INTERNATIONAL JOURNAL OF NEXT-GENERATION ENGINEERING AND TECHNOLOGY (IJNGET)

eISSN: 3087-4319

Volume. 02, Issue. 08, pp. 07-14, August 2025

Securing Multi-Tenant FPGA Clouds: Architectures, Threats, and Integrated Defenses for Trusted Reconfigurable Computing

Sanjay K. Morello
School of Computer Engineering, University of Melbourne, Australia

Article received: 03/08/2025, Article Accepted: 14/08/2025, Article Published: 31/08/2025

ABSTRACT

Background: The rapid adoption of field-programmable gate arrays (FPGAs) within cloud infrastructures has introduced a new class of high-performance, energy-efficient accelerators for datacenter workloads. However, multi-tenant FPGA clouds create unique security, privacy, and trust challenges because hardware bitstreams, shared resources, and physical effects become new attack surfaces. This manuscript synthesizes foundational and contemporary research on FPGA cloud security and related domains — including IP protection, runtime isolation, physical unclonable functions, fault-injection attacks, operating system approaches to reconfigurable computing, and homomorphic encryption accelerators — to present an integrated conceptual framework and prescriptive design guidance for trusted multi-tenant FPGA cloud platforms. Key

contributions: (1) an expansive threat taxonomy tailored to multi-tenant FPGA clouds that links attacks to underlying physical, microarchitectural, and software vectors; (2) a layered architecture for trust that maps defensive mechanisms to threat classes, combining provenance, watermarking, PUF-based attestation, hypervisor and OS level controls, and cryptographic accelerators; (3) a detailed methodology for evaluating trustworthiness that emphasizes measurement-driven experiments and descriptive, non-mathematical analysis; and (4) an agenda for future research that prioritizes measurable isolation primitives, hardware-accelerated privacy techniques, and resilient designs against environmental and fault-based attacks.

Findings: The literature shows that heterogeneous defenses are required: IP watermarking and design fingerprints offer provenance checks but are insufficient alone (Abdel-Hamid et al., 2003); OS-centric and hypervisor approaches such as ReconOS provide strong management abstractions but must be paired with hardware attestation (Agne et al., 2014); PUFs yield compact, device-intrinsic keys appropriate for constrained devices and for bootstrapping trust in tenants and IP (Ahmed et al., 2020); fault and side-channel attacks like RAM-Jam demonstrate that environmental manipulation can violate multi-tenant isolation unless physical resource contention and environmental sensing are monitored (Alam et al., 2019); hardware accelerators for encrypted computation, including FPGA implementations for fully homomorphic encryption, offer promising paths toward preserving confidentiality even when sharing raw compute fabric (Agrawal et al., 2022).

Conclusions: A defensible multi-tenant FPGA cloud must combine provenance, device-rooted trust, runtime enforcement, and privacy-preserving computation. The paper synthesizes extant evidence to propose an integrated blueprint for secure deployment, operational monitoring, and research priorities for resilient, trustworthy FPGA clouds. All claims draw from the supplied corpus of references and directly map to the cited works.

Keywords: Multi-tenant FPGA security, trusted IP, physical unclonable functions, FPGA cloud, hardware watermarking, homomorphic encryption

INTRODUCTION

Field-programmable gate arrays (FPGAs) have evolved from niche prototyping devices into mainstream datacenter accelerators, driven by their capacity to deliver orders-of-magnitude improvements in latency, throughput, and energy efficiency for specialized kernels. The adoption of FPGAs by hyperscalers and cloud providers has enabled new classes of services where tenants request reconfigurable hardware instances to accelerate workloads such as machine learning inference, network processing, and cryptographic operations (Abel et al., 2017). This migration of reconfigurable compute into multi-tenant environments brings profound security and trust challenges because the classical boundary between hardware and software is blurred: bitstreams and partial reconfiguration define computation, and the underlying fabric contains shared resources that can be exploited by an adversarial tenant. The evolving landscape necessitates a unified, deeply technical treatment that weaves together IP protection techniques, device-rooted hardware trust, OS-level resource management for reconfigurable computing, and cryptographic strategies for protecting confidentiality in shared execution contexts.

This paper undertakes such a synthesis. The motivation is twofold. First, to provide a rigorous, conceptually integrated view of threats and defenses specific to multi-tenant FPGA clouds, grounded in peer-reviewed and conference literature. Second, to translate that view into prescriptive architectural components and evaluation methods that practitioners and researchers can adopt to strengthen trust. Prior research has examined individual elements — for example, IP watermarking approaches (Abdel-Hamid et al., 2003), reconciling high-level operating system ideas with reconfigurable fabrics (Agne et al., 2014), or using physical unclonable functions (PUFs) for device authentication (Ahmed et al., 2020). But the field lacks a comprehensive, layered framework that explicitly maps threats to combinations of defenses while asserting practical evaluation methodologies for multi-tenant cloud deployments. This paper fills that gap by integrating findings from the provided references into an extended, theoretically rich treatment of architecture, methodology, and future research directions.

Problem statement: Multi-tenant FPGA clouds must reconcile three conflicting objectives: (1) high resource utilization and low latency for tenants; (2) cryptographic and provenance assurances for IP and data confidentiality; and (3) robust isolation against physical, side-channel, and fault-injection threats that exploit shared hardware and environmental interactions. Existing approaches often address one or two aspects but fail to deliver a unified, empirically grounded framework for production systems. For example, IP watermarking mechanisms can identify provenance but do not inherently prevent runtime extraction or side-channel leakage (Abdel-Hamid et al., 2003). ReconOS offers an OS abstraction for scheduling and management but requires secure binding to hardware roots of trust and environmental monitoring to be effective in hostile, multi-tenant scenarios (Agne et al., 2014). PUFs provide strong device identity primitives suitable for IoT and constrained devices, but their use within large, reconfigurable datacenters is underexplored in combination with runtime enforcement (Ahmed et al., 2020). Fault attacks such as RAM-Jam demonstrate that memory collisions and environmental manipulation can break isolation guarantees if architecture designers do not explicitly consider such vectors (Alam et al., 2019). Thus, the literature indicates an urgent need for an integrated, evidence-based design that blends provenance, device attestation, runtime isolation, and privacy-preserving computation.

Literature gap: While individual studies have advanced the state of knowledge in their niches, there is sparse literature that systematically synthesizes these components into an actionable pathway for cloud operators. Existing surveys and comparisons (Abdel-Hamid et al., 2003) are valuable for IP protection techniques, but they predate the contemporary cloud FPGA paradigm and do not accommodate multi-tenant dynamics. ReconOS (Agne et al., 2014) provides a starting point for OS abstractions, yet does not fully consider adversarial tenants or the cryptographic primitives required for remote attestation in cloud environments. Recent work has begun exploring trustworthy IP and multi-tenant FPGA concerns (Ahmed et al., 2022), but comprehensive architectures that operationalize layered defense — including provenance, PUF attestation, environmental monitoring, cryptographic acceleration, and fault-attack mitigations — remain missing. This paper addresses that multifaceted gap by articulating an integrated architecture and by proposing rigorous, descriptive evaluation methods grounded in the referenced corpus.

METHODOLOGY

This study is conceptual and integrative: it constructs an architecture and evaluation methodology by synthesizing evidence from the provided literature and mapping threats to defenses. The approach intentionally avoids mathematical formalism and instead provides descriptive, experiment-oriented evaluation strategies suitable for practitioners. The methodology consists of four sequential steps: (1) exhaustive literature mapping; (2) threat taxonomy derivation; (3) layered defense architecture design; and (4) descriptive evaluation protocol. Each step draws on multiple references and highlights practical implications for cloud operators and researchers.

Literature mapping: The referenced corpus encompasses IP watermarking surveys (Abdel-Hamid et al., 2003), FPGA cloud platforms for hyperscalers (Abel et al., 2017), OS approaches for reconfigurable computing (Agne et al., 2014), accelerators for cryptographic privacy (Agrawal et al., 2022), empirical attack investigations such as RAM-Jam (Alam et al., 2019), PUF-based security for constrained devices (Ahmed et al., 2020), and trust proposals for multi-tenant FPGA platforms (Ahmed et al., 2022). Additional references on cloud security, authentication, cryptographic approaches, and access control provide broader context (Lu et al., 2018; Lamport, 1981; Al-Assam et al., 2019; Kim & Lee, 2018; Sachdev & Bhansali, 2013; Batista et al., 2017; Reddy et al., 2015; Aluvalu et al., 2016; Atayero & Feyisetan, 2011). These works collectively inform the taxonomy and architecture.

Threat taxonomy derivation: We classify threats along orthogonal axes: origin (insider vs. outsider tenant), vector (bitstream, side-channel, fault injection, environmental manipulation), target (IP confidentiality, tenant data, hardware integrity), and stage (design time, provisioning, runtime). This multidimensional taxonomy is derived by mapping observed attacks (e.g., memory collision induced faults in RAM-Jam) to their root causes and to the mechanisms that existing defenses target (e.g., watermarking targets IP provenance but not side channels). Each mapping references primary studies to ensure the taxonomy is empirically grounded.

Layered defense architecture design: The architecture is structured into five concentric layers: device root (PUF and hardware attestation), provenance (IP watermarking and fingerprinting), runtime enforcement (OS and hypervisor level controls, e.g., ReconOS and trusted IP solutions), environmental monitoring and fault mitigation (techniques to detect and respond to RAM-Jam–like attacks), and privacy-preserving compute (hardware accelerators for homomorphic encryption). The architecture prescribes how each layer interlocks: device attestation anchors trust; provenance enables later dispute resolution; runtime enforcement constrains sharing and enforces policies; environmental monitoring closes physical attack channels; and encrypted computation reduces the need to trust co-located fabrics.

Descriptive evaluation protocol: Rather than prescribing formulas, the methodology outlines stepwise descriptive experiments: provenance validation via watermark detection; attestation benchmarking using PUF challenge-response protocols under temperature and voltage variation; runtime isolation experiments using OS scheduling scenarios with mixed benign and adversarial tenants; fault injection campaigns that emulate RAM-Jam conditions; and encrypted workload performance evaluation using FPGA-accelerated homomorphic kernels (Agrawal et al., 2022). Each experiment emphasizes observable metrics (detection rate, false positives, latency overhead, throughput reduction, and resilience under environmental perturbation) and includes qualitative observations about operational tradeoffs.

Ethical considerations and reproducibility: The methodology stresses safe and ethical experimentation with fault-injection and side-channel measurement, recommending controlled lab environments and nonproduction fabrics. Reproducibility is encouraged by recommending standardized experiment descriptions and by drawing on published methods (Alam et al., 2019; Agne et al., 2014) as procedural templates.

RESULTS

This section synthesizes descriptive findings from the literature and from thought experiments derived via the methodology. Because the study is integrative and conceptual rather than experimental, the results are presented as interdependent observations that emerge when the layered architecture is applied to the threat taxonomy and when prior experimental findings are considered in a multi-tenant context.

Observation 1 — Provenance mechanisms provide valuable forensic evidence but are insufficient for runtime protection. IP watermarking techniques can embed metadata and fingerprints into hardware designs that survive synthesis and partial reconfiguration, enabling later verification of ownership and provenance (Abdel-Hamid et al., 2003). These techniques are instrumental for intellectual property disputes and for post-incident attribution; however, watermarking operates primarily at the design and deployment boundary and does not prevent a malicious tenant from extracting sensitive functionality or from performing side-channel analysis during execution. In practice, watermarking must be combined with runtime controls and cryptographic protections to preserve confidentiality (Abdel-Hamid et al., 2003).

Observation 2 — OS and hypervisor abstractions tailored for reconfigurable fabrics enable stronger management but need hardware-anchored attestation. ReconOS demonstrated the practicality of an operating system approach for reconfigurable computing, offering scheduling and resource management primitives that resemble traditional OS functions yet apply to hardware tasks and software threads (Agne et al., 2014). Such abstractions allow cloud operators to allocate fabric slices, enforce CPU-like scheduling, and provide standardized interfaces for tenants. However, without hardware-anchored attestation — a device-rooted identity and measurement — OS-level controls cannot reliably detect or prevent malicious bitstreams or compromised firmware (Agne et al., 2014). Thus, ReconOS-like approaches are necessary but not sufficient: they must be bound to device roots of trust.

Observation 3 — Physical unclonable functions (PUFs) present a practical anchor for device identity and key generation in reconfigurable contexts, but environmental sensitivity must be managed. PUFs exploit manufacturing variations to derive device-unique secrets that are difficult to clone and can support authentication and key provisioning for resource-constrained devices (Ahmed et al., 2020). In the cloud FPGA context, PUFs could be used to bind bitstreams to particular hardware instances, enabling attestation of provenance at boot and preventing bitstream migration between untrusted fabrics. However, PUF response stability across temperature and voltage variations is a known challenge; robust PUF deployment requires error-correcting helper data and protocols that tolerate environmental noise while preserving security (Ahmed et al., 2020). Deploying PUFs in hyperscale fabrics thus demands careful engineering and monitoring.

Observation 4 — Fault and environmental attacks such as RAM-Jam exploit shared physical resources and can subvert isolation unless detected by environmental sensing and scheduling policies. RAM-Jam demonstrated that remote temperature and voltage manipulation combined with memory collision patterns can induce faults in FPGA fabrics, compromising cryptographic operations and other critical tasks (Alam et al., 2019). In multi-tenant clouds, such attacks could be orchestrated by a malicious tenant to influence co-located workloads. Mitigation requires both hardware measures (e.g., thermal sensors, voltage monitors, conservative guard bands) and software policies (e.g., avoiding placement of mutually distrusted tenants in contiguous regions, throttling suspicious memory patterns). The empirical evidence strongly recommends integration of continuous environmental monitoring and reactive placement/scheduling logic (Alam et al., 2019).

Observation 5 — Trusted IP solutions for multi-tenant FPGA platforms combine hardware and software primitives to limit exposure of third-party designs but performance and expressivity tradeoffs persist. Trusted IP research for cloud FPGA platforms articulates mechanisms to control the lifecycle of IP blocks — from authenticated provisioning and binding to runtime monitoring and revocation — thereby reducing the risk of IP theft and misuse in multi-tenant settings (Ahmed et al., 2022). These approaches leverage attestation, secure bitstream channels, and runtime guards. Nonetheless, constraints on flexibility and developer productivity arise when strict trust controls limit dynamic partial reconfiguration or impose heavy cryptographic overheads. Practical deployment requires balancing security objectives with the dynamic needs of cloud tenants (Ahmed et al., 2022).

Observation 6 — Hardware acceleration of cryptographic primitives, and particularly homomorphic encryption (HE), offers a path to reduce the trust required in shared fabrics by enabling computation over encrypted data. Recent FPGA-based accelerators for fully homomorphic encryption show that bootstrappable HE can be significantly accelerated with reconfigurable logic, making privacy-preserving workloads more realistic in datacenter settings (Agrawal et al., 2022). While HE remains computationally intensive, domain-specific accelerators can reduce overhead sufficiently to make encrypted pipelines tractable for certain workloads. Consequently, integrating HE accelerators into multi-tenant fabrics can materially reduce data exposure risk even when tenants share physical resources (Agrawal et al., 2022).

Observation 7 — Hyperscaler FPGA platforms emphasize integration, scale, and management simplicity but must incorporate security and trust as first-class design requirements. Practical deployments by hyperscalers demonstrate the operational benefits of offering FPGA instances as cloud services, yet these platforms also reveal that scale introduces new attack surfaces that are absent in single-tenant or on-premises settings (Abel et al., 2017). Hyperscalers must therefore prioritize secure provisioning pipelines, hardware roots of trust, and isolation mechanisms that scale to thousands of devices (Abel et al., 2017).

Observation 8 — Classic cryptographic and authentication approaches remain foundational for cloud security, but novel combinations with hardware roots of trust and PUFs produce stronger guarantees in reconfigurable contexts. Foundational protocols such as Lamport's password authentication remain important for secure channels, while higher-level models like attribute-based and hierarchical access control can help structure tenant privileges. Combining cryptographic key management (e.g., AES and Blowfish based approaches) with PUF-derived keys and secure hardware boot chains can improve both security and operational usability (Lamport, 1981; Sachdev & Bhansali, 2013; Reddy et al., 2015; Aluvalu et al., 2016). These combinations require careful attention to key lifecycle, helper data protection, and scalability of issuance.

Observation 9 — Quality-of-service (QoS) and security are intertwined in the cloud: performance-driven placement and scheduling choices can exacerbate or mitigate security risks. QoS models that optimize for latency and throughput must be augmented with security-aware placement policies; otherwise, optimizing for performance alone can place high-risk workloads adjacent to sensitive tenants, thereby increasing exposure to side channels and fault attacks (Batista et al., 2017). A security-aware QoS framework must therefore incorporate hardware-level constraints and environmental monitoring signals.

Observation 10 — Access control and key management at scale remain pressing operational issues; attribute-based models and hierarchical key distribution offer scalable templates but require robust trust anchors. HASBE and similar hierarchical models provide efficient ways to manage access at scale, but their effectiveness depends on secure key distribution and the integrity of the key management infrastructure (Aluvalu et al., 2016). In multi-tenant FPGA clouds, integrating HASBE-like models with hardware attestation and secure provisioning pipelines can reduce the attack surface for key compromise (Aluvalu et al., 2016).

DISCUSSION

This section synthesizes the observations into an interpretive analysis, exploring theoretical implications, counterarguments, practical tradeoffs, and recommended engineering pathways.

Interlocking defenses: The central theoretical proposition is that no single defensive primitive suffices for multi-tenant FPGA clouds because attacks span multiple layers of abstraction — from design-time IP theft to runtime side channels and environmental fault injection. This multi-layered reality implies that security must be built as an interlocking set of complementary mechanisms: device-rooted attestation (PUFs and secure boot) anchors identity and integrity; provenance mechanisms (watermarking) allow ownership verification and forensic analysis; OS and hypervisor constructs (ReconOS-style) enable standardized, enforceable resource controls; environmental monitoring detects and mitigates physical attack vectors (e.g., RAM-Jam); and privacy-preserving compute (HE accelerators) reduces the semantic sensitivity of data executed on shared resources (Abdel-Hamid et al., 2003; Agne et al., 2014; Ahmed et al., 2020; Alam et al., 2019; Agrawal et al., 2022). Each layer compensates for limitations in others: provenance cannot prevent runtime leakage but can help resolve disputes; PUFs provide attestation but require helper data mechanisms and monitoring to compensate for environmental drift; HE reduces data exposure but needs acceleration to be practical at scale.

Counterarguments and limitations: A reasonable counterargument is that introducing many layers of defense increases system complexity and could hurt usability, performance, and cost. Indeed, strict hardware binding of bitstreams and pervasive attestation can complicate dynamic reconfiguration and long-tail tenant workflows. Moreover, PUFs and helper data schemes introduce reliability challenges and additional hardware design constraints (Ahmed et al., 2020). There are also cost and energy tradeoffs when integrating HE accelerators — they provide privacy but consume die area and power and may still be insufficient for extremely latency-sensitive workloads (Agrawal et al., 2022). The recommended approach is therefore not to apply every defense universally but to adopt a risk-based, workload-aware policy: sensitive data pipelines receive the full suite (hardware attestation, provenance, HE), while less sensitive batch workloads may accept lighter measures. Hyperscalers can codify these options into service tiers to balance security, cost, and flexibility (Abel et al., 2017).

Operational implications: The layered architecture implies concrete operational changes for cloud operators. First, provisioning pipelines must incorporate attestation checks and binding of bitstreams to hardware identities; this requires secure channels, key management, and provenance registries to which watermarking traces can be registered. Second, the scheduler must be security-aware, integrating environmental telemetry and workload trust attributes to avoid risky co-placement. Third, incident response must include forensic capabilities that combine watermark verification, PUF-based device logs, and runtime telemetry to reconstruct breaches. These operational shifts are substantial but feasible: Reconstruction of OS-level abstractions for reconfigurable workloads (Agne et al., 2014) and trusted IP proposals (Ahmed et al., 2022) provide pragmatic starting points for cloud operators.

Research priorities and future scope: Several priority research threads emerge from the synthesis. One is the robust integration of PUFs into hyperscale FPGA clouds: designing helper data schemes that minimize key leakage and maximize stability under wide environmental ranges is crucial (Ahmed et al., 2020). Second, developing lightweight, practical watermarking techniques that can survive aggressive optimization and partial reconfiguration without substantial overhead remains an open problem (Abdel-Hamid et al., 2003). Third, advancing hardware accelerators for privacy primitives — especially homomorphic encryption — to the point where they can support generalizable workloads with acceptable overheads is a promising direction (Agrawal et al., 2022). Fourth, understanding the dynamics of environmental attacks at scale (e.g., RAM-Jam) and designing cloud-scale countermeasures such as thermal zoning, adversarial placement, and continuous environmental attestation requires more empirical work (Alam et al., 2019). Finally, more research is needed on composable security policies that reconcile QoS, performance, and security in real time, perhaps borrowing from access control models such as HASBE and combining them with device attestation (Aluvalu et al., 2016; Batista et al., 2017).

Evaluation and measurement: The descriptive evaluation protocol proposed in this paper emphasizes measurable observables rather than formal proofs, reflecting the empirical nature of hardware-level attacks and defenses. In practice, operators should track detection rates for provenance checks, error rates for PUF responses under stress, latency and throughput impacts of HE accelerators, environmental anomalies indicative of RAM-Jam-style attacks, and policy adherence metrics for the resource manager. Importantly, evaluation must be contextual — the acceptable thresholds vary by workload sensitivity and by service tier. The literature suggests that combined measurement of detection fidelity and operational cost will be the most informative for adoption decisions (Agrawal et al., 2022; Alam et al., 2019; Ahmed et al., 2022).

Ethical and regulatory considerations: Deploying FPGA clouds with stronger security controls affects both tenants and regulators. From a tenant perspective, stronger attestation and provenance may be welcomed by IP holders but could introduce friction for independent developers and researchers who require dynamic reconfiguration. Regulators may view hardware-anchored attestation favorably because it reduces data leakage risks, but they will require transparency about key management and privacy protections for helper data. The design choices should therefore be guided by principles of least privilege, measurable accountability, and clear contractual protections for tenant IP and data.

CONCLUSION

This manuscript synthesizes contemporary and foundational literature to present an integrated architecture and evaluation approach for securing multi-tenant FPGA clouds. The principal insight is that trustworthy reconfigurable cloud infrastructures require a layered defense strategy: device-rooted attestation (e.g., PUFs), provenance mechanisms (IP watermarking), runtime enforcement (OS and hypervisor controls such as ReconOS and trusted IP frameworks), environmental monitoring and fault mitigation (countering attacks like RAM-Jam), and privacy-preserving compute (hardware accelerators for homomorphic encryption). Each layer compensates for the weaknesses of others and collectively forms a resilient posture against diverse adversaries.

The proposed methodology emphasizes descriptive, measurement-driven protocols for evaluating trustworthiness in real systems. The literature demonstrates both the potential and the limitations of each defensive primitive: watermarking supports forensics (Abdel-Hamid et al., 2003), ReconOS provides OS-style control for reconfigurable fabrics (Agne et al., 2014), PUFs offer compact device identities suitable for attestation (Ahmed et al., 2020), RAM-Jam highlights the vulnerability of shared physical resources (Alam et al., 2019), and FPGA-based homomorphic encryption accelerators enable privacy-preserving computation (Agrawal et al., 2022). Practical adoption requires engineers to trade off security, performance, and flexibility using risk-aware policies and to implement continuous monitoring and agile incident response.

Future work should focus on robust PUF deployment at scale, resilient watermarking techniques that withstand optimization and partial reconfiguration, practical HE accelerators optimized for realistic datacenter workloads, and cloud-scale strategies for detecting and mitigating environmental attacks. Additionally, the design and study of security-aware schedulers that integrate environmental telemetry, provenance signals, and workload sensitivity represent fertile ground for impactful research.

In closing, multi-tenant FPGA clouds promise transformative performance and efficiency gains, but realizing that promise responsibly demands a holistic, evidence-based approach to security and trust. By aligning device roots of trust, provenance mechanisms, runtime enforcement, environmental defenses, and privacy-preserving computation, cloud operators and researchers can create reconfigurable infrastructures that are both powerful and defensible.

REFERENCES

1. T. Abdel-Hamid, S. Tahar, and El Mostapha Aboulhamid. 2003. IP watermarking techniques: Survey and comparison. In Proceedings of the 3rd IEEE International Workshop on System-on-Chip for Real-Time Applications, 60–65. DOI: https://doi.org/10.1109/IWSOC.2003.1213006

2. Francis Abel, Jagath Weerasinghe, Christoph Hagleitner, Beat Weiss, and Stephan Paredes. 2017. An FPGA platform for hyperscalers. In Proceedings of the IEEE 25th Annual Symposium on High-Performance Interconnects (HOTI ’17), Vol. 10, 29–32. DOI: https://doi.org/10.1109/HOTI.2017.13

3. Andreas Agne, Markus Happe, Ariane Keller, Enno Lubbers, Bernhard Plattner, Marco Platzner, and Christian Plessl. 2014. ReconOS: An operating system approach for reconfigurable computing. IEEE Micro 34, 1, 60–71. DOI: https://doi.org/10.1109/MM.2013.110

4. Rashmi Agrawal, Leo de Castro, Guowei Yang, Chiraag Juvekar, Rabia Yazicigil, Anantha Chandrakasan, Vinod Vaikuntanathan, and Ajay Joshi. 2022. FAB: An FPGA-based accelerator for bootstrappable fully homomorphic encryption. arXiv:2207.11872. Retrieved from https://doi.org/abs/2207.11872

5. Muhammed Kawser Ahmed, Sujan Kumar Saha, and Christophe Bobda. 2022. Trusted IP solution in multi-tenant cloud FPGA platform. In Proceedings of the IEEE 8th World Forum on Internet of Things (WF-IoT ’22), 1–6. DOI: https://doi.org/10.1109/WFIoT54382.2022.10152167

6. Muhammed Kawser Ahmed, Venkata P. Yanambaka, Ahmed Abdelgawad, and Kumar Yelamarthi. 2020. Physical unclonable function based hardware security for resource constraint IoT devices. In Proceedings of the IEEE 6th World Forum on Internet of Things (WF-IoT ’20), 1–2. DOI:https://doi.org/10.1109/WFIoT48130.2020.9221357

7. Md Mahbub Alam, Shahin Tajik, Fatemeh Ganji, Mark Tehranipoor, and Domenic Forte. 2019. RAM-Jam: Remote temperature and voltage fault attack on FPGAs using memory collisions. In Proceedings of the Workshop on Fault Diagnosis and Tolerance in Cryptography (FDTC ’19), 48–55. DOI: https://doi.org/10.1109/FDTC.2019.00015

8. Lu, X., Cao, L., Du, X., et al. 2018. A tag-based protection method for multi-tenant data security. International Conference on Cloud Computing and Security, Haikou, China, 553–565.

9. Lamport, L. 1981. Password authentication with insecure communication. Communications of the ACM, 24(11), 770–772.

10. Al-Assam, H., Hassan, W., Zeadally, S. 2019. Automated biometric authentication with cloud computing. In Biometric-Based Physical and Cybersecurity Systems, 455–475.

11. Kim, S. H., Lee, I. Y. 2018. IoT device security based on proxy re-encryption. Journal of Ambient Intelligence and Humanized Computing, 9(4), 1267–1273.

12. Sachdev, A., Bhansali, M. 2013. Enhancing cloud computing security using AES algorithm. International Journal of Computer Applications,

67(9), 0975–8887.

13. Batista, B. G., Ferreira, C. H., Segura, D. C., et al. 2017. A QoS-driven approach for cloud computing addressing attributes of performance and security. Future Generation Computer Systems, 68, 260–274.

14. Reddy, T. B., Chowdappa, K. B., Reddy, S. R. 2015. Cloud security using blowfish and key management encryption algorithm. International Journal of Engineering and Applied Sciences, 2(6), 2394–3661.

15. Aluvalu, R., Kamliya, V., Muddana, L. 2016. HASBE access control model with secure key distribution and efficient domain hierarchy for cloud computing. International Journal of Electrical and Computer Engineering, 6(2), 770–777.

16. Atayero, A. A., Feyisetan, O. 2011. Security issues in cloud computing: the potentials of homomorphic encryption. Journal of Emerging Trends in Computing and Information Sciences, 2(10), 546–552.

17. V. Agarwal, N. Verma, S. Saha, and S. Kumar. 2018. Dynamic Detection and Prevention of Denial of Service and Peer Attacks with IP Address Processing. Recent Findings in Intelligent Computing Techniques: Proceedings of the 5th ICACNI 2017, Volume 1, 707, 139.

18. Mishra, M. 2017. Reliability-based Life Cycle Management of Corroding Pipelines via Optimization under Uncertainty (Doctoral dissertation).

19. V. Agarwal and S. Kumar. 2017, October. Intelligent multi-level mechanism of secure data handling of vehicular information for post-accident protocols. In 2017 2nd International Conference on Communication and Electronics Systems (ICCES), 902–906.

20. Hariharan, R. 2025. Zero trust security in multi-tenant cloud environments. Journal of Information Systems Engineering and Management, 10.