Annex B
Provider Attestation Service Architecture

Informative – Not Normative

Purpose and Scope

This annex addresses the integration of Cloud Service Provider confidential computing services, such as those offered by AWS, Azure, GCP, IBM and OCI, with the framework's evidence layers, trust anchor types (Section 4.1.3), and Provider Trust Stances (Section 5.3). The purpose is to give cloud adopters and provider integrators a shared vocabulary for identifying which framework requirements a given provider service contributes to and which stances accept it. Specific commercial services are not named within the body of this annex. Providers and integrators map their own offerings onto these categories using vendor documentation. As CSP participation in the working group expands, the present categorical treatment may evolve into a named service catalogue in future revisions.

B.1 Hardware Verification Services

Provider-operated services that accept raw hardware attestation reports from confidential VMs and return signed tokens validating the report against silicon-vendor PKI. These services contribute to L1 evidence and produce Type B trust anchors.

Typical capabilities include hardware report authenticity verification against silicon-vendor PKI on the relying party's behalf, TCB version validation against vendor advisories, and signed token output suitable for key release and downstream policy decisions (for example, KMS condition keys or attestation-gated key release).

These services focus on L1 verification. Continuous L2, L3, the Publisher protocol, and Attested Transport Binding are produced by complementary in-TEE components that compose with the provider service. A Workload-Depth deployment under the Provider-Inclusive stance can layer in-TEE L2 measurement on top of a hardware verification service to satisfy framework L2 requirements without replacing the provider's L1 path.

Compatible stances: Provider-Inclusive; Provider-Limited variants that accept the CSP as an operating service.

B.2 Bundled Launch Verification Services

Provider-operated services that combine hardware attestation with image-launch identity binding, typically delivered as a confidential platform-as-a-service offering. These services contribute to L1 and partial L2 (image launch state) and produce Type B trust anchors.

Typical capabilities include hardware attestation as in B.1, workload identity bound to the launched container or VM image digest, and token-based access to provider-managed secrets and external services.

Continuous L2 (post-launch image change detection) and L3 (runtime evidence) are produced by complementary in-TEE components. Implementations seeking framework Full Depth on top of a bundled launch service add framework-conformant L2 runtime monitoring and L3 components alongside the provider service.

Compatible stances: Provider-Inclusive; Provider-Limited variants that accept the CSP as an operating service.

B.3 Vendor-Operated Accelerator Attestation Services

Hardware-vendor-operated services that validate accelerator TEE attestation reports, including GPU TEEs used for AI training and inference workloads. These services contribute to L1-GPU evidence and produce Type C trust anchors.

Typical capabilities include accelerator TEE report validation against vendor PKI, accelerator firmware and driver version validation, and signed token output binding accelerator evidence to a session nonce shared with CPU evidence (Section 4.1.2 shared nonce binding).

Compatible stances: Provider-Inclusive; Provider-Limited variants that explicitly accept the hardware vendor as a service operator. Provider-Independent stances treat vendor-service-mediated L1-GPU evidence as N/A and rely on direct accelerator attestation paths where available.

B.4 Hardware-Direct Paths

Implementations that allow guest workloads to retrieve raw hardware attestation reports directly from the silicon and to access silicon-vendor PKI without intermediation. These produce Type A trust anchors when both the report retrieval and the certificate distribution path are accessible without an interposed operating service.

Typical implementations include guest-accessible interfaces for hardware report retrieval, privileged-VM-level agents that mediate between the silicon and the workload while remaining inside the trust boundary, and customer-operated mirrors of silicon-vendor certificate distribution. The availability of full Type A paths on managed CVM offerings is evolving. Implementers verify both raw report retrieval and silicon-vendor PKI access independently for their target platform.

Compatible stances: Provider-Inclusive, Provider-Limited, and Provider-Independent.

B.5 Customer-Operated Verification

On-premise, sovereign cloud, and customer-operated cloud deployments that perform verification using silicon-vendor PKI directly, without a provider service in the path. These produce Type A trust anchors by construction. The framework's roles, evidence layers, and Publisher protocol apply identically to these deployments.

Compatible stances: Provider-Inclusive, Provider-Limited, and Provider-Independent.

B.6 Composition Summary

The following table summarises which categories above produce which trust anchor types at each framework layer. The summary is illustrative and the availability of any specific path depends on the deployment platform and provider service generation.

Layer	Type A path	Type B path	Type C path
L1 (CPU)	Hardware-direct paths (B.4); customer-operated verification (B.5)	Hardware verification services (B.1); bundled launch verification (B.2)	not applicable
L1-GPU	Vendor roadmap; not currently available on managed offerings	not applicable	Vendor-operated accelerator attestation services (B.3)
L2 (Image)	In-TEE measurement (framework-conformant L2 components)	Partial via bundled launch verification (B.2)	not applicable
L3 (Runtime), Publisher, ATB	Framework-conformant in-TEE components	not currently available	not currently available

B.7 Abbreviations and Acronyms

Term	Definition
ATB	Attested Transport Binding
AWS	Amazon Web Services
CPU	Central Processing Unit
CSP	Cloud Service Provider
CVM	Confidential Virtual Machine
GCP	Google Cloud Platform
GPU	Graphics Processing Unit
IBM	International Business Machines
KMS	Key Management Service
L1	Layer 1 Evidence
L1-GPU	Layer 1 GPU Evidence
L2	Layer 2 Evidence
L3	Layer 3 Evidence
OCI	Oracle Cloud Infrastructure
PKI	Public Key Infrastructure
TCB	Trusted Computing Base
TEE	Trusted Execution Environment
VM	Virtual Machine