Whitepaper

URGENT/11

In July 2019, Armis Inc.,[1] an enterprise IoT security company based in Palo Alto, California, announced the discovery of eleven zero-day security vulnerabilities that affect the Wind River VxWorks real-time operating system.[2] Dubbed URGENT/11, the vulnerabilities affect more than two billion devices, many of which are used in mission-critical industrial, medical, or infrastructure systems.[3] This paper explains what the vulnerabilities are, why six of them are highly critical, and how operators can protect their systems, even if updating their VxWorks devices is not an option.

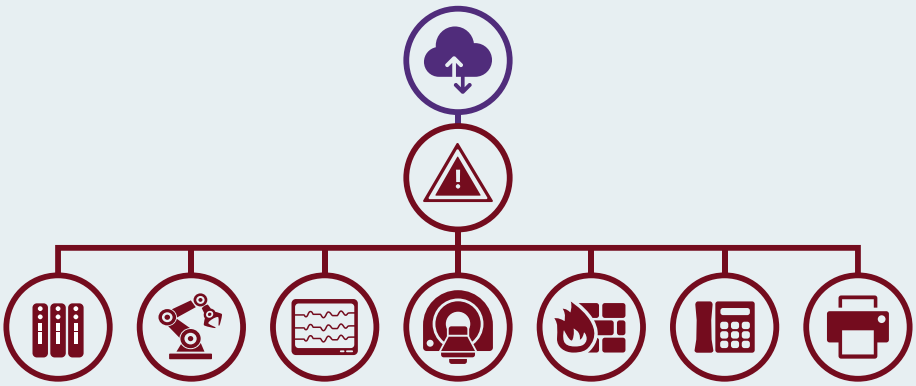

Source: ARMIS

Vulnerabilities in TCP/IP communication

Four of the vulnerabilities (CVE-2019-12255[4], CVE-2019-12260[5], CVE-2019-12261[6], and CVE- 2019-12263[7]) affect a little-known feature of the TCP/IP protocol, sending out-of-band data, also known as urgent data. Although the feature is rarely used in the real world, its implementation, consisting of an "Urgent Flag" and an "Urgent Pointer", is present in the header of every TCP packet. Exploiting these vulnerabilities does therefore not depend on any specific configuration. If a VxWorks device communicates using the TCP protocol, it is vulnerable.
It also does not matter which side initiates a TCP connection. An attacker can exploit the vulnerabilities if the VxWorks device is operated as a server that accepts TCP connections, if the VxWorks device connects to a malicious host operated by the attacker, or as a man-in-the-middle, manipulating a TCP connection between the VxWorks device and a legitimate host.

Vulnerability in source routing

Another vulnerability (CVE-2019-12256[8]) affects the IP protocol. Normally, IP packets pass multiple routers on their way from the source to the destination host, and each router looks up in a table where to send the packet next. However, the IP specification also allows the original sender to specify the routers on the way to the destination explicitly in the IP header. This is known as "Source Routing", and it is again an exotic part of the IP protocol with hardly any use in the real world.

When a device receives an erroneous IP packet, it replies with an ICMP error packet to inform the sender of this situation, and the data in erroneous packet is partially copied into ICMP error packet. VxWorks copies the data without validating it first, but when it sends the ICMP packet, it assumes its content is trustworthy, so it does not validate it at that point either. An attacker can exploit this by sending an erroneous IP packet containing specially crafted Source Routing information.

Vulnerability in DHCP client functionality

Yet another vulnerability (CVE-2019-12257[9]) affects the DHCP client functionality. When a VxWorks device obtains an IP address and further network configuration using the DHCP protocol, an attacker can exploit the vulnerability by sending a specially crafted reply packet to the device. This attack only works if the attacker is on the same LAN as the vulnerable device.

Exploitation opportunities

All of these six vulnerabilities can lead to remote code execution, i.e. an attacker can take over the device and bring it completely under his control.

Five more vulnerabilities (CVE-2019-12258[10], CVE-2019-12259[11], CVE-2019-12262[12], CVE-2019-12264[13], and CVE-2019-12265[14]) affect various network protocols: TCP, IGMP, Reverse ARP, and DHCP. These vulnerabilities can lead to denial-of-service, logical errors or information leaks, but not to remote code execution.

Remedial actions

Updating all devices running VxWorks may not be easy or possible at all. First of all, an operator has to identify all VxWorks devices, which can already be a challenging task. Updates may not yet be available for all VxWorks devices, or the device manufacturer may not provide updates at all. Even if updates are available, installing them may be too risky, or may require a lengthy recertification process. Luckily, the PHOENIX CONTACT mGuard industrial firewall can protect VxWorks devices from exploitation of all six critical vulnerabilities.

How the mGuard can help

The mGuard firewall has a feature to block any TCP packet in which the Urgent Flag is set. Enabling this feature completely protects VxWorks devices from any risk of CVE-2019-12255, CVE-2019-12260, CVE-2019-12261, and CVE-2019-12263. If the mGuard firewall is operated as a router, it terminates the TCP connection when it encounters a TCP packet with set Urgent Flag. If the mGuard firewall is operated in its stealth mode, it drops such TCP packets.

Like many routers, the mGuard firewall drops packets containing Source Routing information, thereby mitigating the risk of CVE-2019-12256. This behavior is active automatically and does not need to be configured explicitly. Within a LAN, the mGuard firewall offers the same protection when operated in stealth mode, i.e. it drops packets containing Source Routing information in this mode as well.

Finally, CVE-2019-12257 cannot be exploited from the internet, but only within a LAN. An mGuard firewall operated in stealth mode can still protect VxWorks devices from attackers within the LAN, by means of its firewall functionality. The UDP ports used by the DHCP protocol (67 and 68) can be blocked or restricted to certain IP addresses.

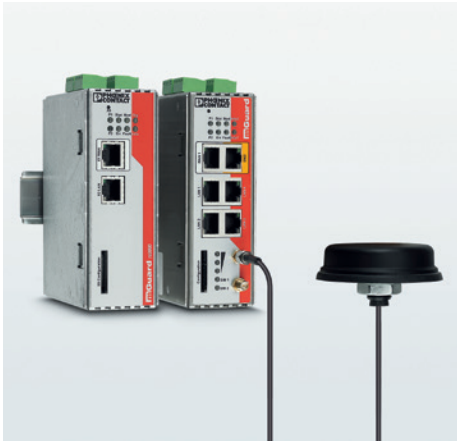

Contact person:

Peter-Jan Deltour
Product Specialist
Industrial Networks & Security
Mail: pdeltour@phoenixcontact.be
Tel.: +32 487 22 45 20

1 https://www.armis.com
2 https://www.armis.com/urgent11/
3 https://go.armis.com/urgent11
4 https://nvd.nist.gov/vuln/detail/CVE-2019-12255
5 https://nvd.nist.gov/vuln/detail/CVE-2019-12260
6 https://nvd.nist.gov/vuln/detail/CVE-2019-12261
7 https://nvd.nist.gov/vuln/detail/CVE-2019-12263
8 https://nvd.nist.gov/vuln/detail/CVE-2019-12256
9 https://nvd.nist.gov/vuln/detail/CVE-2019-12257
10 https://nvd.nist.gov/vuln/detail/CVE-2019-12258
11 https://nvd.nist.gov/vuln/detail/CVE-2019-12259
12 https://nvd.nist.gov/vuln/detail/CVE-2019-12262
13 https://nvd.nist.gov/vuln/detail/CVE-2019-12264
14 https://nvd.nist.gov/vuln/detail/CVE-2019-12265