Physical Security

SL5 Novel Recommendations

NOVEMBER 2025

Luis Cosio

Lisa Thiergart

Yoav Tzfati

Peter Wagstaff

Guy

Philip Reiner

Overview

Executive Summary

"If a bad actor has unrestricted physical access to your computer, it's not your computer anymore." - 10 Immutable Laws of Security

No system can be secure without securing the physical devices that implement it. While physical attacks may be more costly (and dangerous) for attackers than other methods, they still pose a critical risk. The importance of physical security is heightened by the realities of imperfect personnel security. It is not enough to deny facility access to outsiders, SL5 must also guard against improper physical access by compromised insiders. Additionally, SL5 must consider physical channels like acoustic or RF noise that might provide a nearby attacker with compromising information without direct physical access. The physical security of current commercial data centers is insufficient to guard against a nation-state threat. Moving from a compliance-based design to a performance-based design that is effective against the defined level of threat is key. We present five recommendations to help close this gap.

Top 5 Recommendations

1. Enforce two-person integrity (TPI) for all maintenance
2. Mandate body-worn cameras with obstruction alarms (prohibited in SCIF-equivalent spaces)
3. Eliminate out-of-facility maintenance
4. Harden side-channel defenses (LED masking, additional noise generators and shielding display cables)
5. Perform pre-installation component X-ray verification

Enforce two-person integrity (TPI) for all maintenance

Explanation of Recommendation

Conduct maintenance and other activities requiring physical access to systems in pairs. Each pair contains an operator responsible for performing tasks and a witness responsible for ensuring integrity. No one is permitted to enter or remain in sensitive areas alone. Assignment of pairs is done in a manner that cannot be controlled or predicted by either member.

Research Reasoning

Enforcing Two-Person Integrity (TPI) greatly mitigates the risk of single-actor insider threat attacks and forces attackers to compromise and coordinate at least two insiders, significantly increasing difficulty and risk of detection.

In high-consequence domains, TPI is used to prevent a lone actor from performing an incorrect or unauthorized action. The U.S. Air Force applies this rigorously to nuclear surety operations [1], where at least two authorized individuals must concur and act together. Procedures explicitly exist to ensure no lone individual can execute critical steps on weapons or certified components. In COMSEC (communications security) programs, TPI is formally defined as a storage/handling system that prohibits individual access to certain material by requiring at least two authorized persons. This is codified in COMSEC doctrine; see the NIST CSRC glossary entry for “two-person integrity” [12]. In the private sector, PCI DSS (Payment Card Industry Data Security Standard) requires split knowledge and dual control for manual clear-text cryptographic key-management operations, an instance of the same principle [13].

Under SL5, TPI should be required for any task that could materially change availability, integrity, confidentiality, or auditability. Examples include: entering white space/MMR/electrical/mechanical rooms to perform work, IDS/ACS shunt or maintenance-mode activation/deactivation, racking/deracking servers, storage sled moves, media insertion/removal, core/leaf/spine cabling changes, BIOS/BMC or HSM actions at a local console, key ceremonies and cryptographic device (HSM/KMS) rekeys, PDU/UPS changes or EPO testing/disable/enable, connecting diagnostic/analysis tools to classified systems, vendor maintenance under escort.

Pairs are separated into operators and witnesses so there is a clear understanding of who is responsible for completing the task and who is responsible for monitoring it. Without this distinction, both may perform a task with limited awareness of the other. For a given task, the witness must be able to understand what the operator should be doing and is doing to ensure integrity. Therefore, the witness should be at least as qualified to perform the task themselves as the operator. There is a danger that two insider threats may collude to be in the same pair. To prevent this, a system should be developed for assigning pairs that cannot be controlled or predicted by either member.

Cost-Benefit

- **Costs:** Doubles personnel overhead for maintenance, longer maintenance windows, potential morale impact.
- **Benefits**: Eliminates single-point human failure, creates audit trail, deters insider threats.

- **Alternatives**: Traditional badge access with remote monitoring (can be spoofed in AGI scenarios).

Mandate body-worn cameras with obstruction alarms

Explanation of Recommendation

Require all authorized personnel performing maintenance in vital areas (excluding SCIF-equivalent spaces) must wear tamper-resistant, body-worn cameras that continuously record to encrypted storage. Devices must provide obstruction/tamper detection (e.g., lens occlusion, removal, inversion, or recording interruption) with local and logged alarms. A visible recording indicator is mandatory whenever recording/transmitting, consistent with ICD/ICS 705 technical specifications [11]. Body-worn cameras and any video recording devices are prohibited in SCIF-equivalent spaces.

Research Reasoning

Traditional surveillance systems rely upon fixed cameras that have limited perspective and may create blind spots a capable adversary can map and avoid. The 2013 Snowden revelations showed how insider threats can operate undetected even in highly monitored environments [9]. Fixed cameras also watch from a distance. To properly capture maintenance activities, an up close, first-person perspective is required. Body cameras provide this.

Conventional body cameras, most commonly associated with use in law enforcement, are mounted on the chest. Head-mounted body cameras (e.g., helmet-mounted units used in industrial maintenance), while less common, offer a perspective that better tracks activity. One issue body cameras present is that the person they are recording could inadvertently or intentionally obstruct their view. For this reason, body cameras should include obstruction detection and alerting features (e.g., lens-occlusion, inversion, or recording-interruption detection). Similarly, the wearer has physical access to the camera and could tamper with it, so body cameras must implement tamper-evident features (e.g., sealed housings and removal sensors).

A major advantage of fixed cameras is they can easily be wired to a surveillance system. Body cameras by contrast cannot be wired without severely limiting the mobility of the user. Allowing wireless communications is also in tension with other physical security requirements; in SCIF-like spaces, wireless technologies are restricted or prohibited under ICD/ICS 705 and TEMPEST principles [10], [11].

In SCIF-equivalent areas, body-worn cameras—and any video recording equipment—are prohibited. In non-SCIF vital areas, cameras may operate untethered during tasks but must dock at defined intervals to upload footage; missed upload deadlines trigger audible/visible alerts. Audio may be enabled per policy in non-SCIF areas if side-channel considerations are addressed via the noise-masking controls in “Side-channel defense.”

Cost-Benefit

- **Costs:** Camera hardware, encrypted storage infrastructure, battery management logistics, privacy training for personnel.

- **Benefits:** Materially reduces blind spots, provides evidence for incident investigation, creates strong deterrent effect, enables training from real maintenance footage, protects personnel from false accusations.
- **Alternatives:** Fixed CCTV only (creates blind spots), manual logging (unreliable, forgeable), honor system (insufficient for nation-state threats).

Eliminate out-of-facility maintenance

Explanation of Recommendation

Prohibit out-of-facility maintenance, adjustments, and system modifications. Out-of-facility maintenance is defined as any operation initiated remotely (outside the physical facility perimeter), by external vendors via remote connection, or from any network location outside the air-gapped Internal Network dedicated to the facility. All actions that must be performed should be performed through the isolated Internal Network or if it needs to be done physically, by authorized personnel following Two-Person Integrity (TPI) protocols. This includes HVAC systems, power distribution units, UPS, ATS and switchgear, generators, cooling systems, fire detection and suppression, physical security devices, servers, and all infrastructure equipment. For external vendors, no remote access capabilities, network management interfaces, or out-of-band management ports are permitted. Any vendor support requiring remote access must be denied; vendors must either dispatch technicians or provide guidance to on-site staff. If monitoring visibility is required out-of-facility, permit one-way egress only (telemetry only, no command channel) via a hardware data diode or unidirectional gateway with no bidirectional sessions [18], [23], [24], [21]. For guidance on how the Internal Network should be designed, refer to the Network Security memo.

Research Reasoning

Treat remote control planes as systematically unsafe at SL5. Remote channels create **keyboard-to-kinetic** pathways that have been realized in practice (German steel mill; Oldsmar) [4], [6]. Remote Monitoring and Management (RMM) tools are routinely weaponized at scale against enterprises and governments [15]. Out-of-band/management planes (e.g., BMC/IPMI/Redfish/AMT/ISM) have had critical remotely exploitable flaws (e.g., CVE-2017-5689) that bypass host OS controls [14]. Centralized update paths can produce global-scale blast radii via compromised or defective updates (SolarWinds supply-chain compromise; 2024 CrowdStrike content update outage) [16], [17]. In aggregate, any residual remote capability functions as a privileged bypass around physical protections and Two-Person Integrity.

Eliminate inbound control channels rather than “hardening” them to achieve a **qualitative** risk reduction: reducing whole attack classes (RMM abuse, BMC/OOB compromise, remote mass-change cascades) instead of relying on perfect configuration. Authoritative OT guidance (NIST SP 800-82r3) treats remote access as **high-risk** and endorses **unidirectional gateways (data diodes)** for safe, one-way telemetry [18]. IEEE cybersecurity standards for substation automation emphasize strong electronic access controls, including disabling unused remote interfaces and enforcing authenticated access for IEDs (intelligent electronic devices) [23], [24]. While these standards focus on securing remote access, the SL5 framework adopts a more restrictive posture by prohibiting remote maintenance entirely due to the demonstrated and expected capabilities of nation states to easily exploit remote access [18], [23], [24], [21]. For UPS and power-chain equipment specifically, active targeting of internet-connected UPS devices by threat actors demonstrates why external vendor remote management is unacceptable at SL5 [2]. Where communication security on control links is unavoidable in other domains, the field employs cryptographic protections (e.g., IEEE 1711 for substation serial links), but at SL5 we further constrain risk by limiting equipment to one-way telemetry via a hardware data diode [18], [21].

Cost-Benefit

- **Costs:** 24/7 on-site maintenance staff, increased maintenance response time for issues, higher vendor support costs, added operational burden of predictive maintenance workflows, inability to leverage remote expertise.
- **Benefits:** Reduces remote attack surface entirely, forces physical attribution for all changes, creates natural rate-limiting on system modifications, enables visual verification of all maintenance, prevents cascade failures from remote mass-changes. Deletes entire classes of failure (RMM abuse, BMC/OOB compromise, mass-update cascades), introducing human friction and physical attribution by design.
- **Alternatives:** Secured jump boxes, time-limited remote sessions, vendor-managed remote access, and software-defined perimeters still leave exploitable control planes and supply-chain/update risks; if visibility is necessary, use one-way telemetry via a data diode and vendor phone/email guidance to on-site staff with no remote control channel [2], [21], [23].

Harden side-channel defenses (LED masking, additional noise generators and shielding display cables)

Explanation of Recommendation

Guard against unintended leakage of compromising information through physical channels, targeting both attackers trying to compromise the system or an AI trying to exfiltrate itself or establish a channel to the outside. Consider a combination of methods including software hardening, signal attenuation, and noise generation. Consider mediums including acoustic, power, RF, and light. Also consider side channels like caches and timing.

Research Reasoning

Two distinct threats should be considered when defending side-channels. First, an attacker using side channels to indirectly extract confidential information from a machine. Attacks of this variety include extraction of private keys, extraction of algorithmic insights in an unencrypted state, or reconstruction of display content [3], [20]. The second threat to consider is an AI using side channels to establish covert channels to communicate with the outside, or even exfiltrate itself.

Side channels to defend include power, acoustic, RF, optical, timing, and more. In each case, there are generally three strategies that can be used to defend the side-channel. First, a side-channel only reveals sensitive information if it behaves differently based on different internal states of the system. For example, operations on a CPU may cause data-dependent power/frequency behavior [8]. An attacker may combine this information with knowledge of the program being run, for example an encryption algorithm, to guess the internal state. Therefore, a potential defense is ensuring that sensitive functions have consistent operation regardless of inputs and outputs, so leakage contains no useful information. This is often done in practice when such vulnerabilities are demonstrated, because it is far more scalable than eliminating physical leakage [20]. Techniques of this type include **masking and blinding** [20].

Hardening software to prevent compromising information from leaking is not a panacea. It has performance trade-offs, and it may not be possible to prevent all useful information from leaking

[8], [10]. In the context of preventing an AI from establishing a covert channel, this approach may not even be possible, as an AI may deliberately create a leak.

As part of defense in depth, physical leakage should be prevented through **attenuation and noise generation**. Attenuation weakens or eliminates physical leakage. A classic example is RF shielding [10] (e.g., via a Faraday cage).
Something as simple as taping over (or removing) LEDs can remove a high-bandwidth channel [22]. Conversely, some channels may be more difficult to attenuate, like acoustics [22].

Where attenuation is insufficient, it may be necessary to combine it with noise generation. Noise generation involves producing noise on the same channel to drown out whatever signal an attacker might be using. It is perhaps the least attractive option since it has the greatest side effects. These include drowning out legitimate communications over the same channel and inefficient power use [20]. Improperly implemented noise generation can also be overcome by extracting the signal from the noise. However, it can still be an effective approach [10].

We recommend that while still implementing more conventional software-based mitigations, it is necessary—especially to prevent an AI from establishing external communications—to physically attenuate or drown out all possible leakages. There are considerations about at which barriers these leakages should be stopped. In some cases, it should be sufficient to **stop a signal from leaving the building**. In other cases, it may be necessary to **prevent illicit interactions between two systems within the same building**. Special consideration should be given to signals that deliberately leave the building, specifically communications between locations made by inline network encryptors (See the Network Security SL5 Novel Recommendation document).

Cost-Benefit

- **Costs:** Noise generator installation and calibration, LED masking labor, display cable replacement, ongoing verification of countermeasure effectiveness, potential ergonomic impact from noise levels.
- **Benefits:** Defeats sophisticated electromagnetic surveillance, prevents optical exfiltration, masks acoustic signatures
- **Alternatives:** Faraday cage construction (expensive, operationally limiting), TEMPEST-certified equipment (limited availability, extreme cost), selective shielding (incomplete protection), periodic security audits (reactive not preventive).

Perform pre-installation component X-ray verification

Explanation of Recommendation

Subject all hardware components to multi-modal inspection such as **combining X-ray imaging, infrared thermography, and computer vision analysis before installation**. This inspection aims to detect unauthorized modifications including additional or substituted chips, anomalous circuits, and other indications of counterfeiting. Even seemingly benign items like cables must pass inspection to enter inventory.

Research Reasoning

Hardware implants, especially by a sophisticated attacker, can be small and difficult to detect. The 2018 Bloomberg 'Big Hack' story [5], though disputed by Apple, Amazon, and Supermicro, demonstrated the feasibility of chip-scale implants. Other examples include NSA's COTTONMOUTH and FIREWALK implants revealed in the ANT catalog [7]. Such implants can enable covert RF beacons, out-of-band management backdoors, keylogging/sniffing taps, and power-supply manipulation [7]. It is crucial to detect implants, counterfeits, and other illicitly modified hardware before it is installed.

This task calls for **efficient, non-destructive techniques for imaging computer hardware**. X-ray inspection, once sidelined for high cost and low speed, has been increasingly used by manufacturers to inspect chips with increasingly complicated and layered packaging [19].

The technology has seen massive improvements in cost, speed, and non-destructiveness, and can reveal small defects and anomalies within the interior of components.

While powerful X-ray has certain limitations, especially when distinguishing materials of similar density [19], infrared imaging can supplement it by revealing anomalous thermal signatures during operation, which requires a test harness as part of the inspection process. Additionally, computer vision techniques in the visible spectrum can identify subtle visual indicators like probe marks or resoldering evidence.

In addition to routine and efficient techniques for inspecting components at scale, some small percentage of **components may be randomly selected for a more thorough search**. Provided significantly more work and disassembly, it may even be possible for techniques like ptychographic laminography to view the entire structure of chips, down to single-digit nanometer features, a process that would only be practical for a tiny number of chips. It may also be possible for certain simple, highly trusted components like hardware signing keys (see Machine Security memo) to be **deliberately built with larger architectures to enable a thorough scan of each**.

Cost-Benefit

- **Costs:** Multi-modal inspection system, **reference signature database development**, inspection time, specialized operator training, false positive investigation overhead, secure storage for reference samples.
- **Benefits:** Detects sophisticated hardware implants, identifies counterfeit components before deployment, provides forensic evidence trail, deters supply chain attacks through random inspection threat, enables automated verification at scale, protects against probe marks and physical tampering often missed by visual inspection.
- **Alternatives:** Visual inspection only (misses internal modifications and subtle tampering), trust in secure supply chain (vulnerable to compromise), vendor attestation only (insufficient against nation-state supply chain compromise), tamper-evident packaging at receiving (with serialized seals).

Bibliography

[1] "BY ORDER OF THE SECRETARY OF THE AIR FORCE," Air Force Instruction 91-114. [Online]. Available: https://static.e-publishing.af.mil/production/1/af_se/publication/afi91-114/afi91-114.pdf

[2] Cybersecurity and Infrastructure Security Agency (CISA) and U.S. Department of Energy (DOE), "Mitigating Attacks Against Uninterruptible Power Supply (UPS) Devices," Mar. 29, 2022. [Online]. Available: https://www.cisa.gov/news-events/alerts/2022/03/29/mitigating-attacks-against-uninterruptable-power-supply-devices

[3] M. Kuhn, "Compromising emanations: eavesdropping risks of computer displays," University of Cambridge, Computer Laboratory, Tech. Rep. UCAM-CL-TR-577, 2003. [Online]. Available: https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-577.pdf

[4] R. Lee, M. Assante, and T. Conway, "German Steel Mill Cyber Attack," SANS ICS, 2014. [Online]. Available: https://assets.contentstack.io/v3/assets/blt36c2e63521272fdc/blt5bd1acefa6ad7c17/6323756c1ad88e716559ed66/ICS-UseCase2-ICS-CPPE-case-Study-2-German-Steelworks_Facility.pdf

[5] J. Robertson and M. Riley, "The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies," Bloomberg Businessweek, Oct. 04, 2018. [Online]. Available: https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies

[6] F. Robles and N. Perlroth, "'Dangerous Stuff': Hackers Tried to Poison Water Supply of Florida Town," The New York Times, Feb. 09, 2021. [Online]. Available: https://www.nytimes.com/2021/02/08/us/oldsmar-florida-water-supply-hack.html

[7] J. Appelbaum, J. Horchert, and C. Stöcker, "Catalog Reveals NSA Has Back Doors for Numerous Devices," Der Spiegel, Dec. 29, 2013. [Online]. Available: https://www.spiegel.de/international/world/catalog-reveals-nsa-has-back-doors-for-numerous-devices-a-940994.html

[8] Y. Wang et al., "Hertzbleed: Turning Power Side-Channel Attacks Into Remote Timing Attacks on x86," in Proc. 31st USENIX Security Symposium, 2022. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/wang-yingchen

[9] G. Greenwald, "NSA collecting phone records of millions of Verizon customers daily," The Guardian, Jun. 06, 2013. [Online]. Available: https://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order

[10] "TEMPEST: A Signal Problem," NSA, Cryptologic Spectrum, Vol. 2, No. 3, 1972, declassified 2007. [Online]. Available: https://www.nsa.gov/portals/75/documents/news-features/declassified-documents/cryptologic-spectrum/tempest.pdf

[11] Office of the Director of National Intelligence, National Counterintelligence and Security Center, Technical Specifications for Construction and Management of Sensitive Compartmented Information Facilities (IC Tech Spec—for ICD/ICS 705), Version 1.5, NCSC 19-686, Washington, DC, USA, Mar. 13, 2020.

[12] National Institute of Standards and Technology (NIST), “two-person integrity,” Computer Security Resource Center (CSRC) Glossary. [Online]. Available: https://csrc.nist.gov/glossary/term/two_person_integrity

[13] PCI Security Standards Council, Payment Card Industry Data Security Standard: Requirements and Testing Procedures, Version 4.0.1, Jun. 2024. [Online]. Available: https://www.pcisecuritystandards.org/document_library/. (See §3.7.6 on split knowledge and dual control.)

[14] NIST, "NVD: CVE-2017-5689 (Intel AMT/ISM/SBT)," 2017. [Online]. Available: https://nvd.nist.gov/vuln/detail/CVE-2017-5689

[15] CISA, NSA, and MS-ISAC, "Protecting Against Malicious Use of Remote Monitoring and Management Software (AA23-025A)," Jan. 26, 2023. [Online]. Available: https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-025a

[16] CISA, "Advanced Persistent Threat Compromise of Government Agencies, Critical Infrastructure, and Private Sector Organizations via SolarWinds Orion (AA20-352A)," Dec. 17, 2020 (updated Jan. 6, 2021). [Online]. Available: https://www.cisa.gov/news-events/cybersecurity-advisories/aa20-352a

[17] CrowdStrike, "Falcon Content Update Preliminary Post Incident Report," Jul. 2024. [Online]. Available: https://www.crowdstrike.com/en-us/blog/falcon-content-update-preliminary-post-incident-report/ . (See also Microsoft's summary blog and major press coverage.)

[18] NIST, SP 800-82 Rev. 3: Guide to Operational Technology (OT) Security, Sep. 2023. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-82r3. Key guidance: remote access only if justified/limited; unidirectional gateways/data diodes for one-way flows (pp. 122–123, 212–213; 89; 223).

[19] National Institute of Standards and Technology, "NIST Scientists Develop Novel CT Scan Device for Integrated Circuits," NIST News, Apr. 14, 2023. [Online]. Available: https://www.nist.gov/news-events/news/2023/04/nist-scientists-develop-novel-ct-scan-device-integrated-circuits

[20] D. Genkin, A. Shamir, and E. Tromer, "RSA Key Extraction via Low-Bandwidth Acoustic Cryptanalysis — Q&A," 2014. [Online]. Available: https://cs-people.bu.edu/tromer/acoustic/

[21] IEEE Standards Association, "IEEE Trial-Use Standard for a Cryptographic Protocol for Cyber Security of Substation Serial Links," IEEE Std 1711-2010, Nov. 2010. [Online]. Available: https://ieeexplore.ieee.org/document/5715000

[22] M. Guri, "ETHERLED: Sending Covert Morse Signals from Air-Gapped Devices via Network Card LEDs," arXiv:2208.09975, 2022. [Online]. Available: https://arxiv.org/pdf/2208.09975

[23] IEEE Standards Association, "IEEE Standard for Intelligent Electronic Devices Cyber Security Capabilities," IEEE Std 1686-2013, Dec. 2013. [Online]. Available: https://ieeexplore.ieee.org/document/6704702

[24] IEEE Power & Energy Society, "IEEE Standard Cybersecurity Requirements for Substation Automation, Protection, and Control Systems," IEEE Std C37.240-2014, Dec. 2014. [Online]. Available: https://ieeexplore.ieee.org/document/7024885