CONFIDENTIAL COMPUTING CONSORTIUM

A Board-Ready Framework for Adopting Confidential Computing

June 2026

Confidential Computing is not a security science project; it is a way to turn “we can’t use that data” into “we can safely monetize that data”.

Confidential Computing is emerging as one of the few security investments that can both de-risk sensitive workloads and unlock new revenue streams at the same time. By allowing organizations to run analytics and AI on regulated, high-value, or shared data without exposing it to cloud operators, insiders, or partners during processing, it turns previously “off-limits” data into safely usable fuel for growth. For any board, the question is no longer just “Is our data encrypted?” but “Where could stronger protection while data is in use enable faster cloud migration, new data-sharing businesses, and more defensible compliance positions than our competitors can achieve?”

Setting High Level Understanding

Confidential Computing should be treated as a strategic control for high-value workloads, rather than as a general-purpose replacement for encryption, identity, or application security; it is meant to take the defense in depth strategy approach. The strongest business case appears where the enterprise must process sensitive data in shared infrastructure, collaborate across trust boundaries, or demonstrate stronger technical safeguards for regulated data and machine identities.
Three points are essential:

- **Risk lens:** Data-in-use is the phase where secrets, models, credentials, payment data, and personally identifiable information (PII) may be exposed in plaintext to memory scraping, privileged misuse, or platform compromise unless isolated by hardware-backed controls.
- **Control lens:** CC extends protection to runtime by keeping data protected in memory and releasing secrets only after attestation verifies the expected hardware and software state.
- **Governance lens:** CC reduces some trust assumptions, but it does not eliminate residual risks such as weak key management, software vulnerabilities inside the workload, supply-chain risk, or attestation and implementation gaps.

Board framework

A board-ready way to evaluate CC is through five questions:

Board question	Why it matters	Evidence to request
Which data-in-use scenarios create material business risk?	The value of CC is highest where sensitive data must be processed in shared or less trusted environments.	Heat map of workloads using regulated data, cryptographic keys, AI models, payments data, or cross-organization analytics.
What trust assumption is being removed?	CC is most useful when the enterprise wants to reduce dependence on cloud operators, hypervisors, or other privileged layers.	Architecture showing which administrators, providers, or counterparties can no longer access plaintext data.
Is attestation operationalized? Actually, what is attestation?	Attestation is verifiable proof that the trusted execution environment has remained intact. It is central because it lets workloads verify TEE integrity before secrets are released.	Policy describing attestation checks, secret release conditions, and audit logging.
What residual risks remain?	CC narrows exposure but does not remove application bugs, side channels, weak governance, or dependency risk.	Residual risk register, compensating controls, and patch-management metrics for enclave-enabled workloads.
What is the economic payoff?	CC can enable cloud migration, regulated collaboration, and lower breach impact, but costs rise with workload redesign, tooling, and operational complexity.	Business case with avoided-loss assumptions, implementation cost, and effect on time to launch or partner data-sharing opportunities.

Regulatory implications

From a regulatory perspective, NIST frames the problem directly: machine identities and other secrets remain at risk when sensitive *data in use* (think of data in use as, “when I send a question to a chatbot”) is not protected, and Confidential Computing can serve as a practical template for reducing that exposure.

For boards, the implication is not that regulators mandate a single technology, but that stronger runtime protection (aka, data in use) can improve the defensibility of control design in audits, examinations, and incident response reviews. This is especially relevant in highly regulated industries.

Financial implications

The financial case for CC usually rests on four levers:

- Reduced expected loss from breach or insider misuse, because plaintext exposure during processing is narrowed.
- Faster approval of sensitive cloud or partner-hosted workloads, because technical controls can remove some trust barriers.
- New data-sharing and analytics models, including fraud detection and cross-organization collaboration that would otherwise be too risky.
- Higher near-term cost from re-architecture, vendor selection, attestation operations, and specialized skills.

A practical threshold question is whether CC changes a business decision: if it does not unlock a blocked workload, reduce a quantified exposure, or create a new revenue or collaboration path, it may remain an engineering improvement rather than a board-level investment theme.

Competitive implications

CC can become a competitive differentiator when customers, regulators, or partners care about where and by whom sensitive data can be accessed during processing. Providers now position confidential VMs, containers, and ledgers as ways to support sovereign controls, selective data sharing, and auditable protection against operator access, which means buyers should expect the market to compete increasingly on attestation quality, ecosystem maturity, and key-control models rather than on basic encryption claims alone.

This dynamic also raises a strategic sourcing issue for boards: proprietary attestation or key-management dependencies can create concentration risk, while ANSSI's publication cadence and public stance on Confidential Computing indicate growing European attention to trust models in cloud environments. The competitive question is therefore not only "Do we have CC?" but "Can we prove trustworthiness across clouds, partners, and regulators without becoming locked into a single control plane?"

Suggested Board actions (span 12-month period)

1. Direct company leadership to identify the top 10 workloads where data-in-use exposure could create material regulatory, financial, or franchise harm across each line of business.
2. Require a control map showing where CC would remove trust in administrators, providers, or counterparties, and where it would not.
3. Ask for an attestation/verifiable operating model, including secret release policy, evidence retention, and exception handling.
4. Review residual risks explicitly, including software flaws inside enclaves, key custody, and vendor concentration.
5. Approve investment only where management can tie CC to measurable risk reduction, faster regulated deployment, or a defensible collaboration advantage.

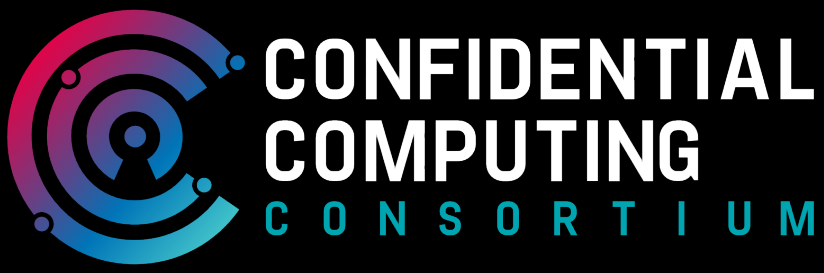
CONFIDENTIAL
COMPUTING
CONSORTIUM