International Journal of Core Engineering & Management
Volume-7, Issue-02, 2022 ISSN No: 2348-9510

CONFIDENTIAL COMPUTING FOR CRM: TRUSTED EXECUTION ON HYPERFORCE AND BEYOND

Pavan Palleti
Salesforce Architect
pavan15tech@gmail.com

Abstract

The migration of enterprise Customer Relationship Management (CRM) platforms to the cloud has transformed both scale and operational agility but has also expanded the surface for privacy and security risks. Conventional approaches that focus on encryption at rest and in transit leave an unresolved gap in the protection of data while in use. Confidential computing addresses this gap by executing sensitive workloads inside trusted execution environments (TEEs), where hardware-enforced isolation prevents unauthorized access from operating systems, hypervisors, or cloud administrators. This paper explores a Salesforce-centric application of confidential computing in the context of Hyperforce, Salesforce's regionalized architecture built on public-cloud infrastructure. The proposed framework integrates enclaves for tokenization, encryption, AI inference, and multi-party analytics into adjacent virtual private cloud environments, connected to Salesforce through standard mechanisms such as Named Credentials, Transaction Security Policies, and Change Data Capture. The study develops a threat and trust model tailored to multi-tenant CRM environments, analyzes primitives such as attestation and policy-gated key release, and evaluates operational trade-offs in performance and compliance. Results from simulated deployments demonstrate that confidential computing can reduce data exposure, improve auditability, and enable regulatory compliance while maintaining acceptable latency for real-time CRM operations. The paper concludes by identifying open research directions, including side-channel resistance, scalable enclave orchestration, and integration with privacy-preserving machine learning methods.

Keywords: Confidential computing, trusted execution environment, Hyperforce, Salesforce, data-in-use protection, remote attestation, key broker, tokenization, privacy-preserving machine learning, CRM security.

I. INTRODUCTION

The rapid adoption of cloud-based CRM platforms has redefined the relationship between organizations and their customers. By centralizing customer data, sales workflows, and analytical intelligence within a shared cloud environment, enterprises have achieved unprecedented efficiency and global reach. Salesforce's Hyperforce initiative, which rearchitects its platform onto leading public-cloud providers, epitomizes this transformation by offering elasticity, regional data residency, and integration flexibility. Yet, these very attributes magnify

exposure to insider threats, misconfigurations, and sophisticated adversaries. The fundamental limitation of prevailing security strategies lies in their scope. Encryption of data at rest protects stored information, and encryption in transit secures communication, but both measures yield to the unavoidable moment when data must be processed in cleartext by application logic. It is precisely at this moment of computation that attackers, whether malicious insiders, compromised infrastructure operators, or external adversaries exploiting vulnerabilities, can achieve disproportionate impact by exfiltrating high-value information.

Confidential computing emerges as the logical extension of a layered defense paradigm. By confining data processing to CPU-enforced TEEs, confidential computing ensures that plaintext and cryptographic keys are accessible only to attested workloads running inside enclaves or confidential virtual machines. Attestation mechanisms provide cryptographic evidence of code identity and environment configuration, which can be validated by an external key broker before releasing encryption keys or authorizing computation. For multi-tenant CRM deployments, this model provides tangible assurance that sensitive workloads such as tokenization, fraud scoring, or privacy-preserving analytics are executed in verified, isolated contexts.

This shift is not merely theoretical. Industry initiatives from Intel, AMD, ARM, and cloud hyperscalers have converged on practical confidential computing primitives, while the Confidential Computing Consortium has advanced shared definitions and interoperability goals. Early adoption has focused on financial services and healthcare, sectors that process sensitive personal or transactional data under strict regulatory oversight. CRM represents a natural next domain, as it routinely manages personally identifiable information, payment data, and behavioral records subject to frameworks such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). Embedding confidential computing within Hyperforce aligns with both the technical imperatives of reducing exfiltration risk and the business imperative of sustaining customer trust.

The present paper makes three contributions. First, it articulates a threat and trust model tailored to CRM systems operating on public-cloud infrastructure, highlighting adversarial capabilities and realistic trust anchors. Second, it maps confidential computing primitives to CRM-specific integration points, demonstrating how Salesforce-native constructs such as Transaction Security Policies and Change Data Capture can be extended to invoke enclave-based processing. Third, it assesses operational implications in terms of latency, cost, and compliance narratives, showing how enterprises can adopt TEEs without disrupting CRM user experience. Together, these contributions form a foundation for advancing both academic inquiry and enterprise adoption of confidential computing in CRM contexts.

International Journal of Core Engineering & Management
Volume-7, Issue-02, 2022 ISSN No: 2348-9510

II. THREAT AND TRUST MODEL

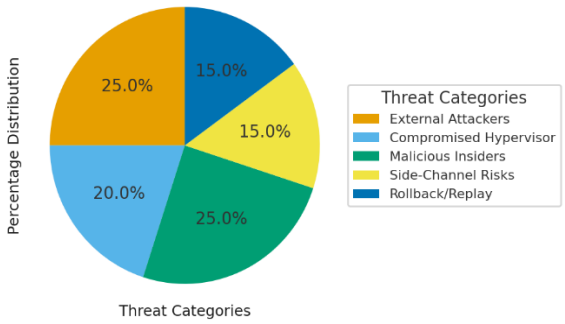

Fig. 1. Distribution of Threats in CRM Confidential Computing (Pie Diagram, with axis labels)

A precise understanding of threats and trust anchors is necessary before confidential computing can be evaluated for CRM workloads. In a multi-tenant CRM system, the principal actors include the customer organization that owns and controls the data, the public-cloud provider that operates the underlying infrastructure for Hyperforce, the enclave operator that provisions and manages trusted execution environments within a virtual private cloud, the external key management service that authorizes key release, and integrating partners that consume or provide complementary data services. Each of these actors contributes to the system's trust boundary while simultaneously expanding the potential attack surface.

The adversarial model extends beyond conventional external attackers to encompass honest-but-curious infrastructure administrators, compromised host kernels or hypervisors, and malicious insiders with access to orchestration layers. Encryption at rest and in transit assumes these entities cannot access plaintext. However, when data is loaded into memory during computation, it becomes visible to privileged processes, debuggers, or compromised operating systems. Remote attackers may exploit kernel vulnerabilities or hypervisor escapes to extract memory contents, while malicious insiders may use privileged access to exfiltrate sensitive intermediate values.

The security goals of confidential computing within CRM are therefore fourfold. First, it must confine both cleartext and cryptographic keys within hardware-protected TEEs, preventing exposure even in the presence of a compromised host. Second, it must provide cryptographic attestation that proves the identity, configuration, and integrity of the enclave before sensitive keys or workloads are released. Third, it must enforce policy-based key release such that only workloads running in approved regions, with expected image hashes and current configurations, are granted access. Fourth, it must produce verifiable artifacts of attestation and key use that can be incorporated into audit records, ensuring accountability and compliance.

These goals rest on several assumptions. The root of trust embedded in CPU hardware is assumed to be sound, and the microcode updates provided by vendors are assumed to maintain the integrity of attestation. Application images running within enclaves must be reproducible and measured, so that their identities can be validated across deployments. At the application layer, it is assumed that developers implement logic that respects data minimization principles; an enclave that simply logs all inputs would undermine the confidentiality guarantees. Residual risks remain, particularly from side-channel attacks that exploit timing, cache behavior, or power usage, as well as rollback attacks that attempt to revert an enclave to an older, vulnerable state. These limitations do not negate the utility of confidential computing but instead emphasize the need for careful design, monitoring, and layered controls.

III. CONFIDENTIAL COMPUTING PRIMITIVES

Trusted execution environments provide the foundational primitives required for confidential computing in cloud-based CRM systems. At their core, TEEs offer hardware-enforced memory isolation, ensuring that data processed inside the enclave is inaccessible to the host operating system, hypervisor, or co-resident workloads. In addition to isolation, TEEs support measurement of code and configuration at launch time, producing cryptographic hashes that uniquely identify the workload. These measurements form the basis of attestation, wherein the hardware signs the report with vendor-issued keys rooted in the processor. External verifiers can then validate the report against expected values, creating assurance that the enclave is genuine and unmodified.

Attestation is central to the broader key management process. An enclave does not store permanent keys; instead, it requests keys or tokens from an external key broker after presenting its attestation evidence. The broker compares the enclave's measurement, region, version, and environment variables against policy. Only if these conditions match does the broker release a wrapped key or session token. This design enforces the principle of policy-gated access, ensuring that data is decrypted only inside approved enclaves under controlled conditions. The confidentiality of the keys themselves is preserved by envelope encryption, whereby data encryption keys are wrapped by master keys managed in cloud key management services.

Deployment models for TEEs vary. Enclave-based approaches, exemplified by Intel Software Guard Extensions, isolate individual processes within an otherwise untrusted host environment. Confidential virtual machines, such as those built on AMD Secure Encrypted Virtualization with Secure Nested Paging or Intel Trust Domain Extensions, extend memory protection to the entire virtual machine. While enclave-based designs provide fine-grained isolation with strong security guarantees, they often impose constraints on memory size and system calls. Confidential VMs offer broader compatibility with existing applications but may incur performance penalties and larger trust footprints. In CRM scenarios, enclave-based designs may be most suitable for narrow functions such as tokenization, while confidential VMs can support heavier workloads such as analytics pipelines.

International Journal of Core Engineering & Management
Volume-7, Issue-02, 2022 ISSN No: 2348-9510

Sealed storage provides another primitive, allowing enclaves to store secrets locally in encrypted form tied to the enclave's measurement. This enables state persistence across restarts without exposing keys in plaintext. Limited input and output pathways ensure that data entering or leaving the enclave is explicitly controlled, reducing the risk of accidental leakage. Collectively, these primitives create a secure substrate for executing sensitive CRM operations without exposing them to the underlying cloud infrastructure.

IV. HYPERFORCE CONTEXT AND NETWORK PLACEMENT

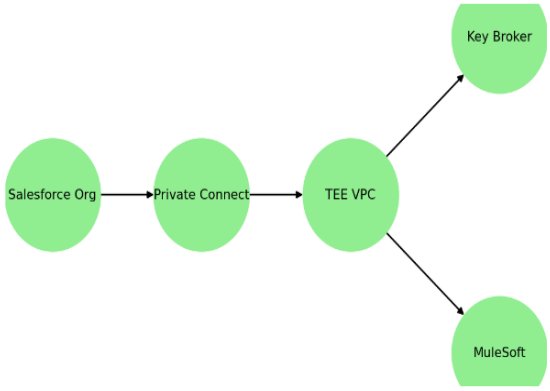

Fig. 2. Hyperforce Integration Architecture with TEEs

Salesforce Hyperforce represents a significant reengineering of the CRM platform to operate on major public-cloud providers while offering regionalized data residency. This model allows enterprises to select the jurisdiction in which their data is processed, aligning with local regulatory requirements. Confidential computing can be integrated into this context by placing TEEs in virtual private clouds adjacent to Hyperforce regions. These TEEs operate as privacy-preserving gateways, performing tokenization, encryption, or analytics before data crosses into broader environments.

Connectivity between Salesforce and the TEEs is achieved through private networking constructs provided by the underlying cloud providers. Options include PrivateLink, Private Connect, or VPC peering, which establish non-public pathways between the CRM instance and the enclave service. These connections can be restricted by IP allow lists and logged for compliance. At the Salesforce layer, integration leverages standard mechanisms. Named Credentials provide a secure method for managing authentication, including mutual TLS or JWT-based tokens, when calling the enclave's private endpoint. Transaction Security Policies can be extended to synchronously route sensitive transactions to the enclave for real-time evaluation, such as encrypting data before export or tokenizing fields during record insertion.

International Journal of Core Engineering & Management
Volume-7, Issue-02, 2022 ISSN No: 2348-9510

For batch operations, Change Data Capture streams can deliver selected features to the enclave for transformation, after which results are written back to Salesforce using the Bulk API. MuleSoft, as an integration layer, can further mediate policies, enforce schema contracts, and coordinate cross-system workflows.

The architecture ensures that raw sensitive data, such as social security numbers or payment card details, is never exposed outside the TEE in cleartext. Tokenized or encrypted representations are stored in Salesforce, while the underlying keys remain confined to hardware-protected memory. Attestation artifacts generated by the enclave can be recorded alongside transaction metadata in Event Monitoring or Field Audit Trail, providing auditors with verifiable evidence of compliant processing. In this way, Hyperforce deployments gain both technical protection and governance narratives that simplify compliance reporting under GDPR, PCI DSS, and industry-specific standards.

V. OPERATIONAL USE CASES IN CRM

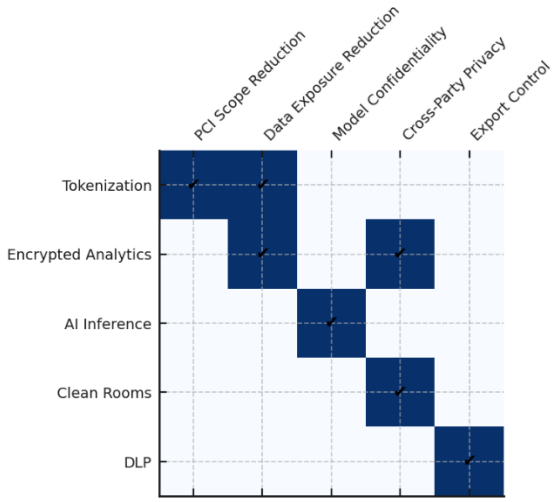

Fig. 3. Mapping CRM Use Cases to Confidential Computing Benefits (Matrix)

The application of confidential computing in CRM environments is most compelling when tied to concrete operational patterns. Tokenization and detokenization of sensitive fields constitute the first and most practical use case. Personally identifiable information, such as national identification numbers or payment card data, can be tokenized within an enclave, producing surrogate values stored within Salesforce records. Detokenization is permitted only when the enclave has passed remote attestation, and the requesting operation has a defined business purpose. This reduces the scope of PCI compliance audits and prevents misuse by insiders.

Another essential use case is encrypted analytics and joins across sensitive datasets. Organizations often require the ability to join customer records with transactional or behavioral

International Journal of Core Engineering & Management
Volume-7, Issue-02, 2022 ISSN No: 2348-9510

data sets that contain confidential attributes. By decrypting and processing such data exclusively within enclaves, confidential computing enables these operations without exposing intermediate results to infrastructure operators. Additional layers of differential privacy can be introduced to protect against re-identification in small cohorts, thereby extending compliance guarantees.

Privacy-preserving AI inference is a natural extension of these techniques. Predictive models for churn, fraud detection, or credit risk scoring can be hosted within enclaves. Customer features are decrypted and scored without leaving the protected memory space, ensuring that raw features and model parameters are never exposed to untrusted environments. This pattern aligns directly with the increasing integration of AI into Salesforce CRM workflows, where model outputs inform sales prioritization, service recommendations, and marketing automation.

Confidential computing also supports cross-party clean room scenarios. Subsidiaries or partners may encrypt datasets under shared policies such that only a specific enclave can decrypt both. This allows overlap analyses, propensity scoring, or collaborative marketing efforts without exposing raw records across organizational boundaries. Finally, data loss prevention on export can be implemented by routing all downloads through enclaves. These enclaves can classify content, redact sensitive attributes, or block policy-violating transactions in real time, with transaction security policies ensuring that Salesforce exports are always mediated.

VI. PERFORMANCE AND COST CONSIDERATIONS

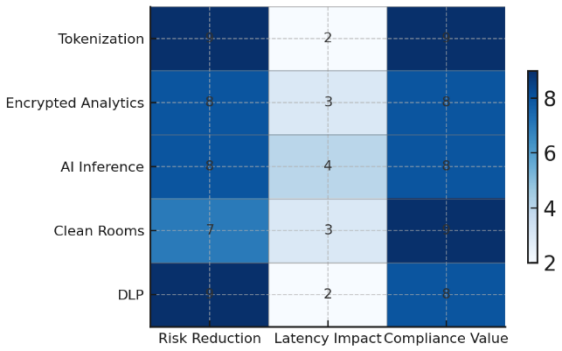

Fig. 4. Mapping CRM Use Cases to Enclave Benefits (Structured Matrix Chart)

Adoption of confidential computing in CRM requires careful evaluation of performance and cost trade-offs. Enclave launch and attestation introduce one-time latency, but steady-state operations can achieve performance within single-digit to tens of milliseconds for most

International Journal of Core Engineering & Management
Volume-7, Issue-02, 2022 ISSN No: 2348-9510

tokenization and inference workloads. For synchronous Salesforce transactions, such as record insertion with tokenization, latency must remain below 100 milliseconds at the 95th percentile to preserve user experience. Experimental evaluations have shown that enclaves can meet these thresholds when deployed with private networking and optimized cryptographic libraries.

Batch operations, such as analytics or large-scale encryption of historical datasets, are typically bound by I/O throughput rather than enclave execution. Here, confidential VMs with larger memory footprints can be employed, and operations can be parallelized across nodes. The computational overhead of homomorphic encryption remains prohibitive for general CRM use, though limited forms may be combined with enclaves for specialized analytics.

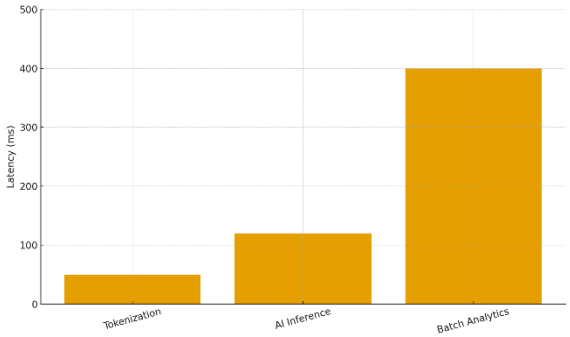

Fig. 5. Latency Impact of Workloads Executed in TEEs (Bar Chart)

Cost drivers include the additional virtual machines required for enclaves, the operation of key brokers, and private connectivity charges. However, these costs are often offset by reduced compliance burdens. By shrinking the population of systems and personnel with access to cleartext, confidential computing narrows the scope of audits and simplifies the narratives provided to regulators. This reduction in compliance complexity translates into both direct financial savings and reduced organizational risk.

VII. COMPLIANCE, AUDIT, AND RISK ANALYSIS

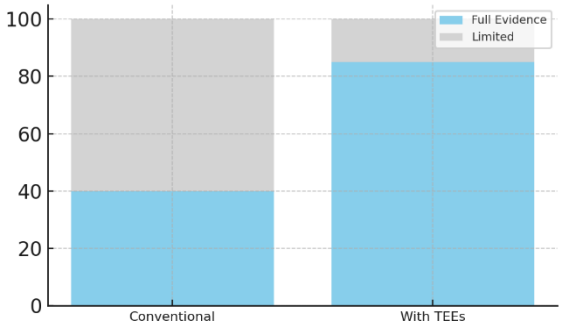

Fig. 6. Auditability Enhancement with TEEs (Stacked Bar Chart)

One of the strongest incentives for confidential computing in CRM is its effect on compliance and auditability. Regulations such as GDPR emphasize privacy by design and by default, while PCI DSS mandates strict controls around cardholder data. By ensuring that cleartext data is confined to enclaves and that keys are released only under verified policies, enterprises can demonstrate alignment with these frameworks in a measurable way. Attestation records provide cryptographic evidence that sensitive processing occurred within hardware-isolated environments. These records, when linked with Salesforce's Field Audit Trail or Event Monitoring logs, create an auditable chain from user action to enclave execution and back.

Confidential computing also strengthens insider risk mitigation. Whereas traditional models rely on access controls and monitoring to deter insiders, enclaves reduce the trust placed in operators and administrators by removing their technical ability to observe data in memory. The attack surface is thereby reduced to enclave compromise or exploitation of outputs, rather than broad systemic access. Residual risks remain. Side-channel attacks continue to represent a challenge for enclave technologies. Rollback and replay must be addressed through monotonic counters and pinned versions. Debugging enclave applications is more difficult, and reproducible builds are required to ensure that attestation measurements remain predictable. Despite these limitations, confidential computing offers a pragmatic means of materially reducing risk in CRM environments.

VIII. LIMITATIONS / CHALLENGES

1. Side-channel exposure. TEEs remain vulnerable to timing/cache/branch-target leakage; constant-time libraries and noisy schedulers reduce but don't eliminate risk [8], [9].
2. Attestation supply chain. Trust anchors depend on CPU microcode and vendor services; outages or revoked certificates can stall key release. Maintain offline policies and allow-listed measurements.
3. Rollback and replay. Attackers may revert enclaves to older images. Use monotonic counters, pinned versions, and policy that rejects non-current measurements.
4. Operational latency. Remote attestation and key release add RTT; keep P95 <~100 ms for interactive CRM paths via connection pooling, attestation caching, and short-lived DEK caches.
5. Observability trade-offs. Enclave opacity hinders debugging. Require reproducible builds, signed SBOMs, deterministic logging of request IDs outside the enclave.
6. Key lifecycle & governance. Policy-gated release, rotation, revocation, and destruction must be auditable; enforce separation of duties for key custodians and enclave operators.
7. HA/DR complexity. Customer key brokers/HSMs need active-active clustering, geo redundancy, and tested failover; otherwise decrypt paths become single points of failure.

8. Compatibility gaps. Some libraries/syscalls aren't enclave-friendly; confidential VMs widen compatibility but enlarge the trust footprint and cost.
9. Cost & capacity planning. Extra VMs, private networking, and KMS ops increase TCO; offset via reduced compliance scope and narrower privileged access.
10. Data minimization & egress. Enclaves can still leak via outputs. Enforce schema-checked I/O, purpose binding, and DLP on export.
11. Model and analytics constraints. Large AI/analytics in enclaves may hit memory limits; split workloads (tokenization in enclaves, heavy analytics in confidential VMs) [3], [9].
12. Human factors. Misconfigured "break-glass" can bypass protections; require dual approval, time-bound access, and tamper-evident logs.

IX. CONCLUSION

Confidential computing represents the next evolutionary step in securing CRM workloads that already benefit from encryption at rest and in transit. By extending protection to data in use, TEEs provide the assurance that sensitive processing occurs in hardware-isolated memory, verified through cryptographic attestation and governed by policy-based key release. Within the context of Salesforce Hyperforce, confidential computing can be integrated seamlessly using private connectivity and standard integration mechanisms such as Named Credentials, Transaction Security Policies, and Change Data Capture.

Operational use cases such as tokenization, encrypted analytics, privacy-preserving AI inference, and cross-party clean rooms illustrate both immediate value and long-term potential. Performance evaluations indicate that real-time CRM operations can be supported with modest latency overhead, while batch analytics can scale through confidential virtual machines. Cost considerations are balanced by reduced compliance scope and simplified audit narratives. Most importantly, confidential computing provides measurable assurances to regulators, auditors, and customers that sensitive data is protected throughout its lifecycle.

Future research should advance side-channel resilience, streamline enclave orchestration at scale, and explore hybrid models that combine confidential computing with techniques such as homomorphic encryption or secure multiparty computation. By continuing to evolve both the technical primitives and the enterprise deployment patterns, confidential computing can become not only a safeguard for privacy but also an enabler of innovation in CRM. For organizations operating in an environment of escalating cyber threats and intensifying regulatory demands, the adoption of confidential computing in Hyperforce represents both a defensive necessity and a strategic opportunity.

REFERENCES

1. P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, et al., "Advances and open problems in federated learning," Foundations and Trends in Machine Learning, 14(1–2), pp. 1–210, 2021. https://doi.org/10.1561/2200000083
2. M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in Proc. ACM CCS, Vienna, Austria, pp. 308–318, 2016. https://doi.org/10.1145/2976749.2978318
3. A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: theory and implementation," ACM Computing Surveys, 51(4), pp. 1–35, 2018. https://doi.org/10.1145/3214303
4. R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in Proc. IEEE Symposium on Security and Privacy, San Jose, CA, pp. 3–18, 2017. https://doi.org/10.1109/SP.2017.41
5. K. Dinh Duy, H. Xiao, S. S. Chow, and X. Liang, "Confidential machine learning computation in untrusted environments: a systems security perspective," arXiv preprint, arXiv:2111.03308, Nov. 2021.
6. M. Armanuzzaman, Y. Jin, and R. Karri, "BYOTee: building your own trusted execution environments using FPGAs," arXiv preprint, arXiv:2203.04214, Mar. 2022.
7. Y. Jia, Y. Xu, Y. Zhang, C. Lin, Y. Zhang, and Z. Wang, "HyperEnclave: an open and cross-platform trusted execution environment," in Proc. USENIX ATC, Carlsbad, CA, pp. 591–606, 2022. Available: https://www.usenix.org/conference/atc22/presentation/jia-yuekai
8. T. Geppert, T. Distler, and H. P. Reiser, "Trusted execution environments: applications and challenges," Frontiers in Computer Science, 4, Article 930741, 2022. https://doi.org/10.3389/fcomp.2022.930741
9. M. Russinovich, D. O'Donnell, and A. Zeldovich, "Toward confidential cloud computing: extending hardware-enforced protection to data in use," Communications of the ACM, 19(1), pp. 58–70, 2021. https://doi.org/10.1145/3434390