DRAFT — FOR EXTERNAL REVIEW

Risk Assessment

Intel ME / vPro Hardware Connected to Corporate Resources Without IT-Controlled Interdiction

Bring-Your-Own-Device and IT-Supplied (Management-Engine-Naive) Procurement Cases

Document title	Risk Assessment: Intel ME / vPro Hardware Connected to Corporate Resources Without IT-Controlled Interdiction
Version	0.9 — Draft for External Review
Date	25 April 2026
Status	Draft for external review. Not approved for distribution outside the named review group.
Prepared for	Chief Information Security Officer
Prepared by	Information Security Architecture
Classification	DRAFT — FOR EXTERNAL REVIEW
Audience	Chief Information Security Officer and supporting security architecture function. External technical reviewers nominated by the CISO.

Document control

This document is a working draft for external review. Comments and proposed amendments should be returned to Information Security Architecture by the date set out in the review brief. The recommendation in Section 17 is conditional on the review process and on the open questions in Appendix B; it is not an organisational decision until accepted by the Chief Information Security Officer.

Version history

Version	Date	Author	Change summary
0.9	25 April 2026	Information Security Architecture	Initial draft for external review. Covers BYOD and IT-supplied (ME-naive) procurement cases. Evidence base across CSME and AMT vulnerability record from CVE-2017-5689 to CVE-2025-20037 plus the INTEL-SA-01138 / 01200 / 01280 cluster. Cisco hardcoded-trust parallel included. Firmware update path constraints documented. Other-radio compromise paths included.

Distribution

This draft is distributed to the named external technical reviewers nominated by the CISO and to the internal security architecture function. It is not for wider distribution while in draft. Final distribution will be set out in the version 1.0 release once external review is incorporated.

Open items at time of issue

- Appendix B records open questions for the reviewer. The two material items at issue are the specific 2026 CSME advisory identifier (the document uses CVE-2025-20037 and the INTEL-SA-01138 / 01200 / 01280 cluster as the most recent public examples), and the fleet-specific hardware generation that determines exactly which evidence-base entries apply most directly.
- Formal scoring (NIST 800-30 likelihood-times-impact, CVSS v3.1 / v4.0 columns) is intentionally absent from the body. Where the sign-off pack requires formal scoring, an annex can be added without re-litigating the technical content.
- Programme cost sizing for the recommended controls is out of scope and is treated as a follow-on costing exercise referenced in Appendix B.

Contents

1. Executive summary

This document evaluates the security posture of Intel vPro laptops whose Management Engine has not been interdicted by the organisation's IT function before connection to corporate resources. Two procurement-provenance cases are in scope and are treated together because they share the same underlying threat. The first case is bring-your-own-device (BYOD), where the user purchased the laptop through any retail or self-funded channel and the organisation has only host-level constraints, including any combination of mobile device management enrolment, conditional access, virtual desktop infrastructure, browser isolation and VPN posture checks. The second case is IT-supplied hardware in environments where the IT team has no procedural awareness of the ME, ships factory-state devices directly to users and treats the ME as if it did not exist.

The conclusion the evidence supports is that connecting an untouched-ME vPro laptop to corporate resources in either case exposes the organisation to a class of compromise that defeats the host security stack in its entirety. The exposed controls include BitLocker full-disk encryption, FIDO2-protected sign-in, endpoint detection and response, the host firewall and the corporate VPN. The defeat occurs because the threat operates below the operating system, in a runtime (the CSME) that the host cannot see, log or contain.

The risk is not theoretical. The PLATINUM nation-state-aligned actor was documented by Microsoft in 2017 using Intel AMT Serial-over-LAN as a covert exfiltration channel that bypassed the Windows host firewall and host monitoring entirely, because the SOL channel is implemented in the CSME and the host TCP/IP stack does not see it (Microsoft Security Blog; Threatpost; The Hacker News). Section 7 sets out the inline cohabitation property under which the same blind spot extends from the host firewall to the corporate perimeter firewall.

The mitigation paths available differ by case. In the IT-supplied / ME-naive case, the answer is procedural: train the IT function, build a small interdiction lab, change procurement to require firmware integrity attestation at receipt and bring the existing fleet into the lab on a controlled cycle. In the BYOD case, no interdiction is available to the organisation, and the answer is access-class constraint at the corporate perimeter: refuse high-sensitivity access classes from BYOD vPro hardware; allow lower-sensitivity access classes only via VDI or browser isolation; document the residual; treat the device as in a permanently untrusted population. Both cases share an irreducible Tier-3 supply-chain residual that no software-only control closes. The recommendation is in Section 17.

2. Trust assumptions this assessment refuses to make

The assessment proceeds by naming the assumptions it will not take on faith. Each is examined in the body. The reader can identify, at any step, the assumption being relied upon and decide whether they accept it.

Hardware integrity in transit. A laptop that has been outside the organisation's chain of custody cannot be assumed to have zero injected Wi-Fi profiles, an untouched MEBx password, an unmodified root certificate hash list or unmodified CSME firmware regions. For the IT-supplied case, the chain of custody runs from contract manufacturer through reseller, courier and customs to the user's desk; for the BYOD case, no organisational chain of custody exists at all. Supply-chain interdiction by a state actor is documented operating doctrine; supply-chain compromise via a paid insider or warehouse handler is documented historical fact. Section 9.6 sets out the public record.

Applet isolation inside the CSME. AMT and the firmware Trusted Platform Module (PTT) run in the same firmware image, on the same microcontroller, sharing the same kernel. The boundary between them is software-enforced. The public record, in particular the Positive Technologies work on CVE-2019-0090, shows the underlying trust foundation has been broken in ways Intel could not patch in the field.

The cryptographic gate at provisioning. The factory root certificate trust list contains hashes for a small number of commercial certificate authorities. CA mis-issuance and CA compromise are recurring documented events; the bar for an attacker is not "compromise impossible cryptography" but obtain mis-issuance from one of approximately a dozen public CAs, where the precise number depends on firmware version and is established by reverse-engineering rather than published list. The list itself is editable through MEBx, which collapses the question to physical access during transit.

A patched advisory closes the surface. A patched memory-corruption advisory in a Ring -3 (sub-OS) network parser is evidence that the surface produces this class of bug and that it remains in production hardware where the next bug has not been disclosed. A reasonable security officer does not gain confidence from "the bug we know about is fixed"; they lose confidence from this surface keeps producing bugs of this severity, across architectures and hardware generations.

A "Denial of Service" classification means DoS only. A reachable, unauthenticated memory corruption in the CSME network parser is, at minimum, an exploit primitive an advanced persistent threat will refine into remote code execution given time and incentive. The vendor classifies what they could prove in a disclosure window.

The HAP "Reserved" bit disables the ME. It does not. It disables most ME functions after the bring-up phase; the CSME runtime continues to exist, because the chipset boot depends on it (Positive Technologies, "Disabling Intel ME 11 via undocumented mode"). Setting it is a meaningful reduction of attack surface, not an elimination of the runtime.

An unprovisioned CSME on battery is electrically silent. Modern thin-and-light laptops in their default Modern Standby (S0ix) state continue to draw small but non-zero power for weeks while ostensibly off, which is the user-visible evidence that low-power listening states are physically maintainable on these platforms. Section 8 covers the publicly-reported behaviour. Whether the listening state is configured to receive AMT traffic on a given device is a firmware policy question, and firmware policy is what the supply chain can change.

The host firewall, host EDR or perimeter network controls can see ME traffic. The host firewall cannot, because the NIC manageability filter delivers AMT traffic to the CSME before the host TCP/IP stack runs. The corporate perimeter firewall cannot disambiguate ME traffic from host traffic at the flow level either, because the CSME shares the host's MAC and IP. Section 7 sets out the implications and the corollary that L3 filters cannot isolate ME-originated traffic on a per-flow basis.

The organisation can apply CSME firmware updates on demand. It generally cannot, in either provenance case. Section 11 sets out the constraints, which include OEM gating of capsule availability, generation-dependent capsule support, OEM signing constraints, downgrade blocking and, for BYOD, the absence of any push authority altogether.

3. Architecture under assessment, by hardware generation

A current-generation Intel vPro laptop ships with three subsystems that operate independently of, and below, the host OS.

The **Converged Security and Management Engine** is a separate microcontroller embedded in the Platform Controller Hub. The runtime has changed across hardware generations and the change matters for the threat assessment, because public attack research is grouped by generation.

Pre-Skylake Intel ME (versions 1.x to 10.x) ran on an ARC core. ME 1.x to 5.x used ARCTangent-A4; ME 6.x to 10.x used ARCompact. The runtime was Express Logic ThreadX, augmented with a Java-style applet engine. The canonical public reverse-engineering work is Igor Skochinsky's two presentations at REcon 2014 and Breakpoint 2014 (REcon 2014, "Intel ME Secrets" PDF; Breakpoint 2014, "Intel ME, Two Years Later" PDF). The Skochinsky work documents the boot ROM verification flow, the FTPR partition layout and the SPI-flash organisation. CVE-2017-5689 spans this generation and the early MINIX-era hardware that followed.

Skylake and later Intel ME (11.x onwards) switched to an Intel Quark x86 32-bit core running an Intel-modified microkernel derived from MINIX 3 (Intel Management Engine, Wikipedia). The most influential public attack work on this generation is Goryachy and Ermolov's 2017 demonstration of unsigned code execution on the CSME (Black Hat Europe 2017, "How to Hack a Turned-Off Computer" PDF). INTEL-SA-00086, CVE-2019-0090, CVE-2021-0146 and the ongoing 2024 to 2026 advisory cadence apply on this generation.

Newer ME or CSE generations (the 15, 16, 17 and 18 family) continue the Quark plus modified-MINIX heritage with incremental hardening, including UMA-region protection added after Positive Technologies demonstrated UMA-page injection, and the deprecation of SHA-1 root certificate hashes for AMT provisioning in ME 16. The advisory cadence has not stopped. CVE-2025-20037 affects CSME firmware before 14.1.77.2497 and is representative of the most recent disclosures in this family (ZeroPath summary).

The cross-generation point is worth stating once. The runtime moved from ThreadX-on-ARC to MINIX-on-Quark; the shape of the attack surface did not. Both architectures have produced exploitable vulnerabilities of the same classes, including unpatchable boot-ROM issues, applet boundary failures, network-stack memory corruption and debug-logic activation flaws.

The CSME has its own memory, its own clock, its own access to the platform NIC over a sideband bus and DMA access to system DRAM. Code running on the CSME is referred to in the literature as Ring -3 (Tereshkin and Wojtczuk, "Introducing Ring -3 Rootkits", Black Hat USA 2009 PDF).

Intel Active Management Technology is an applet that runs inside the CSME. It provides out-of-band remote management, including keyboard, video and mouse redirection, storage redirection (boot a remote ISO), serial-over-LAN and power control. AMT exposes, when enabled, TCP 16992 (HTTP), 16993 (HTTPS), 16994 (Redirection), 16995 (Redirection-TLS), UDP 623 (ASF-RMCP) and KVM on TCP 5900 (Intel AMT, Wikipedia; Intel AMT KVM developer guide).

Platform Trust Technology is Intel's firmware TPM 2.0, also implemented as a CSME applet. Where BitLocker is configured to use "TPM" on a vPro platform without a discrete TPM, the BitLocker volume master key is sealed and unsealed by code running inside the same CSME image as AMT.

The architectural fact on which the rest of the assessment turns is that AMT and PTT live in the same firmware, on the same microcontroller, in the same address space, separated by

software-enforced privilege boundaries inside that firmware. They are not separated by a silicon boundary.

4. Adversary tiers in scope

The question "what can a third party do?" depends on which third party. The assessment uses four tiers.

Tier 1, opportunistic. An attacker on a hostile network such as a coffee-shop or hotel guest LAN. Modest budget. Cannot mis-issue CA certificates, cannot interdict shipments. Wants disruption, credential capture or commodity data theft.

Tier 2, targeted criminal or commercial spyware. Significant budget. Can buy CA-issued certificates with custom OIDs through compromised resellers, can rent or operate hostile networks at venues the target visits, can deploy hostile wireless at distance. Cannot reliably interdict shipments at scale unless an operative is placed at a logistics hub. Wants long-term access to a specific target or to a class of targets.

Tier 3, nation-state cyber service. Operates at the level of compromising firmware in transit, mis-issuing CA certificates via an in-country CA and modifying hardware at customs or courier hubs. The NSA's Tailored Access Operations division was documented through the 2013 Snowden disclosures as treating supply-chain interdiction as routine business, with explicit doctrinal preference for BIOS and firmware implants over disk-level malware (NSA ANT Catalog, EFF document mirror PDF; The Intercept, "Everybody Does It"). Several peer agencies are widely understood to operate similar capability. PLATINUM is a documented Tier-3-class actor that has used AMT against targets.

Tier 4, trusted insider in the supply chain. A reseller employee, a courier, a customs broker, a logistics handler or a factory-floor worker. Limited cryptographic capability and extensive physical access for windows of minutes. The existence of this tier as a realistic threat model is documented by Microsoft's hardware supply-chain risk treatment (Microsoft, "Guarding Against Supply Chain Attacks Part 2").

A control that defeats Tier 1 but is bypassed at Tier 2 still leaves the organisation exposed at Tier 2. The recommendation is calibrated by which tiers are in scope. A note on the BYOD case specifically: the relevant Tier-4 is no longer constrained to a corporate-procurement courier or reseller, because the device may have passed through any retail channel including grey-market resale, refurbishment and second-hand sale.

5. Attack surface

A vPro laptop with an untouched ME is reachable through more channels than the simple "wired Ethernet only, when plugged into AC" framing.

5.1 Wired Ethernet on a hostile network. When a user plugs the device into a wired drop, the CSME initialises its network stack independently of the host OS. Where remote configuration is supported by firmware default, the CSME emits unsolicited "Hello" frames advertising its pre-provisioning state (Intel community discussion of AMT Hello messages and remote config). Anyone capturing the broadcast or crafting a matching provisioning response can begin a provisioning handshake. Whether the handshake succeeds depends on the cryptographic gate in 5.5.

5.2 Wired Ethernet on a hostile home network. The mechanism is identical to 5.1. The likelihood is higher because most home networks are not 802.1X protected and not under organisational monitoring. For BYOD specifically, the home network is the device's primary network most days.

5.3 Wi-Fi. The defence "an unprovisioned ME has no Wi-Fi profiles, therefore the radio is unreachable" is a configuration property of the device, not a hardware property. Three things break it. A supply-chain adversary with brief physical access can boot the device into MEBx, inject a hostile SSID and pre-shared key and re-seal the chassis; the laptop will then associate to the attacker's network whenever the radio is in range and will appear to the user to be sleeping. A future or undisclosed firmware bug in the radio path can in principle bring the radio up without configuration. Vendor-default behaviour can switch the radio on for diagnostic, recovery or zero-touch purposes.

5.4 Inline cohabitation, shared MAC and IP. From AMT 3.0 onwards, the CSME shares a single MAC address with the host operating system (Intel Community thread, "Mapping between Host IP and AMT IP"). When configured for IPv4 with shared addressing, the CSME monitors the host's DHCP exchanges and synchronises onto the same IP address the host receives, or in later releases requests address updates via the Local Manageability Service (Intel manageability documentation, "Configuring the Intel AMT IP Address"; Automatic Synchronization of IP Addresses). The implications for monitoring sit in Section 7.

5.5 The cryptographic provisioning gate. Remote (zero-touch) provisioning is gated by the device's firmware verifying the provisioning server's certificate against three conditions: the issuer's root hash is on the trust list burned into the CSME firmware (Acquiring an Intel vPro Certificate); the certificate carries the AMT provisioning OID 2.16.840.1.113741.1.2.3 in its Extended Key Usage; the certificate Common Name matches the DNS suffix supplied by the local network's DHCP Option 15. The trust list contains hashes for major commercial CAs. The list is editable from MEBx. Whether the gate holds depends on whether any one of those CAs can be induced to mis-issue, on whether Option 15 is supplied by an attacker-controlled network, and on whether the trust list has been altered in transit.

5.6 The firmware itself, independent of provisioning. Even where the gate is not breached, the CSME network parser is listening. A memory-corruption flaw in the parser handling unauthenticated UDP 623, in the AMT HTTP front-end on 16992 or in the radio firmware path becomes an unauthenticated remote code execution flaw below the operating system. Section 9 sets out what the historical record shows about how often that has happened.

5.7 Brief physical access in transit. The MEBx default credential is `admin`, with the user expected to change it on first entry (Intel AMT MEBx documentation); a factory-state device has not yet been entered. A short window with a powered device is sufficient to enter MEBx, change the password, inject Wi-Fi profiles, inject a root certificate hash, enable remote configuration and re-seal the box. F-Secure publicly demonstrated in 2018 that a near-equivalent end-to-end attack on an unprovisioned device takes approximately thirty seconds

of physical access (BankInfoSecurity, "Backdoored in 30 Seconds: Attack Exploits Intel AMT Feature"), which is a useful concrete time estimate for the supply-chain-interdiction operating doctrine.

5.8 Host-to-ME pivot. A compromise of the host operating system at Ring 0 is conventionally treated as the worst-case outcome the host can reach. On a vPro platform it is not the worst case, because the host has a reachable path to the CSME. The Intel-recommended Flash Descriptor write protections, when enabled by the platform manufacturer, restrict who can write to the SPI flash region holding the CSME firmware. Where those protections are not enabled, an attacker who has obtained kernel-level execution on the host can write a modified firmware image through the platform driver and persist into the CSME runtime (Eclypsium, "Firmware Security Realizations Part 2"). The HECI / MEI driver provides the in-band communication channel between host and CSME, and historical advisories have identified flaws in firmware update validation that allow malformed or attacker-controlled images to be accepted. The implication for incident response is that an OS-level compromise on a vPro device cannot, on standard host evidence alone, exclude the possibility that the compromise persisted into firmware. A wipe and reimage of the disk does not clear the CSME. Section 10.6 walks through the resulting threat scenario.

5.9 Other radios accessible from the integrated PCH connectivity stack. Wi-Fi is the most-discussed radio in the public ME literature, but it is not the only one. From the Intel CNVi (Connectivity Integration) generation onwards, the network adapter MAC components, memory, processor and associated firmware are moved inside the chipset Platform Controller Hub, and only the analog and RF functions are left on an external upgradeable Companion RF module in M.2 form factor (CNVi, Wikipedia; Intel, "What Are the Intel Integrated Connectivity (CNVi) and Companion RF (CRF) Module"). The same on-PCH connectivity controller serves both Wi-Fi and Bluetooth. The architectural consequence is that any compromise of the CSME runtime that reaches the connectivity firmware path reaches both protocols on the same module, and a Bluetooth-attached attacker (proximity range of around ten metres for Class 2, longer for Class 1 with directional gear) may be reachable on the same surface that an in-the-clear Wi-Fi attacker would use. The radio MAC, processor and firmware are inside the PCH alongside the ME; the threat surface is not separated by silicon.

WWAN modules on M.2 typically attach via PCIe rather than CNVi (the Fibocom FM350-GL Intel 5G 5000 module, for example, exposes a PCIe Gen 3 host interface), so the WWAN module is not on the CNVi shared-firmware path with Wi-Fi. The WWAN module nonetheless sits on the PCH PCIe fabric the CSME has access to. A device with an active WWAN subscription represents an additional out-of-band reach path that is independent of the user's choice of Wi-Fi network and that may be reachable by the operator of the cellular network or by an attacker with a stingray-style intercept capability. NFC, where present, is generally a host-OS-controlled subsystem and is less directly exposed, but should be confirmed per platform.

The summary for the threat model is that "the radio" is shorthand for several radios that share architectural neighbourhoods with the CSME. Hostile-Wi-Fi reach (5.3) extends to hostile Bluetooth reach on CNVi-equipped platforms, and the cellular path is its own out-of-band reach path on platforms with active WWAN modules.

5.10 Wi-Fi credentials persisted in the CSME after legitimate provisioning. Once a vPro device has been provisioned for AMT-over-Wi-Fi, the CSME stores the corporate 802.1X credentials. A subsequent compromise of the CSME on that device, by any of the paths above, exposes those credentials. The risk is not specific to the BYOD or IT-supplied case in scope here (the device must have been legitimately provisioned for the case to arise) but it should be named for organisations that, on completing the procedural fix recommended in Section 17.1, then operate AMT-over-Wi-Fi at scale.

6. Management infrastructure as a separate (referenced, not detailed) risk

For organisations that operate an enterprise AMT management plane (Intel EMA, MeshCentral, an OEM management service or equivalent), that infrastructure is its own attack surface. A compromise of the management server yields, at scale, the AMT administrator credentials it holds for every device under management, with the realised-attack consequences described for individual devices in Section 10.5. The two cases that this document is scoped to (BYOD and IT-supplied without ME interdiction) typically do not have an enterprise AMT management plane in operation, because by definition the ME has not been provisioned by the organisation. The risk is referenced here for completeness, on the basis that organisations addressing the IT-supplied case via the procedural recommendation in Section 17.1 will, in due course, stand up such a management plane and inherit this surface.

7. Inline cohabitation in detail

A reasonable question is whether the corporate firewall can see ME traffic that the host firewall cannot. The answer is that the corporate firewall is in a similar but not identical position to the host firewall, and that L3 filtering cannot disambiguate ME-originated from host-originated traffic on a per-flow basis.

On the wire there is one MAC address, the host's, and one IP address, the host's, in shared-IP DHCP mode (the default in most enterprise AMT deployments). Inside the NIC silicon a manageability filter examines every inbound frame; frames matching the AMT port set (16992, 16993, 16994, 16995, 623, 5900) are diverted to the CSME via the chipset sideband, and other frames pass to the host TCP/IP stack (Matthew Garrett, "Intel's remote AMT vulnerability").

From the host side, the host TCP/IP stack sees only the frames that were not intercepted. The host firewall, EDR and network monitoring tools operate on those frames. The intercepted AMT-port traffic is delivered to a different processor on the same physical NIC, and host-resident defences cannot see, log or block it.

From the corporate firewall side, the situation is the inline cohabitation property restated. Outbound traffic generated by the CSME egresses through the same NIC as host traffic, with the same source MAC and source IP. To a device-by-device firewall, IDS or NDR sensor at the perimeter, the traffic is "from this laptop". An L3 access control list applied to the laptop's IP address applies equally to host traffic and to ME traffic, because they share an IP address. A signature on the AMT port set at the perimeter can identify AMT-shaped traffic, but the host could in principle generate traffic on those ports too, and a ME-resident exfiltration channel that uses a tunnelled protocol over an arbitrary port (a TLS tunnel back to an attacker server using port 443) does not fingerprint as AMT traffic at all.

The corollary for monitoring is that the perimeter retains some advantages over the host. The perimeter sees flows the host does not, including AMT-port flows and SOL flows. The perimeter can apply destination reputation and TLS metadata controls that the host cannot. But the perimeter cannot, in the general case, attribute a flow to host-versus-CSME origin without out-of-band knowledge. The out-of-band knowledge available is firmware integrity attestation taken at boot or on reconnect, which is the only signal that survives the inline cohabitation property and is discussed in Section 15.

The conclusion to brief the CISO on is that corporate detection of CSME activity requires either network-side signatures targeting AMT-shaped traffic (accepting that a sophisticated attacker will not use AMT ports for exfiltration), network-side anomaly detection on host flows (accepting that the device is the only source of attribution), or deferred attestation. Defence-in-depth arguments that lean on host firewall or perimeter L3 ACLs as the load-bearing control are mis-specified at the architectural level.

8. Why a powered-off laptop drains its battery over weeks

A common observation in any laptop fleet is that a device that has been powered off and stored for several weeks is found, on next use, to have a depleted battery. The observation is consistent with public reporting on Modern Standby (S0ix) behaviour and with Intel's published power state design. It is included here because a security officer evaluating the vPro attack surface will reasonably want to know whether a "powered off" laptop is in fact powered off; the public record indicates that on modern thin-and-light platforms it is not.

Microsoft documents Modern Standby as a state in which the system remains in S0 (working) topology, with hardware and operating system components transitioning to lower-power phases while permitting limited background activity and network connectivity. A Modern Standby idle drain on the order of several percent of battery capacity over twenty-four hours is the design target, which a 55Wh battery yields as a continuous draw on the order of one to two hundred milliwatts (Microsoft, "Modern Standby" learn pages-be-causing-my-laptop-t)). User and press reporting documents that real-world drain frequently exceeds the design target, often attributed to misbehaving wireless drivers and to background SoC activity (XDA Developers, "Modern Standby is draining your Windows 11 laptop battery"; Windows Latest, February 2026 reporting).

The implication for the threat model is twofold. First, the user-perceived state "off" does not correspond to the architectural state "all subsystems unpowered", and the active subsystems include the SoC components on which the CSME runs. Second, the persistence of low-power listening states on these platforms makes claims of the form "the CSME radio is off because it cannot run on battery" empirically unsupported as universal statements; whether the radio is in a Wake-on-Wireless-LAN listening state is a firmware policy question, and on a device whose firmware policy has been altered in transit, the answer cannot be inferred from the visible power state.

The point is not that battery drain proves CSME activity. It is that a small but non-zero baseline draw on these platforms, of the magnitude documented in Microsoft and press reporting, is consistent with low-power listening and is compatible with attacker-injected configurations that the user has no visibility into.

9. The evidence base

The assessment is grounded in named, citable CVE entries, security advisories and published research. A patched status against any individual entry is treated as confirmation that the surface produces this class of bug, not as closure of the residual risk.

9.1 The "admin no password" surface, CVE-2017-5689

The most useful single entry in the public record for framing the resilience of the AMT management plane is CVE-2017-5689, sometimes called "Silent Bob is Silent" after the Embedi disclosure, and popularly described as the "admin no password" bug (NVD; Embedi PoC repository; The Register, "Intel patches remote execution hole hidden since 2010"). The mechanism is more interesting than the popular framing. AMT's HTTP digest authentication compared the user-supplied response field against the computed response using a length-bounded comparison; supplying an empty response field caused the comparison to succeed against zero bytes, which the code accepted as a valid response. The result was that a network attacker reaching TCP 16992 on a provisioned device gained full administrative access to AMT without supplying any password, by sending an empty Authorization header.

Three properties of this bug are worth carrying forward. Firstly, the affected firmware spans Intel ME 6 through 11.6, which covered the ARC / ThreadX architecture and the early Quark / MINIX architecture; the architectural change did not eliminate it. Secondly, the bug shipped in production silicon for approximately seven years before public discovery, on a network-facing service running below the operating system, on every vPro-capable Intel laptop and workstation in that period. Thirdly, the AMT management plane is the same surface that, when not subject to a vulnerability of this kind, provides the legitimate provisioning, KVM, IDE-Redirect and Serial-over-LAN capabilities that make AMT a useful out-of-band management feature; the same surface that delivered "admin no password" also delivers the in-the-wild abuse documented in Section 9.4.

The lesson to brief is not "the bug is patched". It is that this is the ambient quality of code that the architecture exposes to the network, and there is no architectural reason the next bug of this kind has been disclosed yet rather than waiting in the disclosure pipeline.

9.2 Pre-Skylake research on the ThreadX-era ME

The foundational public demonstration that code can run on the Intel ME at a privilege level below SMM and the hypervisor, with full DMA access to system memory and outside the operating system's ability to detect or contain it, is the Tereshkin and Wojtczuk Black Hat USA 2009 paper (PDF). It is not a patched bug; it is the architectural property the rest of the assessment depends on.

The detailed reverse engineering of the ARC / ThreadX-era CSME, including the boot ROM, the FTPR partition layout, the BUP and KERNEL modules and the code-signing flow, is set out in the Skochinsky 2014 work referenced in Section 3. The work established the surface that subsequent research has been mining.

9.3 MINIX-era CSME memory corruption and key extraction

INTEL-SA-00086 (CVE-2017-5705, CVE-2017-5706 and CVE-2017-5707) is a buffer overflow in the CSME's internal MFS file system handler, demonstrated to enable JTAG access to the CSME over the USB DCI debugging interface (a path that requires physical USB DCI access) and arbitrary code execution inside the CSME runtime (Intel; Black Hat Europe 2017 paper PDF). It is the public demonstration of Ring -3 unsigned code execution on a "turned-off" computer.

CVE-2019-0090 is a vulnerability in the CSME boot ROM, in mask-programmed silicon rather than in field-updatable firmware, allowing local extraction of the CSME chipset key

from which a downstream attacker can derive every key the platform protects, including PTT-stored keys (Positive Technologies, "Intel x86 Root of Trust, Loss of Trust"; SecPod; TechTarget). It cannot be patched in the field. It affects all Intel chipsets and SoCs prior to the Ice Point generation. The relevance for this assessment is direct, because CVE-2019-0090 invalidates the assumption that PTT (the firmware TPM) is isolated from AMT inside the CSME so that a CSME compromise cannot extract BitLocker keys; on affected hardware, it can, and there is no patch coming.

CVE-2021-0146 is a debug or test logic activation flaw on Apollo Lake and Gemini Lake-based Pentium, Celeron and Atom processors, also disclosed by Positive Technologies, allowing extraction of the CSME firmware key that secures PTT and Intel EPID (Help Net Security; SecurityWeek). The pattern is that the CSME-as-trust-root model has had to be re-patched because the trust root proved extractable, more than once, across multiple silicon generations.

INTEL-SA-00213 is a buffer overflow in CSME 12.0.0 to 12.0.34 reachable by an unauthenticated user via network access, leading to escalation of privilege, CVSS 9 (Intel). INTEL-SA-00295 covers further buffer-restriction flaws across 12, 13 and 14 (Intel). CVE-2025-20037 is a time-of-check time-of-use race in CSME firmware before 14.1.77.2497, requiring local privileged access (ZeroPath summary). INTEL-SA-01138, INTEL-SA-01200 and INTEL-SA-01280 are recent advisories in the same line, with full details at the Intel Security Center index (Intel Security Center).

9.4 PLATINUM, the realised in-the-wild abuse

Microsoft documented the PLATINUM nation-state-aligned actor in June 2017 using AMT Serial-over-LAN as a covert exfiltration channel (Microsoft Security Blog; Threatpost; BleepingComputer; The Hacker News). The mechanism does not exploit a vulnerability; it exploits a feature, requiring only that AMT be enabled and that valid credentials have been obtained. SOL traffic transits the CSME and the NIC sideband and is not visible to the host TCP/IP stack, so host firewall, host EDR and host network monitoring see nothing. PLATINUM is the single reference point that separates the AMT threat model from the abstract; the actor was real, the use of AMT was real, and the host-side visibility was, as designed, zero.

9.5 Hardcoded trust anchors as a recurring industry pattern, the Cisco parallel

A reasonable reviewer asks whether the CSME's hardcoded root certificate hash list is the kind of trust anchor that, in the public record, has been compromised before. The Cisco product lineage provides the closest enterprise parallel, because Cisco has in repeated cases shipped network and unified communications products with hardcoded credentials, default static SSH keys or signing material that granted unauthenticated remote root access to anyone holding the corresponding key.

CVE-2025-20309 is the most recent maximum-severity example, a static SSH credential for the root account in Cisco Unified Communications Manager 15.0.1 Engineering Special releases, documented by the vendor and reported in industry coverage (Cisco Security Advisory; Arctic Wolf summary; BleepingComputer; CSO Online). CVE-2021-34795 covered a default static debug password in the Cisco Catalyst PON Series switches at CVSS 10. CVE-2021-40119 covered static SSH keys in Cisco Policy Suite installations granting unauthenticated remote root (The Hacker News). The Cisco Nexus 9000 software contained a default SSH key pair allowing remote root connection. The pattern is sufficiently established that Bruce Schneier's "Cisco Can't Stop Using Hard-Coded Passwords" (Schneier on Security) is recognisable industry shorthand.

Two implications follow for this assessment. Firstly, the assertion that a major enterprise vendor has burned cryptographic trust anchors into shipped silicon and firmware in a way that grants a holder of the matching key elevated remote access is not hypothetical, it has happened repeatedly at a major network-equipment vendor, and the pattern is not specific to Cisco. Secondly, the residual risk that the CSME's burned-in factory root hash list contains a hash for which the corresponding signing material is held outside the legitimate certificate authority chain (whether via national security order, insider compromise or undocumented engineering arrangement) is not a credible position to dismiss as paranoid. The residual is named explicitly in Section 16.2.

9.6 Supply-chain interdiction and rogue insider, the realised cases

The Tailored Access Operations doctrine documented in the NSA ANT material describes shipment interdiction, BIOS implantation and firmware tampering as routine operations against hard targets (Wikipedia, ANT catalog; EFF document mirror PDF; The Intercept). The ANT preference for firmware-resident implants over disk malware is explicit, because firmware implants survive OS reinstallation and disk replacement.

Microsoft's hardware supply-chain risk treatment names the trusted-employee-or-vendor-with-physical-access threat model directly (Microsoft Security Blog). Firmware-security industry literature documents the broader pattern (Eclypsium, "Firmware Attacks: An Endpoint Timeline"; Dark Reading, "Firmware Vulnerabilities Continue to Plague Supply Chain"). Operation ShadowHammer in 2019 saw the ASUS Live Update Utility compromised at the vendor's signing infrastructure, pushing malicious firmware updates signed with valid ASUS certificates to thousands of devices (IEEE Spectrum). The Bloomberg "Long Hack" reporting (Bloomberg) is included with the contested-status flag; its relevance is not as evidence of a confirmed incident, but because the capability described, compromise at the contract-manufacturer level, is not in dispute.

The relevance to the assessment is that any laptop reaching a user without organisational interdiction has passed through several parties, including contract manufacturer, finished-goods warehouse, freight forwarder, customs and last-mile courier (in the IT-supplied case), or any retail or grey-market path (in the BYOD case). Each introduces an opportunity for a sixty-second physical-access window. The proper question is not "has this happened to a shipment of ours?", it is "what evidence would we have if it had?", and on a factory-state vPro laptop with no firmware-integrity attestation at receipt, the answer is none.

10. End-to-end threat scenarios

Each scenario is constructed so that the assumption being relied upon at each step is identifiable, and the public evidence for that step is cited above.

10.1 Scenario A, opportunistic on a hostile RJ45 (Tier 1)

The user plugs an unprovisioned vPro laptop into a wired drop at a customer site or hotel. The CSME initialises, requests DHCP and (where zero-touch is enabled by firmware default) emits an unsolicited "Hello" frame.

A Tier-1 attacker on the same LAN segment captures the broadcast. They cannot present a valid AMT-OID certificate, so the cryptographic gate holds. They have, however, confirmed identification of the device, its MAC, its CSME firmware version and its pre-provisioning state, and they can probe the device with any unauthenticated CSME network-stack exploit in their toolkit. On a fully patched device with no current zero-day, the probe lands as a denial of service or fails outright; on a device whose firmware is one cycle behind, it is a foothold. The firmware-currency question is examined separately in Section 11.

10.2 Scenario B, supply-chain Wi-Fi-profile (or Bluetooth) injection followed by drive-by capture (Tier 3 or Tier 4)

A laptop is interdicted in transit, at a logistics hub, customs hold, courier facility, contract-manufacturer floor or, in the BYOD case, at any retail or grey-market handler. A short window with the device powered is sufficient to enter MEBx with the default `admin` credential, inject a hostile SSID and pre-shared key, add a hostile root-CA hash to the trust list, lock MEBx with an attacker-chosen password and re-seal the box.

The user receives the laptop and carries it. At any point within radio range of an attacker-controlled SSID (a parked vehicle, a malicious access point, a long-range directional antenna), the CSME radio firmware associates and brings up an out-of-band link, with no user interaction, no host OS involvement and no visible indication that the laptop is doing anything other than sleeping in a bag. On a CNVi-equipped platform the same compromise reaches the integrated Bluetooth radio, and a Bluetooth-attached attacker in proximity range becomes a viable reach path independent of any Wi-Fi network the device is in range of. The attacker then has Layer-3 reachability to the CSME on a device whose host OS has never touched the corporate network.

Outcome at Tier 3 or Tier 4, complete out-of-band control without ever having access to the corporate physical network. The "must be plugged in to be attacked" framing does not survive at this tier. There is no software fix available on a device that has been outside the trusted chain of custody, because the corruption is below the OS boundary.

10.3 Scenario C, fraudulent provisioning certificate on a hostile RJ45 (Tier 2 or Tier 3)

The attacker possesses a certificate issued by a CA whose root hash is on the firmware trust list, with the AMT OID populated in the EKU, by mis-issuance, CA compromise, reseller-pipeline compromise or coercion under a national security order in the CA's jurisdiction. The attacker controls the network the user plugs into, including DHCP Option 15.

The user plugs in. The CSME requests DHCP. The attacker's DHCP supplies an Option 15 matching the certificate's CN. The CSME accepts the provisioning payload and transitions to Post-Provisioning under the attacker's control. The attacker now has Admin Control Mode, including KVM redirection, IDE-R, serial-over-LAN and full power control.

The attacker can capture every keystroke and every video frame as the user types passwords, BitLocker pre-boot PINs, recovery keys and smart-card PINs; FIDO2 hardware

tokens are immune to this for FIDO2-protected sign-ins, but any password-based fallback the user is permitted to use is exposed. The attacker can mount a remote ISO over IDE-R, persist across host-OS reinstallation and disk replacement, and exfiltrate captured data over Serial-over-LAN, invisibly to the host firewall and EDR, in the manner PLATINUM did in 2017.

The Tier-3 attacker has the option of chaining a CSME firmware exploit on top of the provisioning foothold, which collapses the AMT / PTT applet boundary from inside (Scenario D below).

10.4 Scenario D, pure CSME firmware exploit (Tier 2 or Tier 3)

The attacker has a working exploit for a memory-corruption flaw in the CSME network stack of the class historically represented by INTEL-SA-00086, INTEL-SA-00213, INTEL-SA-00295, CVE-2025-20037 and the ongoing 2024 to 2026 advisory cadence.

The attacker reaches the device via Scenario A, B or C and sends a single crafted UDP 623, TCP 16992 or radio-firmware frame. The frame triggers the bug and they achieve code execution inside the CSME runtime.

From there they are no longer constrained by the AMT applet API; they are running inside the CSME OS, alongside the PTT applet. The applet boundary is enforced in software, by the same firmware they have just compromised. CVE-2019-0090 demonstrated, in public, that the trust foundation underneath this boundary can fail in ways the platform vendor cannot patch in the field, on pre-Ice-Point silicon. On platforms that use PTT (firmware TPM) the worst-case outcome is BitLocker volume master key extraction directly from PTT memory; on platforms with discrete TPM, the same applet-boundary breach does not extract the BitLocker key directly, but provides Ring -3 host-memory access via DMA that can read the host's in-memory representation of decrypted data. Persistence is in CSME flash; reformatting the SSD does not remove it, replacing the SSD does not remove it, and a standard "wipe and reimage" does not remove it.

Outcome at Tier 2 or Tier 3, host compromise transparent to the OS, persistent across reimaging. On PTT platforms with affected silicon, full-disk encryption keys are extractable directly from the security co-processor. On other platforms, in-memory data is reachable via DMA.

10.5 Scenario E, AMT credential abuse (the PLATINUM pattern, Tier 2 or Tier 3)

A device that has been provisioned (whether by a legitimate IT department or by a hostile party in 10.2 or 10.3) sits in the user's normal network environment. The attacker has obtained the AMT administrator credential by phishing, by compromise of the management server, by compromise of the help-desk staff who hold them, by insider abuse or by capturing them during the unprovisioned window in 10.3.

The attacker connects to the device's AMT interface and opens a Serial-over-LAN session. SOL is the covert exfiltration channel for data harvested by other means; the host TCP/IP stack does not see SOL traffic; the host firewall and host EDR do not see SOL traffic; perimeter monitoring sees outbound flows from the laptop's MAC and IP but cannot, on a per-flow basis, distinguish them from legitimate host traffic, because they share that MAC and IP (Section 7).

The attack does not require any vulnerability. It requires that AMT be enabled and that credentials have been obtained. It is realised in-the-wild precedent.

For the two cases this assessment is scoped to (BYOD and IT-supplied without ME interdiction), Scenario E is most acutely a risk after a hostile provisioning event under 10.2 or

10.3; in steady state on an unprovisioned device, Scenario E does not arise, but the device is one provisioning event away from it.

10.6 Scenario F, host compromise pivots into persistent CSME implant (Tier 2 or Tier 3)

The attacker reaches Ring 0 on the host through any conventional path, including a phishing-delivered malware payload that achieves kernel privilege through a Windows local privilege escalation, or a browser-or-application remote code execution chained into a kernel exploit. On a vPro device whose Flash Descriptor write protections are not enabled at the platform level, the attacker uses the HECI / MEI in-band channel, or directly writes to the SPI flash region holding the CSME firmware via the platform driver, to install a modified CSME firmware image (Eclypsium, "Firmware Security Realizations Part 2"; the same Flash Descriptor protection topic is referenced in the INTEL-SA-00086 disclosure).

The host-resident component of the attack can then be removed by the IR team during conventional eradication, and the visible host-side IOCs disappear. The CSME-resident component remains. On next boot the device runs the modified CSME image, which provides an out-of-band channel that is invisible to host EDR, persists across full-disk wipe and reimage, and persists across SSD replacement, because the implant is not on the SSD.

The implication for incident response is that an OS-level compromise on a vPro device does not, on standard host evidence alone, allow the IR team to exclude firmware persistence. The eradication standard described in Section 16.4 is therefore stricter than the standard used for non-vPro fleets. The implication for procurement, on the IT-supplied case, is that the Flash Descriptor write protection setting should be verified at receipt as part of the lab interdiction step. In the BYOD case, the organisation has no procurement leverage and no opportunity to verify the setting; the residual is named in Section 13.

11. The firmware update push problem

A common assumption in security planning is that firmware vulnerabilities are addressed by keeping firmware current. For CSME firmware the assumption is partially supported on current-generation IT-supplied hardware with a cooperating OEM, and is not supported on most BYOD or older IT-supplied hardware. The constraints are matters of public record and warrant a section because they shape the recommendation in Section 17.

CSME firmware is not, in the general case, updateable through standard enterprise software-distribution mechanisms. The firmware ships as part of the OEM BIOS or UEFI capsule, and the organisation's ability to apply a CSME update is gated by whether the OEM has chosen to publish the relevant CSME version inside a signed BIOS capsule for the affected SKU, on the affected platform generation. The Linux Vendor Firmware Service maintainers have discussed publicly that "most vendors package the CSME update as a signed capsule that can be distributed as an update on LVFS" but that "there may be legal or contractual reasons that prevent some OEMs from obtaining ME updates from the ODM as a capsule" (fwupd discussion 9690 on LVFS CSME firmware delivery). The same source notes that capsule-based delivery of CSME has only been routinely possible from the thirteenth-generation Intel platforms onwards, and that for eleventh and twelfth-generation platforms "bundling ME firmware into the main BIOS update was not supported by the BIOS vendor". The implication is that on a non-trivial population of in-fleet hardware, the CSME version is governed by the BIOS the OEM chose to ship at platform launch, and is not changed by routine OS or driver updates.

Where capsule delivery is possible, the practical update path runs through OEM-specific tooling rather than through the central Microsoft Update channel. Dell Command Update, HP Image Assistant, Lenovo Vantage and equivalent vendor utilities are the typical mechanisms, with Microsoft Update providing some coverage where the OEM has registered the device class and the capsule. None of these is a centralised, vendor-neutral, enterprise-controlled push channel. Intel itself, in the INTEL-SA-00086 advisory page, directs end users to "Contact your system or motherboard manufacturer regarding their plans for making the updates available to end users" (Intel-SA-00086 support article), which is the public statement that Intel cannot push the update directly.

The Intel-supplied direct-update tool, FWUpdLcl, runs locally on the device and requires administrative privilege. It will not downgrade the firmware (the public error response is "Error 8758: The image provided is not supported by the platform" on attempted downgrade), and it will not accept images that are not signed by the appropriate combination of Intel and OEM keys (Intel community discussion of FWUpdLcl downgrade behaviour; Intel community on CSME firmware update tools). Where the OEM has cross-signed the firmware partition (typical for CSTXE 3.0 and later), the partition is not updateable outside the OEM's release cadence (Win-Raid forum threads on cross-signing constraints).

Two consequences follow for the present assessment.

In the IT-supplied case, where the IT team is or becomes aware of the ME, the firmware-currency mitigation is not a routine endpoint-management deliverable but a procurement and lifecycle decision. The lifecycle decision includes selecting OEMs that publish capsules promptly, requiring the capsule channel as a procurement criterion, scheduling devices through the lab on a cadence aligned to OEM capsule releases (rather than aligned to Intel's quarterly Platform Update calendar), and accepting that some SKUs will reach end-of-vendor-support before they reach end-of-fleet-life and that those devices will not receive further CSME updates regardless of advisories.

In the BYOD case, the organisation has no firmware-push authority of any kind. The user is responsible for accepting and applying BIOS updates from the OEM, and the organisation has no visibility into whether the user has done so. CSME firmware version, on a BYOD vPro device, is whatever the user has chosen (or, more accurately, has not chosen) to

install, on whatever cadence the OEM happens to publish, against whichever advisories the OEM has elected to address. Treating "keep firmware current" as a control on BYOD vPro is not honest.

Industry telemetry supports the practical reality. Eclypsium reported, on the basis of analysis of production fleets, that the proportion of devices observed vulnerable to specific Intel ME advisories remains a majority of the fleet years after public disclosure, with figures of approximately 72% vulnerable to INTEL-SA-00391 and 61% vulnerable to INTEL-SA-00295 in observed environments; the same reporting documents that the Conti ransomware group developed proof-of-concept Intel ME exploit code with the intent of installing highly persistent firmware-resident implants (CSO Online, "Cybercriminals look to exploit Intel ME vulnerabilities for highly persistent implants"). The data are useful in two ways: they substantiate that "keep firmware current" is, in production, an aspirational control rather than an operational one across the majority of enterprise vPro fleets, and they identify a named criminal actor at Tier 2 with a working interest in the same surface that nation-state actors have already abused.

12. Why each commonly-cited "protection" is conditional, not absolute

The vendor's consolidated security position is set out in Intel's vPro Security Overview white paper, December 2024. Where the section below contests a property the vendor white paper asserts, that assertion is identifiable. The intent is not to dispute the vendor white paper in detail but to identify, item by item, the conditions on which each commonly-cited protection holds. Industry literature (and the vendor's own marketing) makes claims of the form "vPro is fine because the gates are cryptographic". Each of the following protections is real, and each is conditional on assumptions that the public record does not support as universal.

The unprovisioned ME has no Wi-Fi profiles, so it cannot be reached over Wi-Fi. Configuration property of the device, defeated by 5.3, 5.7 and Scenario B. On CNVi-equipped platforms the same compromise extends to Bluetooth (5.9).

The AMT provisioning gate requires a certificate with the AMT OID issued by a CA whose hash is in the firmware. The trust list contains commercial CAs, and CA mis-issuance is a documented recurring event; the gate raises the bar to Tier 2 and above, it does not eliminate it. Section 9.5 records that the wider industry pattern of hardcoded trust anchors granting elevated remote access has been realised at a major network-equipment vendor across multiple advisories.

Once the legitimate IT department provisions the device, the attack surface closes. True only if the legitimate IT department reaches the device first. In the IT-supplied / ME-naive case being addressed here, the legitimate IT department does not provision the device at all, so the gate remains open. If a hostile party has provisioned the device first (Scenarios B or C), the gate is now the attacker's gate, locked behind credentials the legitimate IT team does not hold.

AMT and PTT are isolated by hardware-enforced memory management inside the CSME. Software-enforced privilege boundary inside a firmware runtime; CVE-2019-0090 publicly demonstrated that the trust foundation underneath was extractable on all pre-Gen-10 silicon, INTEL-SA-00086 publicly demonstrated arbitrary code execution inside the CSME OS, and the advisory pattern shows the boundary has been re-broken across multiple firmware generations.

The HAP "Reserved" bit disables AMT, so out-of-band attack surface closes. It disables the AMT applet and a subset of CSME runtime services after the BUP phase, but it does not de-power the CSME silicon. The CSME is required to bring up the chipset.

Modern firmware has fewer bugs than the 2017 firmware. The 2024 to 2025 CSME advisory record (CVE-2025-20037, INTEL-SA-01138, INTEL-SA-01200, INTEL-SA-01280) shows the surface continues to produce bugs at the rate of roughly several per Intel quarterly Platform Update cycle. On BYOD or older IT-supplied hardware, the practical question is not how many bugs the latest firmware contains but how many bugs the firmware actually loaded on the device contains, and that is governed by the constraints in Section 11.

BitLocker plus TPM means data at rest is safe even if the ME is compromised. Relies on the TPM being a separate trust domain from the ME; on platforms using PTT they share a runtime, and CVE-2019-0090 provides direct evidence the trust separation can be defeated. On platforms with discrete TPM, a CSME compromise still has DMA access to host RAM (Scenario D), so the in-memory representation of decrypted data is reachable even where the BitLocker key is not directly extracted.

FIDO2 protects sign-in. FIDO2 is immune to KVM credential capture only on the specific authentications that go through the FIDO2 flow. Any password-based fallback in the user's day, including BitLocker pre-boot PINs, recovery keys, smart-card PINs and an RDP target's Windows password, is captured.

The corporate VPN protects data in motion. The attacker is below the corporate VPN; they are reading the framebuffer the VPN client renders into and the keystrokes the VPN client receives.

The host firewall and host EDR will catch unusual outbound traffic. They will not, by the inline cohabitation property in Section 7.

The corporate firewall and L3 perimeter ACLs will catch unusual outbound traffic. They retain some advantages over the host but cannot, in the general case, attribute a flow to host-versus-CSME origin, because they share an IP address. Section 7 sets out the implications.

13. Residual risk before mitigation

The residual profile differs by procurement provenance and by adversary tier. Likelihoods are stated qualitatively, because precise probability estimation against an adversarial population is not honest.

For IT-supplied hardware shipped factory-state because the IT function is unaware of the ME, the residual on the existing fleet is approximately equal to the residual under a no-mitigation deployment, because no mitigation has been applied. For each scenario in Section 10, the residual is the impact of that scenario weighted by the relevant tier's likelihood. Tier 1 contact is high (any user using a hostile network, repeatedly across the device lifetime); Tier 2 contact is moderate for high-value users; Tier 3 contact is a function of who the target is. The IT-supplied case has the property that the organisation has the ability to change this through procedural and procurement action, so the residual is recoverable.

For BYOD, the residual on each device is similar in mechanism to the IT-supplied case but the organisation has no ability to act on the device. The device's CSME firmware version, AMT provisioning state and Wi-Fi-profile contents are properties of the user's procurement and configuration choices, and the organisation can only control what the device is permitted to reach when it presents itself to the corporate network. The residual is bounded not by what the organisation does to the device, but by what corporate workflows the organisation permits the device to access.

The shared property across both cases is that no software-only control on the device closes the residual to a level that survives a competent adversary. The difference is that in the IT-supplied case the organisation can fix the procedure; in the BYOD case the organisation can only fix the access boundary.

14. Mitigations, by procurement provenance

The mitigations available differ by case. Each control reduces risk; none is a silver bullet.

14.1 Mitigations available in the IT-supplied case once IT becomes aware of the ME

The single most important control is to introduce an interdiction step into the device-issue process. Every laptop is received by IT, not by the user. In a controlled environment the IT team verifies the firmware version and attests it where the platform supports it; enters MEBx and changes the credential; deletes third-party root-CA hashes and replaces them with the organisation's internal CA hash if AMT is to be used (otherwise disables AMT and sets the HAP bit); confirms that the Flash Descriptor write protections are enabled (Section 5.8); disables USB DCI debug; provisions the device under organisational control through Host-Based or Client Control Mode; images the disk; and only then hands the laptop to the user.

The interdiction step closes Scenarios A, C and D substantially against Tiers 1 and 2; closes Scenario B only if firmware integrity is checked rigorously at receipt; closes Scenario F by enabling the Flash Descriptor protection. A Tier-3 supply-chain residual remains, and is the irreducible cost of buying any silicon that ships with a Ring -3 manageability engine.

Cryptographic firmware attestation at receipt takes a known-good measurement of CSME and BIOS regions at the moment of receipt and compares to vendor attestation, reducing Tier 3 and Tier 4 residuals further.

Procurement controls add chain-of-custody constraints, with direct-from-manufacturer purchase, serial-numbered manifests, tamper-evident packaging and drop-shipping to the IT lab rather than to the user; they reduce Tier 4 risk meaningfully but do not eliminate Tier 3.

A scheduled firmware reflash cadence brings devices back through the lab on a cadence aligned to the OEM's capsule release calendar (not Intel's). Section 11 sets out why the OEM's calendar is the binding constraint.

Disabling AMT entirely and setting the HAP bit at first IT touch reduces the attack surface presented by the CSME runtime for most non-managed deployments. Where AMT is required for the deployment use case, the AMT management plane becomes its own asset that is referenced in Section 6 and warrants its own controls.

Discrete TPM in place of PTT, where the platform offers it, separates the TPM by silicon from the CSME, so CVE-2019-0090 and CVE-2021-0146 do not apply identically. This is a procurement decision, not a configuration; a fleet that is already PTT-only cannot be retrofitted.

Eliminating typed-password fallbacks for sensitive credentials, by moving BitLocker to TPM-only or to FIDO2-attested unlock, eliminating password-based fallback for sign-in and forcing smart-card or FIDO2 for VPN, reduces Scenario C credential-capture impact. It does not address Scenario D.

14.2 Mitigations available in the BYOD case

No interdiction is available, and no firmware-push authority exists. The mitigation surface is access-class constraint at the corporate perimeter and the host-level constraints the organisation already imposes.

Conditional access tied to attestation, where the BYOD device is required to produce a host firmware-integrity measurement (TPM PCR baseline, vendor BIOS-attestation token or third-party firmware-monitoring agent measurement) at sign-in and on session refresh, denies access to devices that fail attestation. The control does not eliminate Tier-3 implants that are

designed to subvert the attestation collector, but it closes Tiers 1 and 2 against many configurations.

Refusal of high-sensitivity access classes from BYOD vPro hardware, with corporate-supplied attested hardware required for the affected workflows, is the structural answer. The set of access classes refused is a CISO decision, but the candidates include privileged administrative access, access to source-code repositories, finance and treasury workflows, M&A workspaces, customer personal-data workflows under regulatory regime and any access to systems holding the organisation's own cryptographic key material.

Routing low-sensitivity access through virtual desktop infrastructure or browser isolation places a security perimeter between the BYOD device and the corporate environment. The data the user works on is rendered as pixels on the BYOD device rather than processed on it; a CSME compromise still captures keystrokes and screen content via Scenario C, but the corporate data is not present on the BYOD storage. The control reduces but does not eliminate exposure, and is most useful for read-only or limited-interaction workflows.

Network containment of BYOD vPro traffic to a quarantined VLAN with egress permitted only to the VDI or proxy plane reduces the lateral movement opportunity if a CSME-resident attacker is using the device's authenticated access. Combined with destination filtering, the BYOD device cannot reach internal systems directly even where the user's session credentials have been captured.

Discouraging or prohibiting AMT-over-Wi-Fi on BYOD does not close the unprovisioned-ME risk (the ME is the threat surface regardless of AMT being enabled), but it reduces the scope of credential exposure under Section 5.10 should the user choose to provision AMT for personal management.

User-facing guidance for BYOD users, including a recommendation to apply BIOS updates from the OEM promptly and to disable AMT in MEBx if not used, is hygiene; it has limited bite, because the user is not employed by the organisation in the BYOD case, and the organisation has no enforcement.

14.3 What does not work, in either case

Software-only solutions on the host (MDM enrolment, host-based EDR, host firewall hardening, full-disk encryption with PTT-sealed keys against a Tier 2 or Tier 3 attacker) are valuable for the post-interdiction deployment in the IT-supplied case, and are useful at the policy-enforcement layer in the BYOD case, but they do not fix the architectural problem. The very thing they would protect (the operating system) is the layer the threat model bypasses.

Network L3 ACLs that target the device's IP cannot disambiguate ME-originated from host-originated traffic (Section 7).

Host firewall rules cannot see ME-port traffic (Section 7).

Endpoint EDR cannot see SOL traffic (Section 9.4 and Scenario E).

15. Detection and monitoring, what is actually visible

The honest answer differs between on-corporate-network and roaming, and the standard intuition that host-side defences will catch the threat is wrong for the reasons in Section 7.

15.1 On the corporate network

When the device is on the corporate wired network or associated to a corporate SSID, several signals are obtainable with appropriate instrumentation.

AMT port fingerprints at the perimeter, on a SPAN/TAP or on an NDR sensor are reliable for AMT-shaped traffic. TCP 16992, 16993, 16994 and 16995 carry distinctive HTTP-style or TLS-wrapped management traffic; UDP 623 and the RMCP "Hello" or "Pong" message format fingerprint with high confidence in Suricata, Zeek or commercial NDR. A device talking on those ports without the AMT management plane having initiated the conversation is high-value telemetry.

The unprovisioned-state "Hello" broadcast is a useful detection control. A passive Layer-2 sensor on the broadcast domain catches it, and every "Hello" from a device not on the provisioning manifest is a finding. The control is particularly effective for the IT-supplied case because the manifest is known; for BYOD there is no manifest, and the "Hello" detection becomes a posture signal that the device is in the unprovisioned state.

DHCP-time fingerprinting using Option 60, Option 55 and the OEM hostname pattern in Option 12 identifies factory-state Windows OOBE devices. The same MAC requesting addresses with materially different DHCP fingerprints minutes apart is the inline cohabitation pattern made visible.

TLS metadata on AMT ports (JA3 or JA4 fingerprints of the AMT TLS stack) is stable enough across firmware versions to catalogue and alert on; a connection to TCP 16993 from outside the management plane, or a TLS handshake on 16995 to a destination not in the management VRF, is anomalous.

NAC and 802.1X posture failures act as detection signals. A device that fails 802.1X because it is unprovisioned and lacks the EAP-TLS certificate is quarantined, and the quarantine event is itself a finding. For BYOD the same signal is a posture decision point: the device that cannot attest is the device that is denied access.

A monitoring SSID broadcasting at low power in the office vicinity, mimicking common targets that an injected profile might use (vendor-default SSIDs, "attwifi", "FreeWiFi") and logging association attempts by MAC, catches generic injected profiles. The control is asymmetric: one association from a managed device to a SSID it has no legitimate reason to know is one device that warrants recall and inspection.

Endpoint-side firmware-integrity attestation through Microsoft Defender for Endpoint, Intel CSME Version Detection Tool, vendor BIOS-attestation services and third-party firmware monitoring (Eclypsium, Binarly) reports CSME and BIOS measurements to a corporate collector. Drift between expected and observed is the highest-confidence host-side signal. The agent runs on the host and a Ring -3 implant can suppress its reporting, so the signal is high-confidence against Tier 1 and Tier 2 and one signal among several against Tier 3.

TPM or PTT PCR drift at boot, collected to a remote attestation service, produces a forensic record. A device whose PCRs diverge from baseline is signalling that something below the OS changed since last boot. The signal survives an attacker with Ring -3 capability, because the measurement is taken before the attacker's code regains execution and is delivered to a service the attacker does not control.

DRTM attestation through Windows Virtualization-Based Security on Secured-core PC platforms gives a measured-launch path that is structurally harder for a CSME-resident implant to subvert than the static root of trust.

What remains invisible on-network includes Serial-over-LAN payload content where the attacker uses TLS, traffic that the CSME chooses not to surface to the host (the inline cohabitation property), and devices configured against a hostile profile that have not yet associated.

15.2 When the device is roaming

Real-time monitoring while roaming is essentially impossible, because the corporate network is not in the path. Deferred detection is meaningful and is the gap that lets a roaming compromise stay invisible until it is too late.

Endpoint-side firmware integrity agents that phone home on reconnection capture CSME measurements continuously and forward them to a cloud collector when the device next has internet connectivity.

Hardware attestation on reconnect, with TPM or PTT remote attestation to a corporate verifier on VPN connect, on 802.1X re-auth or on Conditional Access policy evaluation, reveals PCR drift across a roaming window. Where no firmware update event was authorised, drift is direct evidence of unauthorised modification.

Conditional Access policies tied to attestation deny corporate resource access to devices that fail attestation on reconnect. The control does not undo the CSME compromise (data exfiltrated via SOL during the roaming window is already gone), but it stops the next round of corporate exposure and lets the device be removed from service, or, in the BYOD case, denied future access.

Corporate-issued cellular hotspots, where the only sanctioned route to the internet for a high-value user is an MDM-managed LTE or 5G hotspot, collapse "roaming" into "always on corporate transit" for IT-supplied hardware. The control is not available in the BYOD case unless the organisation issues hotspots as a separate corporate asset, which is a procurement decision.

For the IT-supplied case once the procedure is fixed, tamper-evident sealing of the chassis between IT-lab interdiction and user receipt and a forced firmware-reflash cadence convert unknown-state risk into known-clean risk on a schedule. Neither control applies in the BYOD case.

15.3 Note on management infrastructure detection

For the avoidance of doubt, where an organisation operates an enterprise AMT management plane (Intel EMA, MeshCentral, an OEM service or equivalent), that plane is a high-value asset that requires its own monitoring, including session logging, privileged access management and integration with the corporate SIEM. The detection treatment in this document concerns the device-level signals; the management-plane treatment is a separate exercise referenced in Section 6.

16. Risk management framework, treatment, residual, ownership and review

16.1 Treatment per scenario

For each scenario in Section 10, the treatments follow the standard avoid, mitigate, transfer, accept taxonomy. Transfer is of limited applicability because most vendor EULAs explicitly disclaim CSME firmware liability.

Scenario A is mitigated in the IT-supplied case by lab interdiction with AMT disabled or HAP-set, plus on-network detection (15.1). In the BYOD case it is mitigated by attestation-gated conditional access plus access-class refusal (14.2). Owner is IT Security. Review is per Intel Platform Update cycle and per OEM capsule release.

Scenario B is mitigated, with residual accepted, by tamper-evident sealing, firmware-integrity attestation at receipt, attestation on reconnect and scheduled reflash in the IT-supplied case. In the BYOD case the control surface is solely access-class refusal for sensitive workflows and conditional access tied to attestation for permitted workflows. Owner is Procurement and IT Security in the IT-supplied case, IT Security alone in the BYOD case. Acceptance requires written CISO sign-off for any user population in Tier 3 scope.

Scenario C is mitigated in the IT-supplied case by replacing the MEBx root-CA hash list with internal-CA-only at lab interdiction, disabling AMT where not operationally required, and DHCP Option 15 control on corporate infrastructure. In the BYOD case the residual sits at the access boundary, and is bounded by the access classes the BYOD device is permitted to reach. Owner is IT Security. Review is annual, plus on internal-CA incident.

Scenario D is mitigated, with residual accepted, by firmware currency where achievable (Section 11), AMT disabled where not in active use, network-side detection in 15.1, attestation on reconnect in 15.2, and scheduled reflash in the IT-supplied case. In the BYOD case the firmware-currency leg is not available and the access-boundary controls do most of the work. Residual is the unknown N-day. Owner is IT Security. Review is quarterly, with immediate review on any new disclosure of CVE-2019-0090-class severity.

Scenario E is mitigated, where AMT is enabled at all, by strong AMT credential hygiene with rotation, SOL flow detection at the perimeter and conditional access tied to attestation. Where AMT is disabled (the recommended default in both provenance cases for any device not under an active management plane), the scenario does not arise in steady state. The note for sign-off is that the scenario fails the standard host-firewall intuition; the control surface lives at the network and attestation layers.

Scenario F is mitigated in the IT-supplied case by verifying Flash Descriptor write protection at receipt and by eradication standards that include CSME reflash and post-build attestation. In the BYOD case the protection cannot be verified or enforced, and the mitigation reduces to the attestation-and-access-boundary controls in 14.2.

16.2 Residual risk register

Four residuals require explicit acceptance.

A Tier-3 supply-chain interdiction with a firmware-level implant that survives both the receipt-time integrity check and the periodic reflash is the irreducible residual at the top of the register, applicable to both provenance cases. The implant subverts the attestation collector or the reflash path itself, or operates in regions not covered by the attestation. The mitigation ceiling is hardware procurement of platforms with the manageability engine architecturally absent or vendor-attested-removed for highest-value users. Acceptance is explicit and per-affected-user, with the CISO as the named risk owner.

An unknown CSME N-day exploited on a hostile network the user roams onto sits as a permanent residual. Mitigation ceiling is firmware currency where achievable, AMT exposure reduction and attestation on reconnect. In the BYOD case the firmware-currency leg is largely unavailable per Section 11. Acceptance is organisational, with annual re-rating.

A PTT or CSME shared-runtime trust failure of the CVE-2019-0090 class on a hardware generation in the active fleet sits as a residual whose mitigation ceiling is a discrete TPM where the platform offers it. Acceptance governs the choice between BitLocker TPM-only and BitLocker TPM-plus-PIN; TPM-plus-PIN is recommended because the PIN survives a PTT key extraction.

An insider or coercion event at a CA holding the AMT-OID issuance privilege, where the provisioning gate of 5.5 was being relied upon, sits as a residual whose mitigation ceiling, in the IT-supplied case, is replacement of the firmware trust list with an internal CA only at IT lab interdiction. In the BYOD case this control is not available; the residual is bounded by the access boundary.

16.3 Review cadence

Quarterly review of the Intel Security Center advisory feed, re-rating any open scenarios whose mitigation depends on firmware currency, with attestation baselines updated for any in-scope device that received a CSME update in-cycle.

Per-procurement-cycle review of the supply-chain treatment when a new vendor, reseller, courier, customs path or contract manufacturer enters the chain (IT-supplied case).

Per-significant-CVE-disclosure review when a CVE-2019-0090-class issue is disclosed against a generation in the active fleet, with immediate residual re-rating, CISO notification and a decision on whether to accelerate the next reflash cycle (IT-supplied case), or to revoke conditional-access permission classes (BYOD case).

Annual review of whether the Tier 3 in-scope user population has changed (new executives, M&A activity, regulated workloads), with hardware procurement, reflash cadence and BYOD access-class policy adjusted accordingly.

16.4 Incident response, what "we suspect ME compromise" means operationally

Forensic preservation depends on capability that most enterprise IR teams do not currently have in-house. Standard host-forensic toolchains do not extract CSME state. Capturing CSME firmware regions requires a Dediprog or equivalent SPI-flash dumper, a controlled environment and skill in firmware reverse engineering. If Tier 3 is in scope, building or contracting for the capability is part of the cost of being prepared. In the BYOD case, the organisation does not own the device and cannot in general insist on a SPI-flash dump; the IR posture is necessarily containment-and-revocation rather than forensic-and-eradication.

Containment by network-side blackholing of the device's MAC and IP at the corporate perimeter is necessary but not sufficient, because the attacker may already have exfiltrated via SOL during the roaming window. In the IT-supplied case the device is recovered. In the BYOD case the device's access is revoked at the conditional-access plane and the user is informed; subsequent access requires the user to demonstrate a clean device, which in practice means a new device or a verified firmware reflash that the user undertakes themselves.

Treatment of disk data starts from the position that disk contents are exposed. If the suspected compromise includes Scenario D and the platform uses PTT, BitLocker provides no protection. If the suspected compromise is Scenario C only (KVM credential capture), disk contents are exposed via the captured PIN, with the same outcome. In the BYOD case the disk is the user's, but any corporate data that was on the disk is treated as exposed.

Treatment of credentials is rotation across the board: passwords, BitLocker PINs and recovery keys entered on the device during the suspected window; SAM hashes and cached domain credentials. FIDO2 hardware tokens that authenticated against the device are retained for forensics and the cryptography is intact, because the FIDO2 challenge-response protocol does not transmit anything that captures a usable secret; the user can continue using the same hardware token on other devices. Smart-card PINs are rotated.

Eradication in the IT-supplied case is a CSME reflash to a verified-clean image, MEBx reset, full disk wipe and reimage and post-build firmware-integrity attestation against the rebuilt baseline. A device that fails post-eradication attestation is destroyed, not redeployed. In the BYOD case the organisation cannot eradicate; it can only revoke access until the user has independently established a clean device.

Communication to the affected user, the data-protection officer and, where applicable, the regulator cannot be conditioned on host-side IOCs, because the CSME compromise model includes the possibility that the host EDR did not see anything. The standard form is that the architectural exposure existed, the device is not provably clean and notification is being given because the relevant standard (GDPR Article 33 or equivalent) requires it where likelihood of compromise cannot be excluded.

Lessons-learned triggers re-rating of the Tier model in Section 4 and re-evaluation of whether the affected user population should be moved off vPro hardware (IT-supplied case) or off BYOD permissions for the affected workflows (BYOD case).

17. Recommendation, by procurement provenance

The evidence supports a single defensible co nclusion that is then differentiated by what the organisation can do about it. Connecting an untouched-ME vPro laptop to corporate resources, in either provenance case, is not an acceptable security posture without compensating controls, and there is no software-only mitigation on the device that closes the gap.

17.1 For the IT-supplied / ME-naive case

The single largest improvement is procedural: the IT function must become aware of the ME, must accept that factory-state vPro hardware is not equivalent to factory-state non-vPro hardware, and must introduce an interdiction step into the device-issue process. The interdiction step is the gold-standard mitigation set in Section 14.1, in summary MEBx lockdown, AMT disable (or properly provisioned to corporate control), HAP bit set, Flash Descriptor write protection verified, internal CA replacing the factory hash list where AMT is to be retained, USB DCI debug disabled, firmware version validated and host-OS imaged, before user handover.

The procurement function adopts firmware-integrity attestation at receipt as a standard requirement and selects OEMs that publish CSME-bearing capsules promptly as a procurement criterion (Section 11).

The on-network detection programme in Section 15.1, including AMT-port flow signatures, RMCP "Hello" detection, DHCP fingerprinting, NAC posture, honeypot SSID and firmware-integrity attestation, is implemented as a connected programme rather than as individual controls.

The roaming detection programme in Section 15.2, including attestation on reconnect, scheduled reflash, tamper-evident seals and corporate cellular hotspots for high-value users, is implemented for the in-scope user populations.

The incident response capability in Section 16.4, including SPI-flash dumping where Tier 3 is in scope, is stood up. Without it, the controls above are detection without response.

The operational cost of all of the above is budgeted for every device, every refresh cycle. The "ship straight to the user and remote-provision them later" model is the model the assessment recommends against, and the existing fleet is brought through the lab on a scheduled cadence.

17.2 For the BYOD case

The organisation does not control the device, does not control its firmware version, does not control its provisioning state and cannot verify its supply chain. The recommendation is to govern what the device is permitted to reach.

High-sensitivity access classes are refused from BYOD vPro hardware. The set is a CISO decision; candidates include privileged administrative access, source-code repositories, finance and treasury, M&A workspaces, regulated personal-data workflows and any access to systems holding the organisation's cryptographic key material. For these access classes a corporate-supplied attested device is required.

Permitted access classes from BYOD vPro hardware are gated by conditional access tied to firmware integrity attestation, with denial of access on attestation failure. Access is routed where feasible through VDI or browser isolation so that corporate data is rendered rather than processed on the BYOD device.

Network containment of BYOD vPro traffic to a quarantined VLAN with egress permitted only to the VDI or proxy plane reduces lateral movement opportunity if a CSME-resident attacker is using the device's authenticated access.

User-facing guidance, including a recommendation to apply BIOS updates from the OEM promptly and to disable AMT in MEBx if not used, is published as part of the BYOD onboarding pack, while recognising that the organisation has no enforcement.

The organisation accepts that the residual on each BYOD vPro device is bounded by the access boundary rather than by any state of the device, and treats every BYOD vPro device as a permanently untrusted member of a permanently untrusted population.

17.3 Shared residual that requires explicit CISO acceptance

A Tier-3 supply-chain residual remains in both cases and is the irreducible cost of buying any silicon that ships with a Ring -3 manageability engine. For the highest-value users in either case (executives, M&A, legal, finance, infrastructure-engineering admins), procurement of hardware with the manageability engine architecturally absent or vendor-attested-removed, or non-Intel architectures where the threat model is materially different, should be considered.

A vendor-friendly framing that "the gates are cryptographic and unbreakable" should not be relied upon. The evidence does not support it. The residual register in Section 16.2 contains the four positions that the framing avoids.

The summary, in one sentence with no dramatic flourish, is that the cost of acting on this assessment is the cost of a small interdiction lab plus the network-detection and attestation infrastructure (in the IT-supplied case), or the cost of access-class governance and conditional-access tooling (in the BYOD case), and the cost of not acting is one Tier-2 incident at any point in the lifecycle of any laptop in scope, executed via an attack pattern that has already been demonstrated in the wild.

Appendix A. Sources and CVE register

Foundational research: Tereshkin and Wojtczuk, *Introducing Ring -3 Rootkits*, Black Hat USA 2009 (PDF); Skochinsky, *Intel ME Secrets*, REcon 2014 (PDF); Skochinsky, *Intel ME, Two Years Later*, Breakpoint 2014 (PDF); Goryachy and Ermolov, *How to Hack a Turned-Off Computer*, Black Hat Europe 2017 (PDF).

CVE and advisory record: CVE-2017-5689, INTEL-SA-00075, "Silent Bob is Silent", "admin no password" (NVD; Embedi PoC; The Register); INTEL-SA-00086 (CVE-2017-5705/5706/5707) (Intel; Intel-SA-00086 support article on update path); CVE-2019-0090 (PT Security; SecPod; TechTarget); CVE-2021-0146 (Help Net Security; SecurityWeek); INTEL-SA-00213 (Intel); INTEL-SA-00295 (Intel); CVE-2025-20037 (ZeroPath); INTEL-SA-01138, 01200, 01280 and the Intel Security Center index; HAP and AltMeDisable (PT Security).

Realised in-the-wild attack and feature abuse: Microsoft on PLATINUM and AMT SOL (Microsoft Security Blog; Threatpost; BleepingComputer; The Hacker News); F-Secure thirty-second AMT exploit (BankInfoSecurity); CSO Online on Eclypsium fleet-vulnerability statistics and Conti ransomware-group AMT exploit research (CSO Online); ASUS Operation ShadowHammer (IEEE Spectrum); Bloomberg, *The Long Hack* (Bloomberg, reporting contested).

Vendor framing of the same architecture (for comparison): Intel, *vPro Security Overview*, December 2024 (Intel white paper PDF).

Hardcoded trust anchor parallel: CVE-2025-20309 Cisco Unified CM static SSH credentials (Cisco Security Advisory; Arctic Wolf; BleepingComputer; CSO Online); CVE-2021-40119 Cisco Policy Suite static SSH keys (The Hacker News); CVE-2021-34795 Cisco Catalyst PON; Schneier, *Cisco Can't Stop Using Hard-Coded Passwords* (Schneier on Security).

Supply-chain doctrine and risk frameworks: NSA ANT Catalog (Wikipedia; EFF document mirror PDF); The Intercept, *Everybody Does It* (The Intercept); Microsoft, *Guarding Against Supply Chain Attacks Part 2* (Microsoft Security Blog); Eclypsium, *Firmware Attacks: An Endpoint Timeline* (Eclypsium); Eclypsium, *Firmware Security Realizations Part 2* (Eclypsium); Dark Reading, *Firmware Vulnerabilities Continue to Plague Supply Chain* (Dark Reading).

Architecture and AMT mechanics: Intel Management Engine (Wikipedia); Intel Active Management Technology (Wikipedia); Intel AMT KVM developer documentation (Intel); Intel AMT provisioning certificate guide and OID 2.16.840.1.113741.1.2.3 (Intel software manageability); Intel AMT IP-address sharing and DHCP cohabitation (Intel software manageability; IP synchronisation; Intel community thread on shared MAC); Matthew Garrett, *Intel's remote AMT vulnerability* on host-firewall-bypass (mjg59); Intel MEBx documentation including default credential (Intel support).

Integrated radio architecture: CNVi (Wikipedia); Intel, *What Are the Intel Integrated Connectivity (CNVi) and Companion RF (CRF) Module* (Intel support).

Firmware update path constraints: Linux Vendor Firmware Service discussion 9690 on CSME firmware delivery via LVFS (fwupd); Intel community discussion of FWUpdLcl downgrade behaviour (Intel community); Intel community on CSME firmware update tools (Intel community); Win-Raid forum threads on cross-signing constraints (Win-Raid).

Modern Standby and battery behaviour: Microsoft, *Modern Standby* learn pages (Microsoft Learn-be-causing-my-laptop-t)); XDA Developers, *Modern Standby is draining your Windows 11 laptop battery* (XDA Developers); Windows Latest, February 2026 reporting (Windows Latest).

Appendix B. Open questions

The most recent CSME advisory in the public record at the time of writing (CVE-2025-20037 and the INTEL-SA-01138, 01200, 01280 cluster) is used as the representative example of the ongoing advisory cadence. If the organisation has knowledge of a more recent specific advisory not yet captured in industry coverage, it should be inserted into Section 9.3 in addition to the existing references.

The fleet's hardware generation (ME or CSE 16.x, 17.x, 18.x, presence of discrete versus firmware TPM, whether HAP is settable in BIOS, whether USB DCI is fused off) determines exactly which paragraphs of Section 9 apply most directly. A second pass with the actual model and firmware version in hand is worthwhile but not required for the recommendations in Section 17 to stand.

The mitigation recommendations in Section 17 are calibrated to a