WIND RIVER

The Intelligence in the Internet of Things

SECURITY IN THE INTERNET OF THINGS

Lessons from the Past for the Connected Future

By AJ Shipley, Senior Director, Security Solutions, Wind River

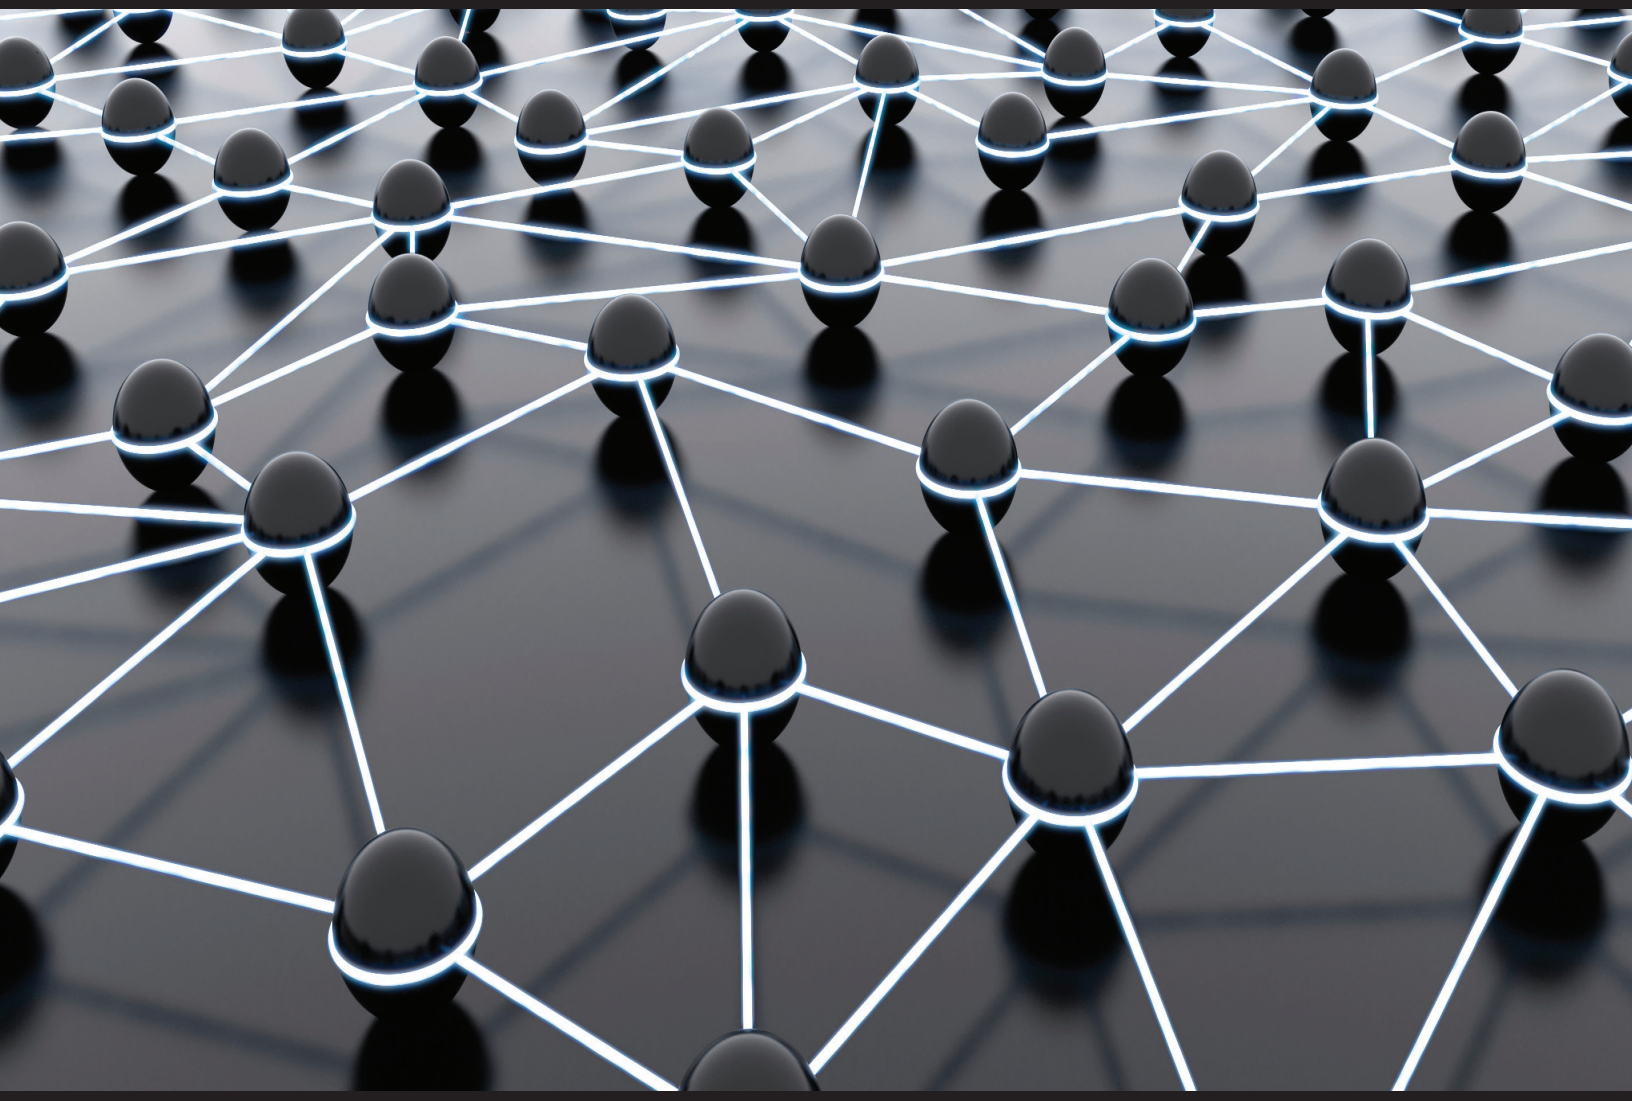

EXECUTIVE SUMMARY

Although it has been with us in some form and under different names for many years, the Internet of Things (IoT) is suddenly *the thing*. The ability to connect, communicate with, and remotely manage an incalculable number of networked, automated devices via the Internet is becoming pervasive, from the factory floor to the hospital operating room to the residential basement.

The transition from closed networks to enterprise IT networks to the public Internet is accelerating at an alarming pace—and justly raising alarms about security. As we become increasingly reliant on intelligent, interconnected devices in every aspect of our lives, how do we protect potentially billions of them from intrusions and interference that could compromise personal privacy or threaten public safety?

As a global leader in embedded technology solutions, Wind River® has been deeply involved since its inception in securing devices that perform life-critical functions and comply with stringent regulatory requirements. This paper examines the constraints and security challenges posed by IoT connected devices, and the Wind River approach to addressing them.

SEARCHING FOR THE SILVER BULLET

As every player with a stake in IoT is well aware, security is paramount for the safe and reliable operation of IoT connected devices. It is, in fact, the foundational enabler of IoT.

Where there is less consensus is how best to implement security in IoT at the device, network, and system levels. Network firewalls and protocols can manage the high-level traffic coursing through the Internet, but how do we protect deeply embedded endpoint devices that usually have a very specific, defined mission with limited resources available to accomplish it? Given the novelty of IoT and the pace of innovation today, there seems to be a general expectation that some entirely new, revolutionary security solution will emerge that is uniquely tailored to IoT—that we can somehow compress 25 years of security evolution into the tight time frame in which next-generation devices will be delivered to market.

Unfortunately, there is no "silver bullet" that can effectively mitigate every possible cyberthreat. The good news, though, is that tried-and-true IT security controls that have evolved over the past 25 years can be just as effective for IoT—provided we can adapt them to the unique constraints of the embedded devices that will increasingly comprise networks of the future.

HOW WE GOT HERE: THE EVOLUTION OF NETWORK SECURITY

Protection of data has been an issue ever since the first two computers were connected to each other. With the commercialization of the Internet, security concerns expanded to cover personal privacy, financial transactions, and the threat of cybertheft. In IoT, security is inseparable from safety. Whether accidental or malicious, interference with the controls of a pacemaker, a car, or a nuclear reactor poses a threat to human life.

Security controls have evolved in parallel to network evolution, from the first packet-filtering firewalls in the late 1980s to more sophisticated protocol- and application-aware firewalls, intrusion detection and prevention systems (IDS/IPS), and security incident and event management (SIEM) solutions. These controls attempted to keep malicious activity off of corporate networks and detect them if they did gain access. If malware managed to breach a firewall, antivirus techniques based on signature matching and blacklisting would kick in to identify and remedy the problem.

Later, as the universe of malware expanded and techniques for avoiding detection advanced, whitelisting techniques started replacing blacklisting. Similarly, as more devices started coming onto corporate networks, various access control systems were developed to authenticate both the devices and the users sitting behind them, and to authorize those users and devices for specific actions.

More recently, concerns over the authenticity of software and the protection of intellectual property gave rise to various software verification and attestation techniques often referred to as trusted or measured boot. Finally, the confidentiality of data has always been and remains a primary concern. Controls such as virtual private networks (VPN) or physical media encryption, such as 802.11i (WPA2) or 802.1AE (MACsec), have developed to ensure the security of data in motion.

NEW THREATS, CONSTRAINTS, AND CHALLENGES

Applying these same practices or variants of them in the IoT world requires substantial reengineering to address device constraints. Blacklisting, for example, requires too much disk space to be practical for IoT applications. Embedded devices are designed for low power consumption, with a small silicon form factor, and often have limited connectivity. They typically have only as much processing capacity and memory as needed for their tasks. And they are often "headless"—that is, there isn't a human being operating them who can input authentication credentials or decide whether an application should be trusted; they must make their own judgments and decisions about whether to accept a command or execute a task.

The endless variety of IoT applications poses an equally wide variety of security challenges. For example:

- In factory floor automation, deeply embedded programmable logic controllers (PLCs) that operate robotic systems are typically integrated with the enterprise IT infrastructure. How can those PLCs be shielded from human interference while at the same time protecting the investment in the IT infrastructure and leveraging the security controls available?
- Similarly, control systems for nuclear reactors are attached to infrastructure. How can they receive software updates or security patches in a timely manner without impairing functional safety or incurring significant recertification costs every time a patch is rolled out?
- A smart meter—one which is able to send energy usage data to the utility operator for dynamic billing or real-time power grid optimization—must be able to protect that information from unauthorized usage or disclosure. Information that power usage has dropped could indicate that a home is empty, making it an ideal target for a burglary or worse.

BUILDING SECURITY IN FROM THE BOTTOM UP

Knowing no one single control is going to adequately protect a device, how do we apply what we have learned over the past 25 years to implement security in a variety of scenarios? We do so through a multi-layered approach to security that starts at the beginning when power is applied, establishes a trusted computing baseline, and anchors that trust in something immutable that cannot be tampered with.

Security must be addressed throughout the device lifecycle, from the initial design to the operational environment:

1. **Secure booting:** When power is first introduced to the device, the authenticity and integrity of the software on the device is verified using cryptographically generated digital signatures. In much the same way that a person signs a check or a legal document, a digital signature attached to the software image and verified by the device ensures that only the software that has been authorized to run on that device, and signed by the entity that authorized it, will be loaded. The foundation of trust has been established, but the device still needs protection from various run-time threats and malicious intentions.
2. **Access control:** Next, different forms of resource and access control are applied. Mandatory or role-based access controls built into the operating system limit the privileges of device components and applications so they access only the resources they need to do their jobs. If any component is compromised, access control ensures that the intruder has as minimal access to other parts of the system as possible. Device-based access control mechanisms are analogous to network-based access control systems such as Microsoft® Active Directory®: even if someone managed to steal corporate credentials to gain access to a network, compromised information would be limited to only those areas of the network authorized by those particular credentials. The principle of least privilege dictates that only the minimal access required to perform a function should be authorized in order to minimize the effectiveness of any breach of security.
3. **Device authentication:** When the device is plugged into the network, it should authenticate itself prior to receiving or transmitting data. Deeply embedded devices often do not have users sitting behind keyboards, waiting to input the credentials required to access the network. How, then, can we ensure that those devices are identified correctly prior to authorization? Just as user authentication allows a user to access a corporate network based on user name and password, machine authentication allows a device to access a network based on a similar set of credentials stored in a secure storage area.
4. **Firewalling and IPS:** The device also needs a firewall or deep packet inspection capability to control traffic that is destined to terminate at the device. Why is a host-based firewall or IPS required if network-based appliances are in place? Deeply embedded devices have unique protocols, distinct from enterprise IT protocols. For instance, the smart energy grid has its own set of protocols governing how devices talk to each other. That is why industry-specific protocol filtering and deep packet inspection capabilities are needed to identify malicious payloads hiding in non-IT protocols. The device needn't concern itself with filtering higher-level, common Internet traffic—the network appliances should take care of that—but it does need to filter the specific data destined to terminate on that device in a way that makes optimal use of the limited computational resources available.
5. **Updates and patches:** Once the device is in operation, it will start receiving hot patches and software updates. Operators need to roll out patches, and devices need to authenticate them, in a way that does not consume bandwidth or impair the functional safety of the device. It's one thing when Microsoft sends updates to Windows® users and ties up their laptops for 15 minutes. It's quite another when thousands of devices in the field are performing critical functions or services and are dependent on security patches to protect against the inevitable vulnerability that escapes into the wild. Software updates and security patches must be delivered in a way that conserves the limited bandwidth and intermittent connectivity of an embedded device and absolutely eliminates the possibility of compromising functional safety.

IT STARTS IN THE OS

Security cannot be thought of as an add-on to a device, but rather as integral to the device's reliable functioning. Software security controls need to be introduced at the operating system level, take advantage of the hardware security capabilities now entering the market, and extend up through the device stack to continuously maintain the trusted computing base. Building security in at the OS level takes the onus off device designers and developers to configure systems to mitigate threats and ensure their platforms are safe.

As a pioneer in deeply embedded operating systems, Wind River understands what it takes to ensure functional safety in trusted devices, delivering software that performs tasks on which everyday lives depend. Often the only difference between safety and security considerations is the intent behind them. Wind River is uniquely positioned to implement and deliver security for IoT because of where our products reside in the device software stack. Wind River products and solutions support secure booting with hardware roots of trust, various access control mechanisms, secure package management and software updates, firewalling and IPS, and integration with network management and event correlation products.

THE END-TO-END SECURITY SOLUTION

Security at both the device and network levels is critical to the operation of IoT. The same intelligence that enables devices to perform their tasks must also enable them to recognize and counteract threats. Fortunately, this does not require a revolutionary approach, but rather an evolution of measures that have proven successful in IT networks, adapted to the challenges of IoT and to the constraints of connected devices. Instead of searching for a solution that does not yet exist, or proposing a revolutionary approach to security, Wind River is focusing on delivering the current state-of-the-art IT security controls, optimized for the new and extremely complex embedded applications driving the Internet of Things.

WIND RIVER