Building Comprehensive Hardware Security

A Lattice Semiconductor White Paper.

May 2019

[Four Million Vehicles Recalled After Jeep Is Hacked](#)
[FDA Recalls Pacemakers Due to Hacking Vulnerability](#)
[Largest DDOS Attack Ever Exploits IoT Devices](#)
[At Least 3 Billion Chips Have Security Hole](#)

What's the common denominator in all of the above news stories? In each case hackers exploited vulnerabilities in an Internet-connected device to steal data or hijack devices. In 2018 alone, attacks on hardware rendered over 3 billion chips in systems of all types open to data theft, improper operation and other security threats. Unsecured hardware threatens system reliability and performance and can undermine the customer experience. It also exposes OEMs to financial and brand damage that can quickly impact a company's reputation and its bottom line.

Learn more:

www.latticesemi.com

Contact us online:

www.latticesemi.com/contact
www.latticesemi.com/buy

TABLE OF CONTENTS

OEMs are interested in developing secure hardware that addresses a number of security threats including data theft, data corruption, equipment hijacking, cloning and design theft. Moreover, security threats are no longer confined to systems in active use. Attackers target components anywhere in the product lifecycle, from initial component manufacturing and shipment to a contract manufacturer, to system integration and on through its entire operating life. Accordingly, OEMs need a robust security solution that protects hardware from these threats across every stage of a system's lifecycle.

How can OEMs address this problem? They must establish one or more hardware root-of-trust (RoT) devices to be used as a platform to provide cryptographic capabilities that secure their systems. These include data encryption, data authentication, firmware authentication, system authentication and code/configuration encryption.

A RoT device is the first link in a chain-of-trust that protects the entire system. Once designers have identified the first trusted device (usually a PLD, FPGA or MCU), it can serve as as the foundation that enables the cryptographic functions required to secure system hardware. RoT devices must contain the hardware necessary to verify their own configuration and should be the first digital devices to boot at power up and the last to shut down at power off.

What kind of security architecture do system designers need when both the number and sophistication of threats is constantly rising? First and foremost, any solution must be robust enough to protect against new and existing threats to firmware. To help designers measure the capability of their solution, the National Institute of Standards and Technology (NIST) recently defined a new uniform security mechanism. The NIST SP 800 193 Platform Firmware Resilience (PFR) guidelines were designed to comprehensively ensure a RoT is established to all system firmware.

Developers of the new specification were driven by three guiding principles:

- Protection: protect non-volatile firmware memory through access control
- Detection: cryptographically detects and prevents booting from malicious code
- Recovery: in case of corruption the system recovers to the latest trusted good firmware

Engine Options

Ideally an engine for providing hardware security would consume little power, offer a high degree of design flexibility, be scalable and occupy a small physical footprint. MCUs offer excellent computational resources, but typically don't offer the comprehensive capabilities needed to help boot other system processors or components. Furthermore, once an MCU is running it's hard for it to also monitor its own boot memory.

Field Programmable Gate Arrays (FPGAs) offers a significant advantage relative to MCUs. Often an FPGA is used as the first device to power up and coordinate system boot and the last device to shut down after coordinating system shut down. This position as the first-on, last-off device makes them ideal for simply establishing RoT. Designers can exploit the parallel nature of FPGAs to check multiple memories in parallel, which can lead to significant boot time improvements. And unlike MCUs, FPGAs can protect non-volatile storage by providing real-time monitoring. Lastly, they provide the logic and interfaces necessary to enable firmware recovery in case of system corruption.

Lattice MachXO3D – The Root-of-Trust FPGA for Comprehensive Hardware Security

To address this growing need for firmware security across a wide range of applications, Lattice recently announced the MachXO3D FPGA, the first small, low power FPGA for system control applications designed to secure system firmware across a wide range of applications including computing, communications, industrial control and automotive. This new device helps OEMs protect against data theft, data modification, design theft, product cloning, overbuilding, device tampering and hijacking by simplifying the implementation of a comprehensive, flexible and robust hardware security system across the product lifecycle.

Pin compatible with Lattice's popular MachXO3 devices used in a wide variety of control PLD applications, the MachXO3D establishes Lattice's product line as the control PLD of choice for secure firmware applications.

Simplifying Integration

Ensuring the simple implementation of firmware security was a high priority in the design of the MachXO3D. Lattice's designers wanted to ensure that developers could easily take advantage of the new device. Since over 50 percent of all communications systems and servers use control PLDs based on the MachXO architecture, the new device was expressly designed to be pin compatible with the legacy architecture. That allows developers to simply retrofit or add these new security capabilities to their existing control solutions. Given the rapidly escalating demand for these new security capabilities and the established popularity of the XO architecture, it's not surprising more than five of the leading server OEMs are already working directly with Lattice on MachXO3D designs. And since developers often use a MachXO3 device as the first on and last off component, they can quickly build root-of-trust and chain-of-trust capabilities without worrying whether some other digital part lies ahead of their Lattice PLD.

Enables Flexible Security and Maintains System Integrity

- MachXO architecture used to develop the majority of control PLDs for critical infrastructure

- MachXO3 and MachXO3D are pin compatible

- MachXO3D is platform's first-on/ last-off device for simple chain of trust implementation

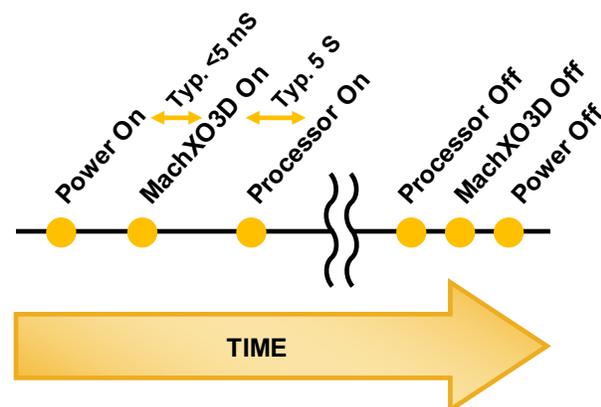- 5+ server OEMs working on MachXO3D designs

Figure 1

Comprehensive Security

To address growing hardware security concerns, Lattice augmented the MachXO3D FPGA's control PLD functionality with an Embedded Security Block that provides the hardware RoT and hardened cryptographic capabilities developers need to address multiple security threats.

- To ensure design security and protection against theft of IP, the MachXO3D features bit stream encryption.

- To protect the integrity of OEM revenue streams and brands, the device adds Secure Device ID that can be used in conjunction with other security functions to provide device / platform authentication.

- Elliptic curve cryptography, public key private key functions and AES encryption and decryption help guard against data theft.

- Elliptic curve authentication and signature generation providing the building blocks to verify authenticate firmware and general purpose data.
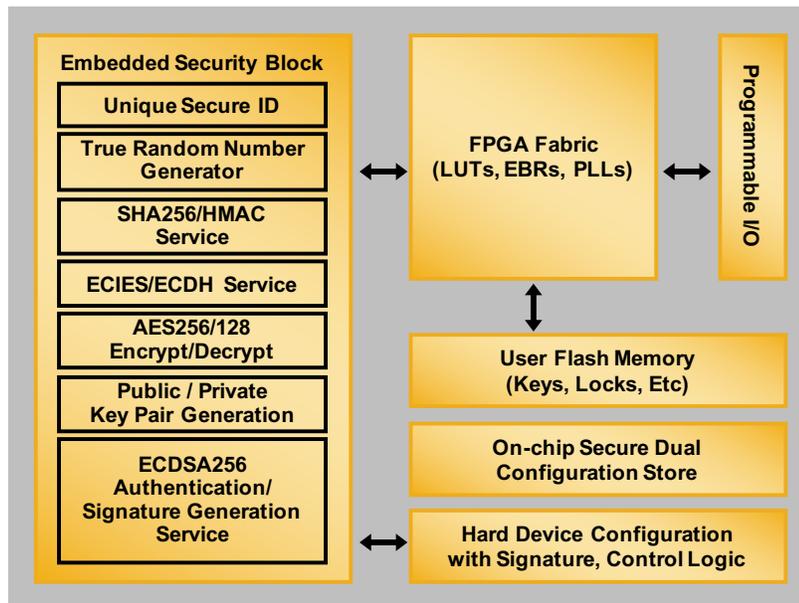

Figure 2: MachXO3D Architecture

Robust NIST Compliant Implementation

Indicative of its robust design, the MachXO3D is the industry's first control-oriented FPGA that complies with the NIST SP 800 193 Platform Firmware Resiliency (PFR) guidelines. As such, the device protects non-volatile memory through access control, cryptographically detects and prevents boot up from malicious code, and in the event of corruption, recovers to the latest trusted firmware. Moreover, the MachXO3D can dynamically reconfigure I/O ports at any time to minimize the system's attack surface.

Flexible Implementation

Design flexibility was also a key concern. Most Lattice customers appreciate the ability to deploy system enhancements to the XO architecture after equipment is deployed to the field. That reprogrammability allows for dynamic control of the attack surface. And it allows users to easily update the FPGA to render recent firmware attacks obsolete. Accordingly, Lattice's designers wanted to offer robust security capability without compromising programmability.

To meet these conflicting demands the MachXO3D adds two key features. As part of a hardened configuration engine, the device supports code authentication to confirm that each configuration loaded has the appropriate digital signature. At the same time, the MachXO3D features additional on-chip flash memory to store two configurations of the device at any time. This dual-boot capability allows the system to default to the on-device back-up configuration in the event of a compromise.

Typical Applications

The MachXO3D is designed to serve a wide range of applications across multiple markets. Potential applications include 5G wireless communications equipment like switches and routers, servers and enterprise computers, and factory automation and industrial IoT devices.

The block diagram below depicts a typical implementation of the MachXO3D in a secure server, which includes a Board Management Controller (BMC), a main CPU and a variety of additional CPUs or FPGAs. Normally a small FPGA referred to as the Control PLD manages all the resets and power supply controls for the board. All the processors boot out of SPI or Quad SPI memory. Developers can now upgrade their server security by using the MachXO3D as that small FPGA and adding a small switch. This configuration allows the FPGA to boot itself, verify each of the SPI memories and release those components to boot (assuming the memories are correct and appropriately signed). If the SPI memories are not correct, the MachXO3D takes additional steps such as shutting the system down or attempting to reconfigure from another source depending on the customer's preference. After the system has booted the MachXO3 also monitors access to the various SPI memories to prevent unauthorized writes.
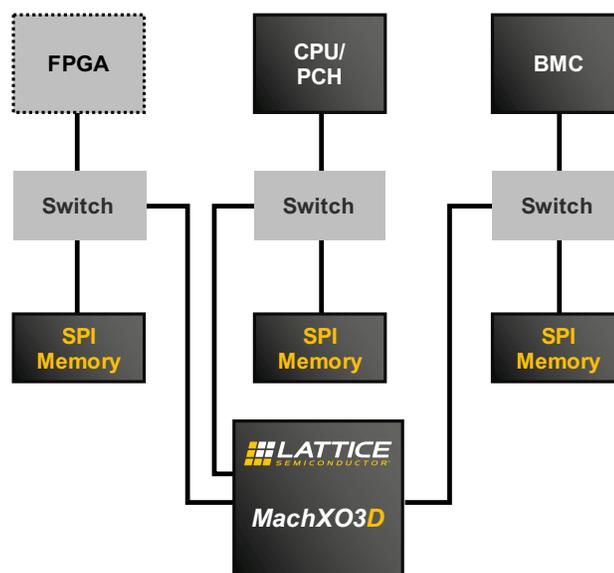

Figure 3

Security Across the Life-Cycle

To meet the growing need for increased security across the entire product lifecycle, Lattice revised its manufacturing device test integration flow to support the use of public key encryption for the secure programming of devices in an insecure environment. This ensures each device remains secure across its entire lifecycle. The new capability assures customers that their devices will remain secure from the moment they leave the Lattice factory all the way through their end-of-life.

Conclusion

Today's digital systems are under attack like never before. Hackers are exposing system vulnerabilities and using them to steal data and designs, tamper with or hijack products, or create clones. These attacks occur across the entire product lifecycle. In 2018 alone attacks on unsecured firmware left billions of ICs in computing, communications, industrial control and automotive systems vulnerable to these security risks. Ultimately, vulnerable hardware exposes OEMs to financial risk and a negative brand reputation.

How can designers mitigate this threat and protect their systems? By implementing comprehensive, flexible and robust security systems based on hardware root-of-rust and chain-of-trust techniques. Lattice's new MachXO3D FPGA lets designers enhance secure control functionality with hardware root-of-trust and dual boot capabilities. At the same time, the new device promises to dramatically simplify implementation of security solutions that cover the entire component lifecycle.