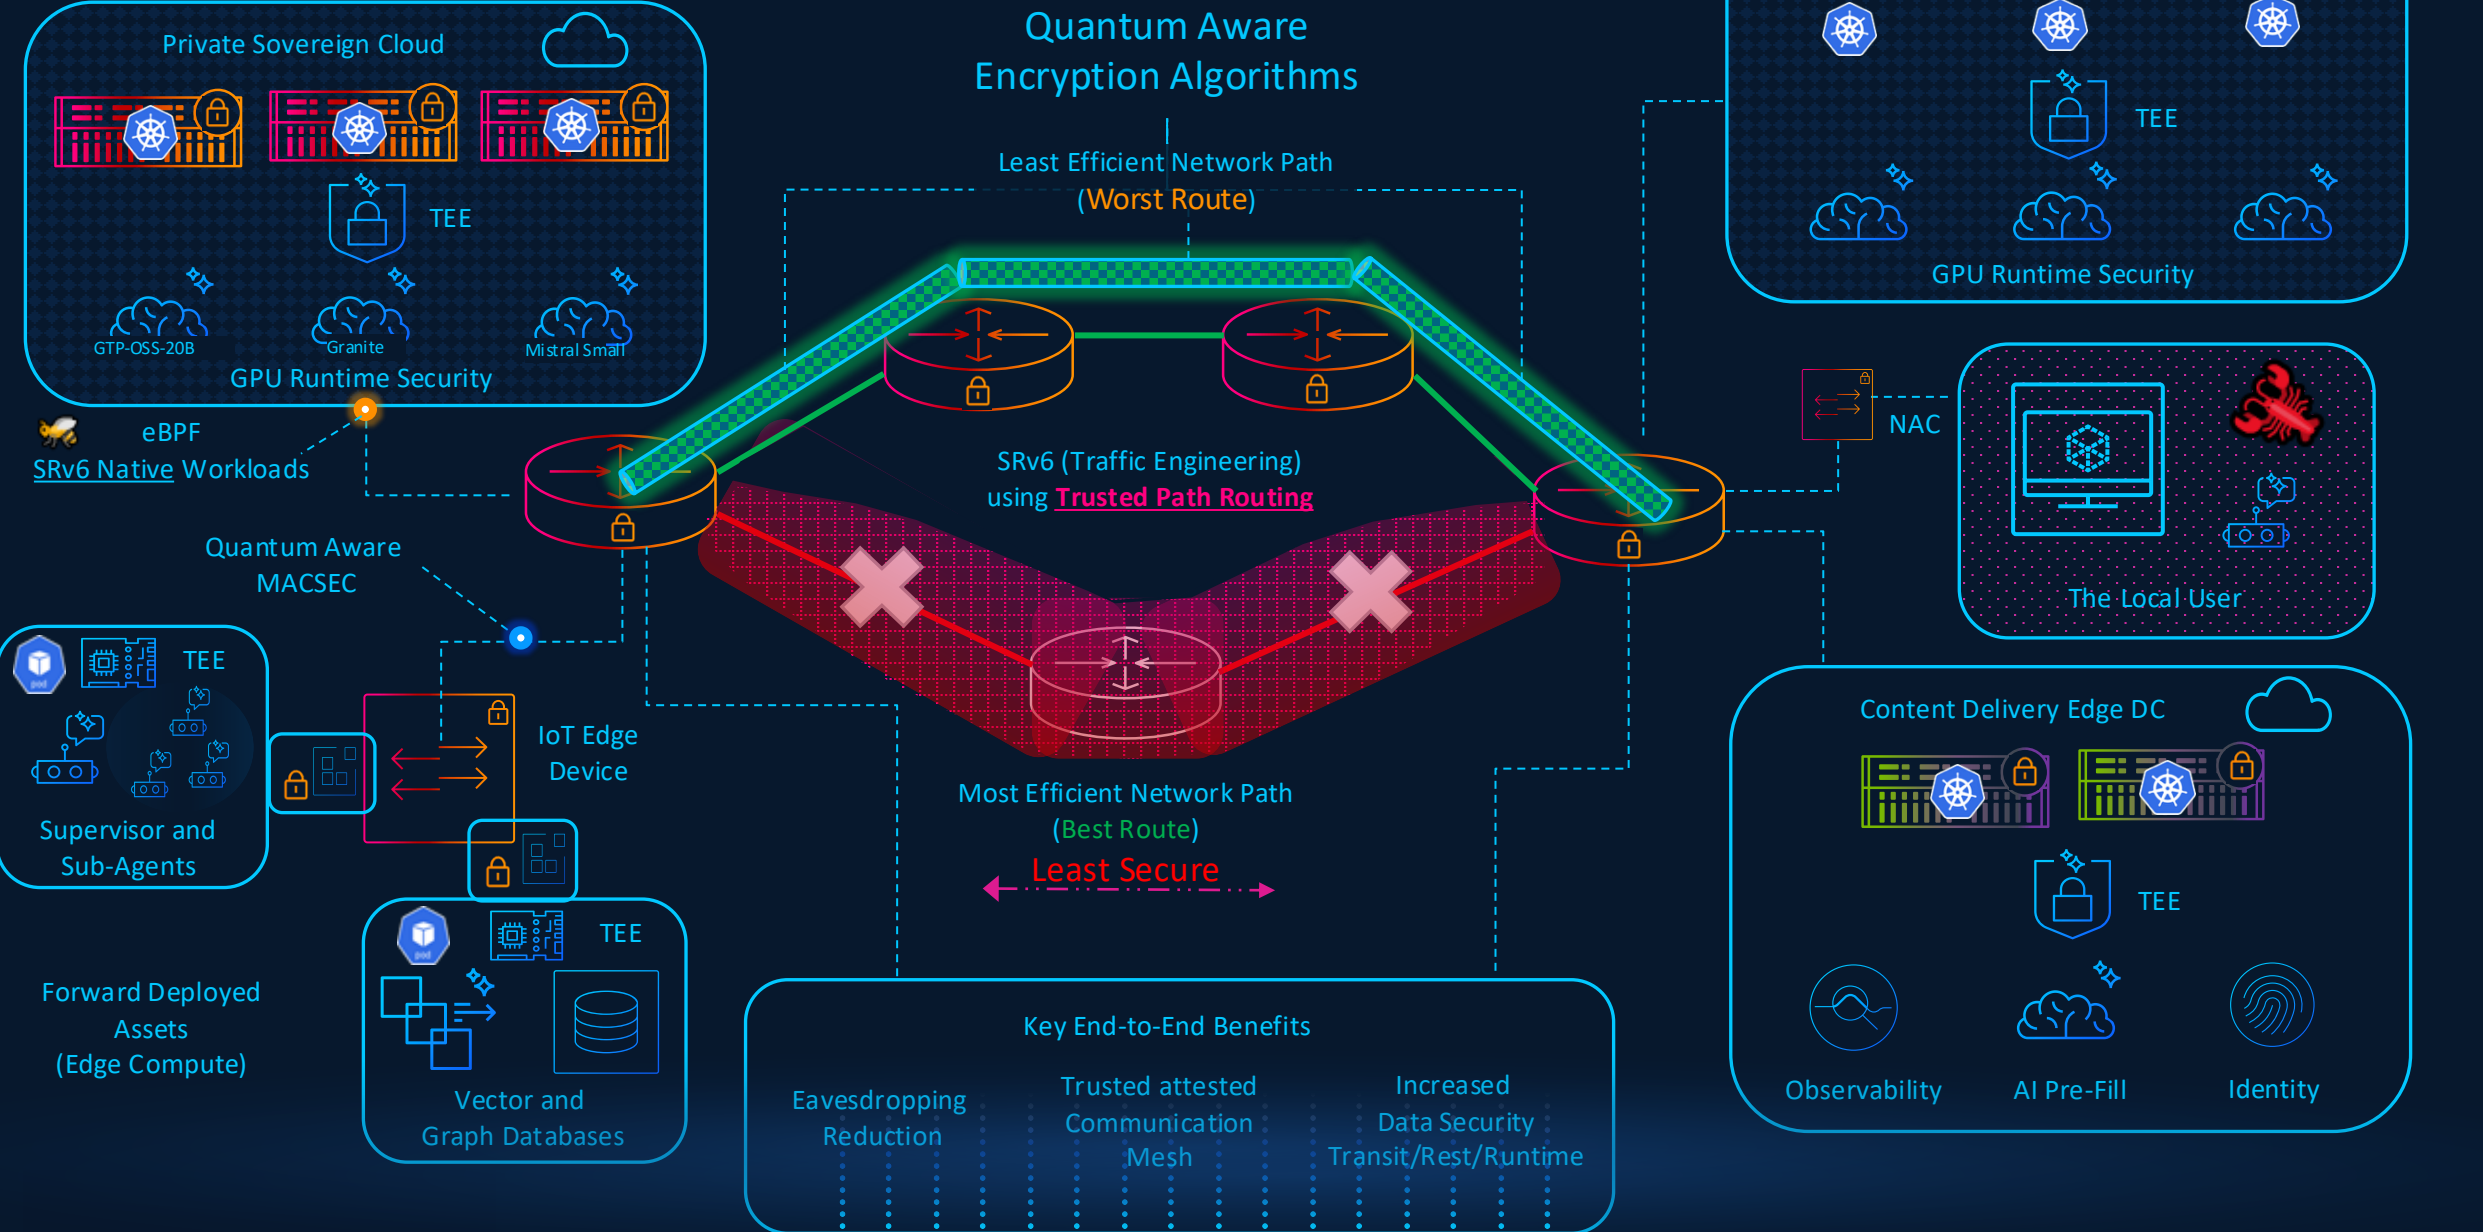
Platform Configuration Registers

Network Device:



PCR Index	Measurement
0	Firmware / BIOS / boot ROM code
1	Firmware / platform configuration data
2	Option ROM / expansion code
3	Option ROM configuration/data
4	Boot loader (MBR / IPL) code
5	Boot loader configuration & GPT
6	Platform state transitions / events
7	Secure Boot policy & keys
8+	OS / kernel measurements (platform-specific)

Achieving Sovereign Security as Scale



Recap

Agentic Security Doesn't Stop Here!

Agent Identity	Governance
Tool Authorization	Explainability
Guardrails	Benchmarking (validation)

What we learnt along the way....

- Zero Trust focuses on removing the weakest link
 - GPU
 - CPU
 - Memory
 - Transit
 - Rest



Thank You!!!