Security Now! #661 - 05-01-18
Securing Connected Things

This week on Security Now!

This week we discuss Win10 got a new spring in its step, Microsoft further patches Intel microcode, even the UK's NHS plans to update, another hack of modern connected autos, Oracle's botched WebLogic patch, an interesting BSOD-on-demand Windows hack, a PDF credentials theft hack (which Adobe won't fix), your Echo may be listening to you, a powerful Hotel keycard hack, a bit of errata and feedback, and a discussion of another Microsoft-driven security initiative.

Our Picture of the Week

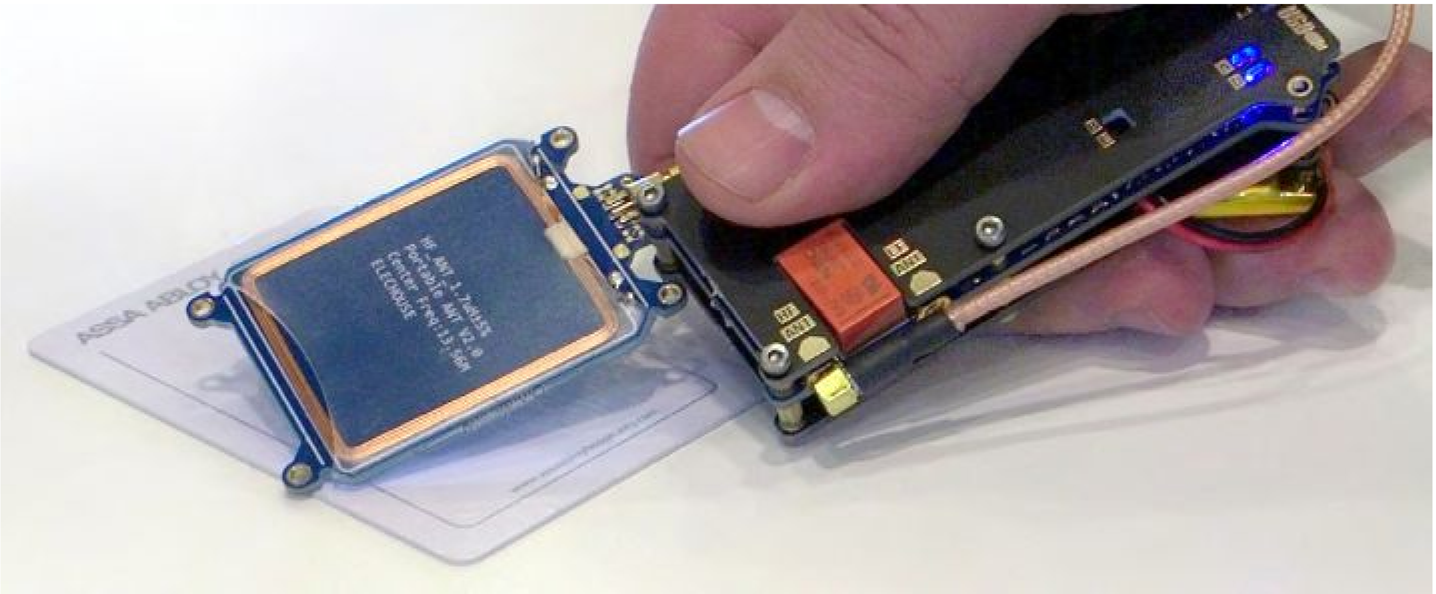

Security News

The Windows 10 April 2018 Update arrived yesterday -- what's new?
https://www.bleepingcomputer.com/news/microsoft/the-windows-10-april-2018-update-arrives-today-heres-whats-new/

Win10 version 1803

Focus Assist -- Alarms Only and 2:01am to 2:00am (always!)

TAR and CURL are now built in.

And lots more...

Microsoft updates Windows Updates with new Spectre mitigations
https://www.ghacks.net/2018/04/25/windows-10-microcode-updates-kb4090007-kb4091663-kb4091664-and-kb4091666/

- KB4090007 - Win10 v 1709 (1st release 3/13. Updated 4/24)
- KB4091663 - Win10 v 1703 (1st release 3/14. Updated 4/24)
- KB4091664 - Win10 v 1607 (1st release 3/14. Updated 4/24)
- KB4091666 - Win10 v 1507 (1st release 4/24)

Ghacks notes that:

There is no update for Windows 10 version 1511.

Yesterday's microcode updates cover most processor families that Intel wants to support with microcode updates. If you check the master list that Intel released, you will notice that some products are still missing and that some (older) processors won't receive the updates at all.

Ghacks' Martin Brinkmann: "It is likely that Microsoft will update the updates to integrate support for processor families that are not supported yet. I suggest you monitor the relevant KB article pages so that you know when updates are released."

I have a little Win10 machine with a CPUID of 406C4 - Intel has patches but it's still missing from Microsoft's patched CPUs lineup.

https://newsroom.intel.com/wp-content/uploads/sites/11/2018/04/microcode-update-guidance.pdf

And... Speaking of Windows 10...
UK Health Agency Switches to Windows 10 Citing WannaCry Ransomware Outbreak

https://www.bleepingcomputer.com/news/government/uk-health-agency-switches-to-windows-10-citing-wannacry-ransomware-outbreak/

The UK Department of Health and Social Care has announced that it will transition all National Health Service (NHS) computer systems to Windows 10.

Officials cited the operating system's more advanced security features as the primary reason for upgrading current systems, such as the SmartScreen technology included with Microsoft Edge and Windows Defender, Microsoft's sneakily good antivirus product.

WannaCry outbreak played a role

Department officials didn't ignore the elephant in the room, and also referenced the damages caused by the WannaCry ransomware outbreak last year as one of the reasons for upgrading their infrastructure.

The NHS was one of the first WannaCry victims last year, and one of the most harshly hit.

In a report published last year, NHS officials said WannaCry hit more than a third of all NHS trusts and led to the cancelation of over 6,900 medical appointments across the UK, including critical operations, albeit there was no loss of human life because of the cyber-attack.

According to Kaspersky and Microsoft telemetry, over 98 percent of all WannaCry victims were Windows 7 users.

By moving its infrastructure to Windows 10, NHS officials hope to leverage the plethora of new security features added in Windows 10 to safeguard NHS networks from similar future incidents.

The Connected Car: Ways to get unauthorized access and potential implications
https://www.computest.nl/wp-content/uploads/2018/04/connected-car-rapport.pdf

It's called the CAN bus... which is convenient, since, once you gain access, you CAN make it do anything you want.

The CAN bus appeared and became standardized during the 90's. It is now present in every automobile built since and because every component in today's autos is hooked up to this single bus, it is used to control everything in the vehicle from steering to unlocking the doors to the volume of the radio.

The CAN protocol is quite straight forward. Each message has an arbitration ID and a payload. There is no authentication, authorization, signing, or encryption. It's wide open... and once you are on the bus you can send arbitrary messages which will be received by all parties connected to the same bus. There is no sender or recipient information and each component decides for itself if a specific message applies to it.

As we often ask on this podcast: WHAT could POSSIBLY go wrong?

Naturally, any attacker who might somehow gain access to a vehicle's CAN bus would control the car. They could impersonate the front radar for example to instruct the braking system to make an emergency stop due to a near collision or take over the steering.

The attacker only needs to find a way to actually get access to a component that is connected to the CAN bus... and many, being radio-based, offer enticing targets for attack.

We've seen many local attacks. These researchers wanted to explore the possibility of Internet-based attacks.

Today's cars also typically have multiple CAN busses separating convenience systems from critical systems for safety. But most also break the rule of total isolation and incorporate CAN bus gateways to allow inter-bus interactions.

They wrote:

<quote> We started this research with nine different models, from nine different brands. These were all lease cars belonging to employees of Computest. Since we are not the actual owner of the car we asked permission for conducting this research beforehand from both our lease company and the employee driving the car.

We conducted a preliminary research in which we mapped the possible attack vectors of each car. Determining the attack vectors was done by looking at the architecture, reading public documentation and by a short technical review.

Things we were specifically searching for:

- cars with only a single or few layers between the cellular connection and the high-speed CAN bus;
- cars which allowed us to easily swap SIM cards (since we are not the owner of the cars, soldering, decapping etc. is undesirable);
- cars that offered a lot of services over cellular or Wi-Fi.

From here we choose the car which we thought would give us the highest chance of success. This is of course subjective and does not guarantee success. For some models getting initial access might be easier than others, but this does say nothing about the effort required for lateral movement.

We finally settled for the Volkswagen Golf GTE as our primary target. We later added the Audi A3 e-tron to our research. Both vehicles share the same IVI-system which, on first sight, seemed to have a broad attack surface, increasing the chance of finding an exploitable vulnerability.

The MIB version installed in the Volkswagen Golf has the possibility to connect to a Wi-Fi network.

A quick port scan on this IP shows that there are many services listening:

```
$ nmap -sV -vvv -oA gte -Pn -p- 192.168.88.253
Starting Nmap 7.31 ( https://nmap.org ) at 2017-01-05 10:34 CET

Host is up, received user-set (0.0061s latency).
Not shown: 65522 closed ports
Reason: 65522 conn-refused
PORT STATE SERVICE REASON VERSION
23/tcp open telnet syn-ack Openwall GNU/*/Linux telnetd
10123/tcp open unknown syn-ack
15001/tcp open unknown syn-ack
21002/tcp open unknown syn-ack
21200/tcp open unknown syn-ack
22111/tcp open tcpwrapped syn-ack
22222/tcp open easyengine? syn-ack
23100/tcp open unknown syn-ack
23101/tcp open unknown syn-ack
25010/tcp open unknown syn-ack
30001/tcp open pago-services1? syn-ack
32111/tcp open unknown syn-ack
49152/tcp open unknown syn-ack
Nmap done: 1 IP address (1 host up) scanned in 259.12 seconds
```

Port 49152 has a UPnP service listening.

After further research, we found a service on the Golf with an exploitable vulnerability. Initially we could use this vulnerability to read arbitrary files from disk, but quickly could expand our possibilities into full remote code execution. This attack only worked via the Wi-Fi hotspot, so the impact was limited. You have to be near the car and it must connect with the Wi-Fi network of the attacker. But we did have initial access:

```
$ ./exploit 192.168.88.253
[+] going to exploit 192.168.88.253
[+] system seems vulnerable...
[+] enjoy your shell:
uname -a
QNX mmx 6.5.0 2014/12/18-14:41:09EST nVidia_Tegra2(T30)_Boards armle
```

8. Conclusions

Internet-connected cars are rapidly becoming the norm. As with many other developments, it’s a good idea to sometimes take a step back and evaluate the risks of the path we’ve taken, and whether course adjustments are needed. That’s why we decided to pay attention to the risks related to internet-connected cars. We set out to find a remotely exploitable vulnerability, which required no user interaction, in a modern-day vehicle and from there influence either driving behavior or a safety feature.

With our research, we have shown that at least the first is possible. We can remotely compromise the MIB IVI system and from there send arbitrary CAN messages on the IVI CAN bus. As a result, we can control the central screen, speakers and microphone. This is a level of access that no attacker should be able to achieve. However, it does not directly affect driving behavior or any safety system due to the CAN gateway. The gateway is specifically designed to firewall CAN messages and the bus the IVI is connected to is separated from all other components. Further research on the security of the gateway was consciously not pursued.

We argue that the threat of an adversary with malicious intent was long underestimated. The vulnerability we initially identified should have been found during a proper security test. During our meeting with Volkswagen, we had the impression that the reported vulnerability and especially our approach was still unknown. We understood in our meeting with Volkwagen that, despite it being used in tens of millions of vehicles world-wide, this specific IVI system did not undergo a formal security test and the vulnerability was still unknown to them. However, in their feedback for this paper Volkswagen stated that they already knew about this vulnerability.

Oracle WebLogic Servers in the crosshairs:
Hackers are scanning the web for still-vulnerable Oracle WebLogic Servers... after Oracle's patch failed to fix the entire problem.

We previously covered this highly critical JAVA deserialization flaw which was found late last year by a security researcher. The flaw exists in Oracle's WebLogic Server of its Fusion Middleware and allows attackers to gain complete control over any vulnerable Oracle servers... over the Internet through TCP port 7001.

In Java, "serialization" and "deserialization" is the process of packing up for storage or transmission and then unpacking after retrieval or receipt, Java's structured storage objects. Not surprisingly, the Java deserializer is an interpreter... and interpreters are notoriously difficult to get correct.

So earlier this month, Oracle patched this highly critical remote code execution vulnerability, but apparently not completely.

Now, a security researcher, who is @pyn3rd on Twitter, claiming to be part of the Alibaba security team, has found a way to remotely bypass the security patch and exploit the WebLogic vulnerability once again.

With any luck, the only thing attackers will do is mine some coin. WebLogic Server admins should consider themselves lucky.

Windows 10 NTFS Crash DoS
https://github.com/mtivadar/windows10_ntfs_crash_dos
https://github.com/mtivadar/windows10_ntfs_crash_dos/blob/master/doc/ntfs_crash.pdf

So... a smallish (10mb) NT Filesystem image was hand-crafted which causes Windows to barf and collapse horribly when it attempts to mount the file system.

Microsoft replied that (and I'm paraphrasing here), while yes, having Windows barfing and collapsing horribly is always a bit embarrassing, the attack, such as it is, as presented, requires physical access to the machine and for connected drives to be automatically enumerated and mounted.

For his part, the discoverer of this problem was mostly concerned that this could be done even to a machine that was in a locked state.

What's of potentially greater interest is what more someone who works to further exploit whatever-is-going-on might be able to do.

"Interpreters are incredibly difficult to get right."

NTLM Credentials Theft via PDF Files
This all dates back to early decisions made for the IBM/Microsoft LAN Manager system.

The major weaknesses of LAN Manager authentication protocol are:

- Passwords are not case sensitive. All passwords are first converted into uppercase before hashing.
- Password characters are limited to a subset of the ASCII character set.
- Passwords are limited to a maximum of 14 characters.
- And... is everyone sitting down for this??... A 14-character password is broken into two 7-character substrings, EACH hashed separately!
- And, (you really cannot make this up) if the password is 7 characters or less, the second half's "null hash" is always (0xAAD3B435B51404EE). Therefore, if the length of password is less than or equal to 7 characters, then a password length of 7 characters or less can be identified visibly without using tools.
- Oh... and note that that was 16 Hex characters, which is 8 bytes, which is 64 bits. The LAN Man hashes are 64 bits long.

The hash value is sent over networks to authenticate their users, without salting, making it susceptible to man-in-the-middle attacks such as replay the hash.

So...

Our personal computing history is chock full of creative exploits and ways that a very weakly protected user authentication can be leveraged... and now we have another...

https://research.checkpoint.com/ntlm-credentials-theft-via-pdf-files/

Check Point Research writes: A few days after it was reported that malicious actors could exploit a vulnerability in MS outlook using OLE to steal a Windows user's NTLM hashes, the Check Point research team can also reveal that NTLM hash leaks can also be achieved via PDF files with no user interaction or exploitation.

According to Check Point researchers, rather than exploiting the vulnerability in Microsoft Word files or Outlook's handling of RTF files, attackers take advantage of a feature that allows embedding REMOTE documents and files inside a PDF file. The attacker can then use this to inject malicious content into a PDF and so when that PDF is opened, the target automatically leaks credentials in the form of NTLM hashes.

(The user's NTLM credential is famously sent out with the request for a remote document -- over the Internet to wherever -- in order to authenticate the user's authorization to receive the remote resource.)

Check Point believes that any and all Windows PDF viewers can be so induced to exfiltrate the user's NTLM credentials in this fashion.

And Adobe, always quick off the starting line with security and home of Adobe FLASH responded to Check Point's private disclosure, saying: "Thank you for checking in on this case. Microsoft issued an OPTIONAL security enhancement late last year that provides customers with the ability to disable NTLM SSO authentication as a method for public resources. With this mitigation available to customers, we are not planning to make changes in Acrobat."

<sarcasm> "Right... because, after all... we're still publishing FLASH, so we obviously don't care about you at all." </sarcasm>

https://portal.msrc.microsoft.com/en-US/security-guidance/advisory/ADV170014

Amazon's Echo turned into a persistent eavesdropper
"She's got skillz!"

https://info.checkmarx.com/wp-alexa
https://thehackernews.com/2018/04/amazon-alexa-hacking-skill.html

As we know, Amazon's "EchoSystem" supports the use of add-on "Skills" for their devices with an API used to teach her new tricks.

Security researchers at the cybersecurity firm Checkmarx created a proof-of-concept voice-driven 'skill' for Alexa that, once launched, spins off a secondary background non-terminating task which causes any skill-enhanced Amazon voice assist device to indefinitely record all surrounding audio, eavesdropping on users' conversations, while also exfiltrating all the audio to a third-party website.

Whoopsie!

Checkmarx reported the issue to Amazon, and the company has addressed the problem, though not very robustly in my opinion: They now regularly scan for malicious skills that "silent prompts or that listens for unusual lengths of time", then kicks them out of their official store.

This feels like an immature and not-well-thought-out API. No skills should have such any access to the system's microphone. Rather, any skill should register its microphone usage and access policy and needs as part of an unhackable policy packet. Then the device's API should grant the skill's requests -- filtered through and enforced by the skill's previously declared and approved policy.

F-Secure reverse engineers the most popular Hotel card-key locking system...
... with devastating effect
https://www.f-secure.com/en/web/business_global/electronic-lock-systems-are-vulnerable
https://www.f-secure.com/documents/10192/406831/f-secure-electronic-lock-systems-are-vulnerable-faq-en.pdf

Device Can Generate Master Keys From Valid or Expired Hotel Keys

"Ghost in the Locks"

The design flaws discovered in the smart lock system's software -- known as "Vision" by VingCard -- and used to secure millions of hotel rooms worldwide -- prompted the world's largest lock manufacturer, Assa Abloy, to issue software updates with security fixes to mitigate the issue.

The researchers simulated the attack with an ordinary electronic key to the target facility. Using information on the key, they were able to create a master key that can open any door using the same lock system in the facility. The key doesn't even have to be a working key. Even one that's long expired, discarded, or is used to access spaces such as a garage or closet could be used. The attack can be performed without being noticed.

Sweden's Assa Abloy is the world's largest lock manufacturer and "Assa Abloy Hospitality" is the division of the company that provides lock systems for hotels, cruise ships and other industries.

However, the "Vision" software which is impacted by the attack is used in hotels and cruise ships only.

https://www.assaabloyhospitality.com/en/aah/com/solutions/

In F-Secure's FAQ about their successful exploit:

WHY DID THE RESEARCH TAKE OVER A DECADE TO COMPLETE?

Figuring out the complexities of how the lock system, software and keys worked was very complicated. Building, and breaking, an electronic access control system is very difficult because there are many facets to get right. Assa Abloy is a highly reputed lock manufacturer and aside from the seemingly innocuous security oversights in the software, their products are well designed.

These security oversights were not gaping obvious holes. It took a thorough understanding of the design of the whole system to be able to identify small flaws in the system. The researchers then creatively combined these flaws to produce the attack.

The takeaway for our listeners: NEVER (never never never) assume that the door of your hotel room is locked. (Never never.) So... inconvenient as it is, you must =ALWAYS= place valuables which will fit into the room's safe in the safe, carry laptops what will not fit away with you, or turn them over for safer-keeping to the front desk and their hopefully larger safe.

Errata

Microsoft's Azure Sphere -- not QUITE as much "choice" as I stated last week:

"Azure Sphere gives you choice. You can connect data from any cloud, proprietary or public, or even to your on-prem infrastructure to the Azure Sphere Security Service."

"Add data and telemetry" is what can be connected to non-Microsoft services.

But the use of Microsoft's "Azure Sphere Security Service" appears to be an inseparable component of their system.

And, assuming that they price it all affordably, this offloads ALL of the responsibility of maintaining the security and integrity of IoT devices from hardware manufacturers who just want to crank out hardware without investing in huge after market maintenance and security.

SpinRite

Lachlan Gabb in Sydney, Australia
Subject: Another SpinRite Success Story
Date: 27 Apr 2018 22:14:29

Hi Steve,

Just wanted to share another SpinRite success story. I am a security analyst living in Sydney Australia and also the go to guy for computer problems within my family and close neighbours. A couple of weeks ago, one of my neighbours called me to take a look at their computer that had been running very slow recently.

Before I had even arrived to take a look at the system, it had blue screened and was now refusing to boot at all, stating that no operating system could be found. I asked, but predictably there was no backup and the computer contained important photos of their grandchildren and documents dating back over five years.

I removed the drive and connected it to my computer to see if I could recover any data, however Windows would not recognise the drive at all. Being an avid Security Now listener I immediately thought of SpinRite, which I had been looking for an excuse to try out for a couple of months anyway. I went over and purchased a copy and let it run on the drive overnight. I think we know how the rest of this goes.

In the morning, not only could I copy all the data from the drive, but it would even boot again. I promptly copied the data to a fresh drive and reinstalled it in the computer, along with a much needed lecture on the importance of backups. I was given immense gratitude and even a container of freshly baked cookies for saving the system. I wanted to pass on the gratitude to you, but I ate the cookies myself.

Thanks for the wonderful product and i'm looking forward to the next episode of Security Now.
Kind regards from Down Under,
Lachlan

Closing The Loop

Nathan Boeger in Tampa, FL / Subject: Drupal updates / Date: 01 May 2018 06:37:15

Steve: What are your thoughts on Drupal developers or a 3rd party patching publically accessible Drupal servers with open remote code execution vulnerabilities? Are you aware of any legal precedent? A "do-gooder" could write a worm that propagates and patches. I'm guessing that person would be bearing risk without much possible benefit.

2nd question - are you aware of Malware that drops an agent then patches the server behind itself? It seems like the natural progression for Drupalgeddon malware, particularly if they have an out of band mechanism for persistent C2.
Thanks! A regular listener.
Nathan Boeger

Securing Connected Devices

Microsoft's IoT and Industrial Control security project:
Trusted Cyber Physical Systems (TCPS)

https://az835927.vo.msecnd.net/sites/iot/Resources/documents/TCPS-WP.pdf

Executive summary

Attacks against critical infrastructure, like the Triton breach of critical safety systems of an energy plant, have become more frequent globally. While the financial toll of cyber-attacks has become an unfortunate norm, attacks that can damage connected systems, and put human lives and property at risk are emerging with new regularity. This document describes how cyber-physical systems (CPSs), also known as Internet-of-Things (IoT) in an industrial context, can be securely controlled, monitored, and audited throughout the IoT infrastructure, including cloud services, compute devices, and microcontrollers (MCUs), down to the pins that provide power to, e.g., open and close a valve in a water plant. The key to this approach is that all actions and messages to and from a CPS device, all the way down to the hardware I/O pin, are cryptographically secured. Even if the OS on a cyber-physical system itself is compromised, an attacker will not be able to operate the valve nor tamper with the activity log of the valve. Furthermore, even the OS vendor or ISV cannot access private data nor send unauthorized commands or software updates. This will allow a clear separation between the authorized operators of solutions and the software vendors, hardware vendors and solution providers.

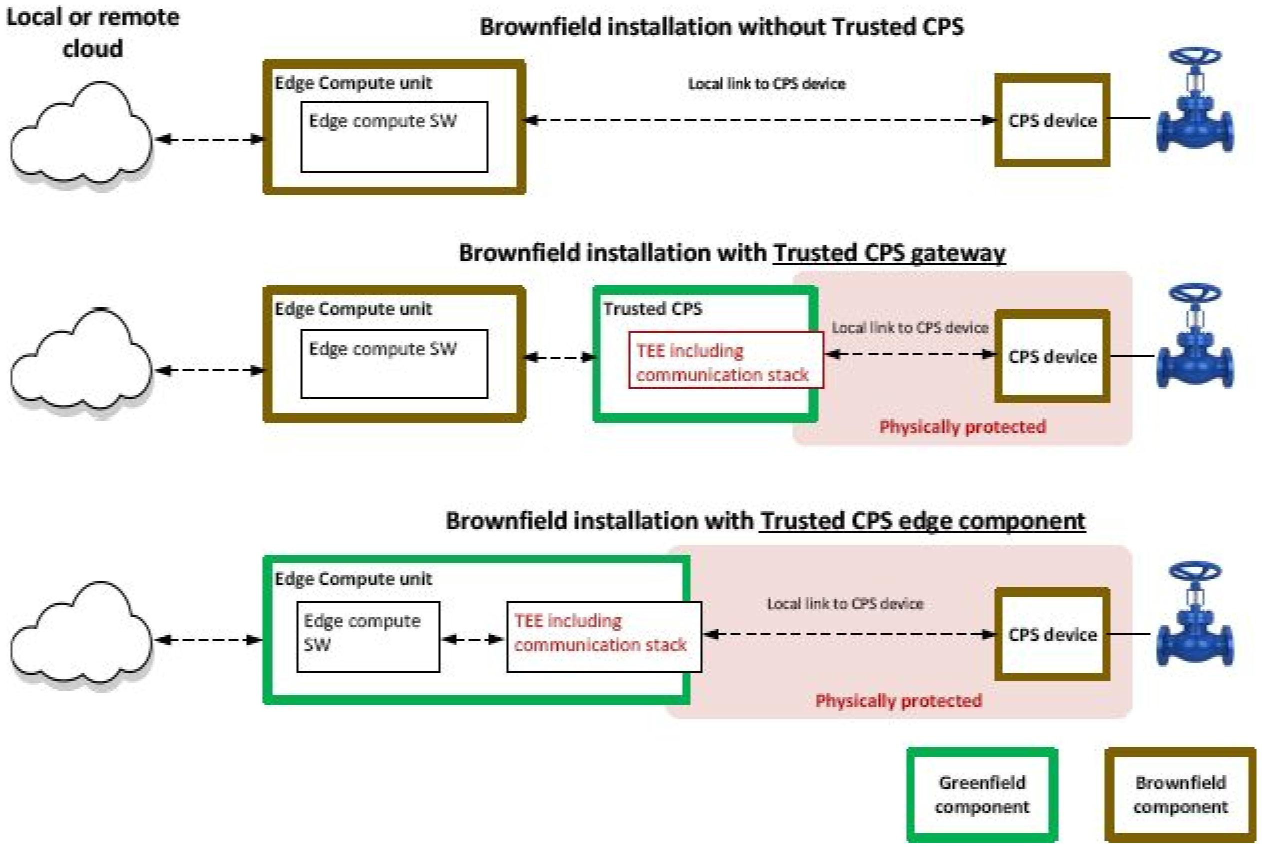

Local or remote cloud
Brownfield installation without Trusted CPS
Edge Compute unit
Edge compute SW
Local link to CPS device
CPS device
Brownfield installation with Trusted CPS gateway
Edge Compute unit
Edge compute SW
Trusted CPS
TEE including communication stack
Local link to CPS device
CPS device
Physically protected
Brownfield installation with Trusted CPS edge component
Edge Compute unit
Edge compute SW
TEE including communication stack
Local link to CPS device
CPS device
Physically protected
Greenfield component
Brownfield component