Electronic Design

By DAVID W. ARCHER, CTO, Niobium

11 Myths About Fully Homomorphic Encryption

No longer a futuristic concept, fully homomorphic encryption enhances data privacy while providing security.

Fully homomorphic encryption (FHE) has emerged as a powerful tool for improving data privacy, yet its rise has been accompanied by a host of misconceptions. This article dispels the notion that FHE is purely theoretical and explores the shortcomings of traditional data-security methods, providing an informed perspective on why FHE isn't just a concept but a viable solution for secure data processing today.

Whether you're skeptical about FHE's real-world applications or curious about its efficiency, this discussion will offer deeper insights into the technology that will reshape data privacy.

1. FHE is a new technology.

While FHE might seem novel, it's been in the research phase for decades, starting with the discovery of partially homomorphic systems. The technology took another significant step toward its full capability in 2009. The implications were significant: third parties could process data without risking privacy breaches. However, the computational power required to implement FHE remained prohibitively high, confining its use to academic research.

This landscape has changed. As the data economy grows and regulations like GDPR and CCPA demand stronger se-

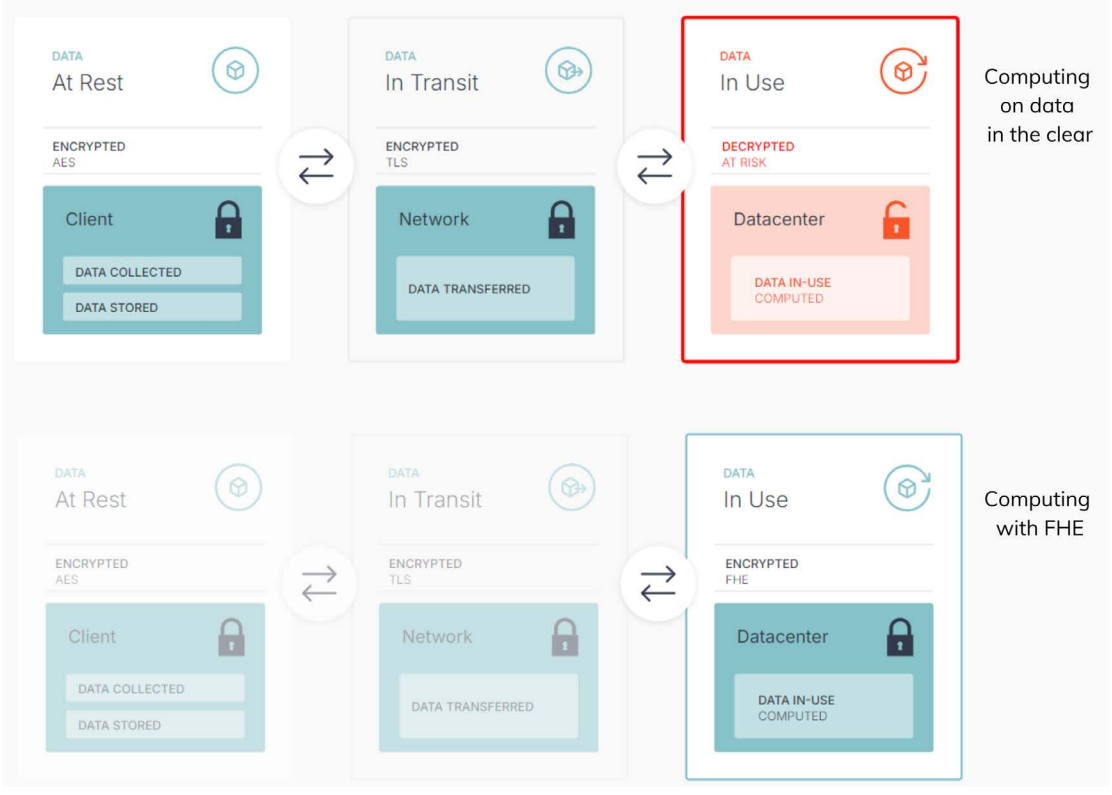

Traditionally, data must be unencrypted to be computed on. With FHE, data in use stays encrypted. This image shows data in use with and without the use of FHE.

curity measures, advances in FHE software and hardware have made it commercially viable. Today, specialized FHE accelerators are bringing encrypted computing into practical applications across industries *(see figure)*.

2. FHE is only theoretical and not ready for practical use.

The belief that FHE is purely theoretical needs to be updated. Niobium's first-generation FHE accelerator demonstrates significant performance improvements, making FHE viable for real-world applications. Although (some outrageous claims aside) encrypted computation will always be slower than traditional computation, recent advances have closed the gap considerably, making its adoption a reality.

This breakthrough has unlocked practical applications for FHE in sectors such as finance, insurance, and healthcare, where data security and privacy are critical. A range of new and previously infeasible use cases is now possible based on early adopters' feedback. Companies can now securely process and analyze sensitive data without compromising confidentiality, paving the way for a safer and more profitable data economy.

3. Confidential computing and trusted execution environments are good enough.

Confidential computing and trusted execution environments (TEEs), such as Intel's Software Guard Extensions (SGX) and Trust Domain Extensions (TDX), AMD Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV), and Arm's TrustZone, have been widely adopted as solutions for securing sensitive data during processing.

However, the reality is that these technologies, while useful, fall short of providing the comprehensive security guarantees needed in today's increasingly complex threat landscape. Unlike FHE, which offers cryptographic proof of security, TEEs don't have such guarantees, making them vulnerable to both side-channel attacks and direct breaches, placing the companies that rely on them at risk.

The fundamental problem with TEEs is that they lack cryptographic proof of security, arguably because the *root of trust* of these systems is too large and complex to tackle with such proofs. These systems do their best to provide isolation, but that "best effort" has repeatedly failed in practice.

Researchers have uncovered numerous side-channel attacks and direct breaches of SGX. For instance, an architectural flaw in certain Intel CPUs made it possible for attackers to bypass SGX's security features, exposing sensitive data. This isn't an isolated incident; academic researchers continue to exploit weaknesses in TEEs, demonstrating that these environments can leak sensitive information. It thus leaves organizations with limited assurances in their quest for robust protection.

In contrast, FHE keeps data encrypted throughout computation and provides a comprehensive security definition and proof, eliminating risks inherent to TEE-based security models. In a world where data privacy is paramount, relying solely on TEEs for secure data processing is increasingly risky.

FHE provides a robust alternative, offering mathematically guaranteed security without the vulnerabilities that TEEs cannot avoid. As an added benefit, FHE is fully *portable* between processor architectures, while TEEs (and their related software) are not, requiring substantial adaptation to move TEE-based applications from one platform to another.

4. FHE doesn't have real-world applications.

Contrary to the belief that FHE is purely theoretical, FHE is already making an impact across various industries, proving its practical value in enhancing data privacy and security.

For instance, Fhenix offers a fully homomorphic encryption platform that enables secure data computation across industries like finance and healthcare. Their solutions enable companies to analyze sensitive data collaboratively without exposing raw data, ensuring compliance with strict privacy regulations while unlocking valuable insights. Similarly, Apple's Swift Homomorphic Encryption, announced as an open-source library, allows developers to incorporate homomorphic encryption directly into their applications.

This democratization of FHE tools empowers businesses to integrate secure computing into day-to-day operations, from privacy-preserving health analytics to confidential financial modeling. Other industries poised to benefit from FHE:

- **Healthcare:** Enables secure patient data analysis while maintaining privacy.
- **Finance:** Supports secure fraud detection and risk modeling without exposing sensitive information.
- **Cloud computing:** Ensures encrypted processing, eliminating data breach risks.
- **Machine learning:** Allows for the application of certain machine-learning techniques to process encrypted data, enhancing privacy in applications like secure customer segmentation.

5. FHE needs to be as fast as "computing in the clear" to offer commercial value.

A common misconception is that FHE must match the speed of unencrypted computing to be commercially viable. However, this overlooks the true value of FHE in applications where data privacy and accuracy outweigh real-time processing.

Historically, the transition from HTTP to HTTPS provides a valuable analogy. When HTTPS was first introduced, it offered a significant improvement in security—but it came with a performance cost. Initially, this performance delta

deterred adoption. However, as the gap in speed diminished to an acceptable range, the value of added security began to outweigh the perceived loss of performance: secure operation became *fast enough*. This shift eventually led to HTTPS becoming the standard for web communication.

Similarly, while FHE doesn't yet match the speed of "computing in the clear," its performance is improving to a point where the value of data privacy outweighs the operational inefficiencies in selected early adopter applications. For example, in fields like machine-learning inference and image recognition, the integrity and confidentiality of data are more critical than achieving instantaneous results, and processing is emerging to the "fast enough" stage even as we speak.

Niobium is focused on narrowing this gap. Although the first generation of its FHE accelerator chip hasn't been evaluated or commercially deployed yet, it aims to bring FHE speeds closer to those of clear-text computation. The goal is to achieve a performance level where the benefits of privacy and security far exceed the perceived tradeoffs, paving the way for mainstream adoption much like the move to HTTPS did in the past.

6. FHE hardware is just repurposed AI hardware.

Attempts have already been made to pivot from AI-focused hardware design to apply the same technology to FHE. However, it simply doesn't work as claimed. FHE requires specialized hardware that differs significantly from AI-centric architectures.

Purpose-built FHE accelerators are designed from the ground up to handle FHE's unique demands, including large vector operations on prime-modular fields, frequency-domain transforms, and very large memory bandwidth and memory volume demands. Unlike repurposed AI hardware, these FHE-specific designs avoid the performance penalties associated with adapting architectures not intended for encrypted data processing, ensuring optimized throughput and efficiency in FHE operations.

7. FHE solves all privacy problems.

Though FHE secures data during processing, it doesn't solve all privacy challenges, such as ensuring the correct program is executed or assuring that the output of a computation reveals no sensitive information about inputs to the computation. To fully protect data, FHE should be combined with other technologies in a Zero Trust Compute (ZTC) framework, which ensures data privacy and integrity across storage, transmission, and processing.

8. FHE can be easily implemented on any hardware.

FHE requires unique mathematics, memory access patterns, and large vector operations—most general-purpose hardware architectures aren't optimized for that. Niobium's chip is designed to handle these specific requirements efficiently.

9. FHE is slow and inefficient.

FHE offers unparalleled data privacy and security, but this comes at a cost, particularly in terms of efficiency. FHE-encrypted data expands significantly, and its computations are more complex than processing unencrypted data.

However, recent advances in FHE hardware acceleration are changing the game. By significantly boosting processing speeds, purpose-built FHE accelerators have made FHE viable for real-world use on selected applications.

While still slower than "in the clear" computing, FHE can now operate at speeds close enough to "in the clear" to meet organizational workflow throughput demands. With ongoing enhancements, FHE is increasingly viable for applications in finance, healthcare, and beyond, demonstrating its growing efficiency and practicality.

10. FHE is too complex for developers to implement.

Many assume that implementing FHE requires extensive cryptographic expertise, making it inaccessible to average developers. While FHE does involve sophisticated mathematics, the development of high-level libraries and frameworks has dramatically simplified implementation.

Tools like Microsoft SEAL, IBM's HElib, and Apple's Swift Homomorphic Encryption provide developer-friendly interfaces that abstract away much of the underlying complexity. These libraries come with comprehensive documentation and examples, enabling developers with standard programming skills to integrate FHE into their applications.

Is there more work to do on usability? Absolutely. Today, emerging efforts such as Google's HEIR framework are developing the next round of usability improvements and application development automation. Still to do: We need to bring FHE application development into the undergraduate training regimen of computer scientists and programmers. Partnerships between the FHE community and the educational system will be an important step moving forward.

The industry has also made strides in interoperability and standardization through initiatives like the FHE Technical Consortium for Hardware (FHETCH). This international, industry-led consortium brings together hardware and software developers, application owners, and end-users to advance FHE technology's commercial adoption.

FHETCH plays a crucial role in defining open APIs and standards that ensure seamless integration across diverse platforms. By developing practical performance benchmarks and fostering innovation in hardware acceleration, FHETCH is making it significantly easier for developers to implement FHE solutions without requiring deep cryptographic expertise.

These collaborative efforts are creating a more accessible ecosystem where developers can leverage standardized interfaces, consistent benchmarks, and improved tooling. Organizations interested in implementing FHE can now join this growing community rather than developing solutions in isolation, further reducing the technical barriers to adoption.

The learning curve, while steeper than that for conventional encryption, is becoming increasingly manageable as the ecosystem matures, educational resources expand, and industry standards become established through bodies like FHETCH.

11. FHE is only relevant for large enterprises with sensitive data.

FHE benefits businesses of all sizes. As data privacy regulations expand, small- and medium-sized enterprises must also comply with security requirements. FHE enables businesses to collaborate on analytics while protecting sensitive information.

Startups, healthcare providers, and financial institutions are already adopting FHE to meet compliance needs and unlock new opportunities. As adoption grows, FHE will become increasingly accessible to a wider range of organizations.

FHE is no longer a futuristic concept—it's a rapidly advancing technology with real-world applications across multiple industries. While challenges remain, improvements in FHE hardware, software, and standardization efforts are making it more practical than ever.

As privacy concerns grow and regulatory pressures increase, FHE is set to become a foundational component of secure data processing. Organizations that embrace FHE now will be well-positioned for the future of data security.

David Archer, Ph.D., is Chief Technology Officer of Niobium Microsystems, maker of cutting-edge hardware acceleration for privacy-enhancing technologies (PETs), and the first company to develop ASIC-based accelerators for Fully Homomorphic Encryption. He has over 40 years of research and development experience in system hardware and software architecture, secure computation, cryptography, and data-intensive systems.

Dr. Archer leads and has led 12 years of research in PETs, applied cryptography, and information security totaling over $75M in investment by the U.S. Government and private industry. He's a member of the U.S.-wide Expert Panel for the U.S. National Secure Data Service; a founding member of the United Nations Privacy Preserving Task Team; served as one of seven judges nationwide in the U.S. PETs Prize Challenge; and serves on the cryptographic advisory board of cutting-edge company Callisto Project.

Dr. Archer holds a PhD in Computer Science from Portland State University, and an MS in Electrical Engineering and BS in Computer Engineering from the University of Illinois at Urbana-Champaign.