What is secure communication in a substation?

Author: Gaurav Jagadeesh

Substation communication is hard to secure because a single protocol does several jobs at once. Separate those jobs, give each its own boundary, and the secure design follows.

MMS (Manufacturing Message Specification, ISO 9506) carries client–server traffic in IEC 61850 substations. It was last revised in 2003. In *MMS: A Protocol That Refuses to Retire*, we showed that every modern need, from encryption to enterprise integration, is added around MMS rather than inside it.

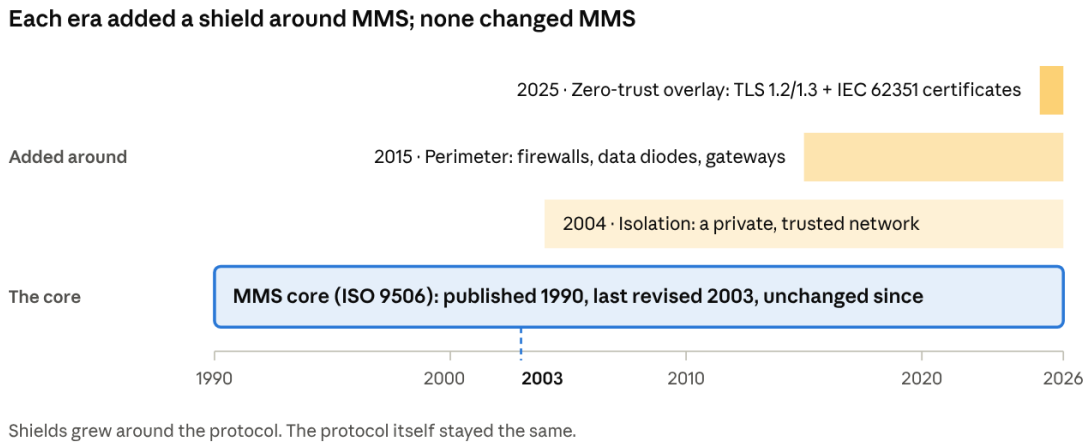


Figure 1. Security added around MMS by era; the core never changed

This post starts from a blank page. It states three principles and applies them in order to derive the design. Scope: client–server communication only. But before designing anything, we need to know what we are defending against.

What are we defending against?

Attacks on power grids have so far mostly aimed to disrupt supply, not to destroy equipment, and they used commands the network accepted as legitimate.

- **Ukraine, December 2015:** about 225,000 customers lost power. Attackers phished staff, stole VPN credentials (the VPN had no two-factor authentication), passed a firewall that allowed remote administration, and opened breakers from the operators' own HMIs, then wiped workstations and corrupted substation device firmware to slow recovery (POWER, CCDCOE).
- **Ukraine, December 2016:** Industroyer cut power to part of Kyiv for about an hour. It spoke IEC 61850, IEC 60870-5-101/104 and OPC DA, and could make Siemens SIPROTEC relays unresponsive (ESET).

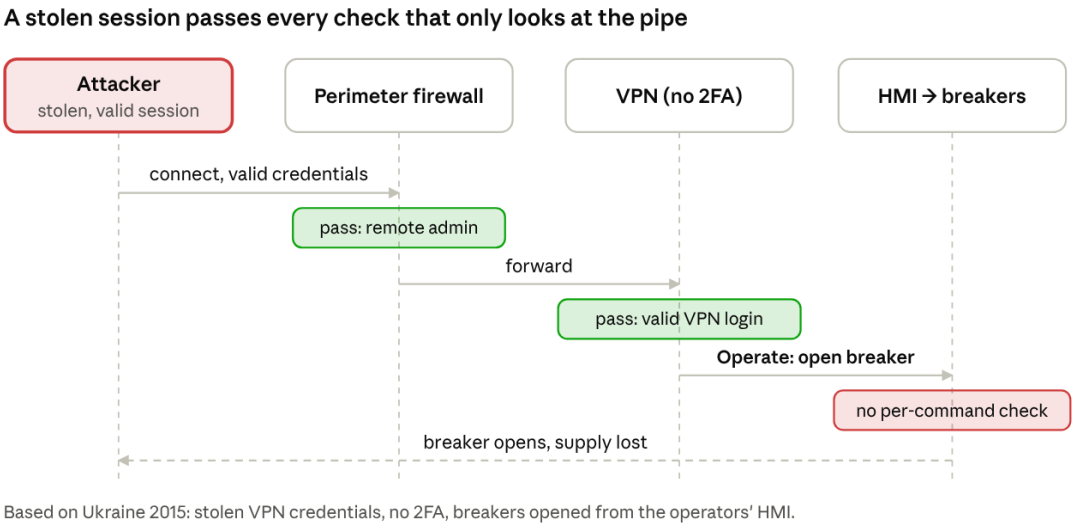


Figure 2. Sequence: a stolen session opens breakers (based on Ukraine, 2015)

The trend is moving toward destruction: Dragos reports destructive wiper operations against energy and infrastructure targets in 2025 (Dragos).

Threat model for this post: an attacker already inside the substation network, holding a valid session, sending commands the network will accept. Perimeter defenses do not cover this case. To see why, we first need to look at what a single command is made of.

Start from zero: four jobs in every command

Any command between a control center and a substation device answers four separate questions. Example: an operator opens a breaker.

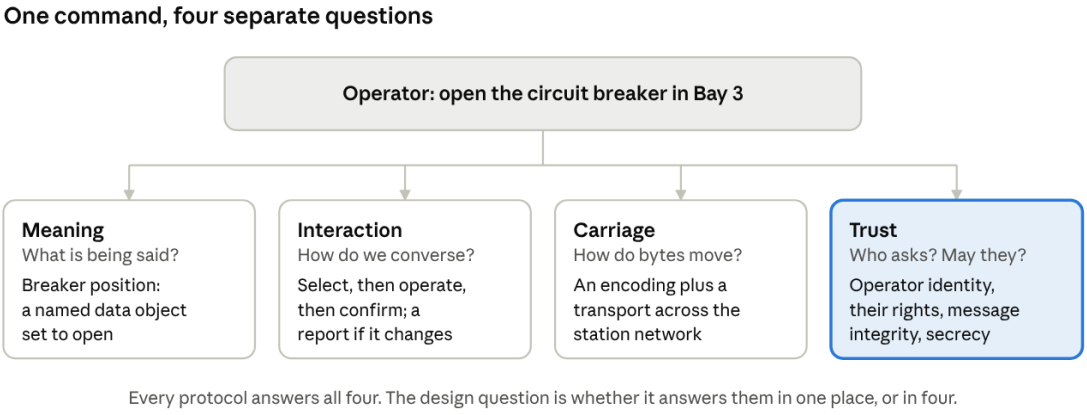


Figure 3. One command, four concerns

- **Meaning:** what is being said. In IEC 61850, the data model.
- **Interaction:** how the conversation runs: read, write, report, select-before-operate. In IEC 61850, ACSI (Abstract Communication Service Interface).

- **Carriage:** how the request becomes bytes and travels. Today, MMS over TCP.
- **Trust:** who is asking, whether they may, and whether the message was altered. Today, added from outside.

On paper, IEC 61850 already separates meaning and interaction from carriage: ACSI (IEC 61850-7-2) is abstract, and IEC 61850-8-1 maps it onto MMS. In products and tools that separation collapses, because they are built directly on MMS. And trust was never one of the separated concerns. That gap is where the first principle starts.

Principle 1: A concern we cannot name, we cannot secure

If trust has no place of its own, there is nowhere to secure it. We can only secure what we can point to.

Take MMS, the dominant client–server stack in utility operation. It combines meaning, interaction and carriage in one stack, so in practice trust is added outside it.

Security for MMS does exist. IEC 62351-4 (2018) adds authentication, integrity and encryption at the application layer, and IEC 62351-8 carries role-based access rights. Both work only when every IED and every client implement them, and most of the installed base does not.

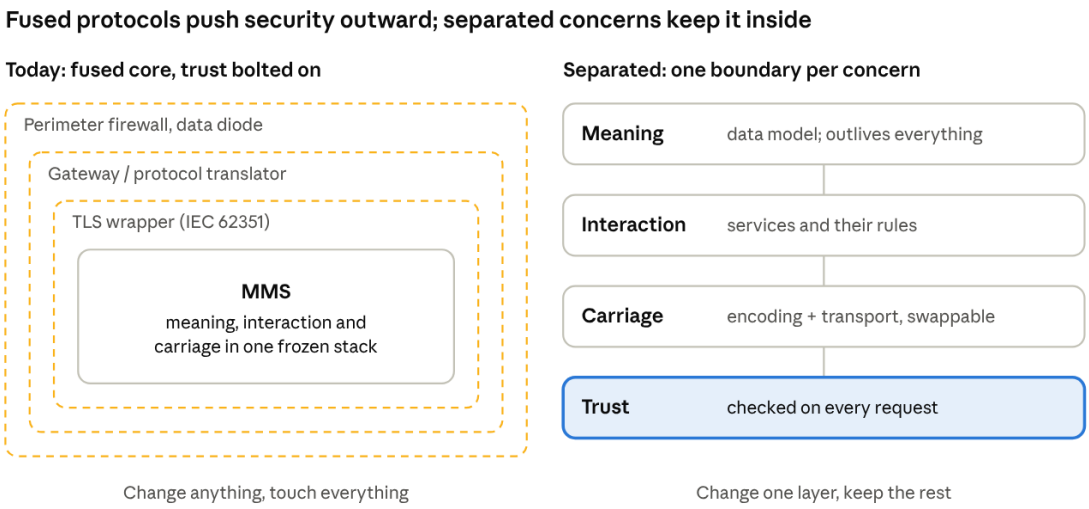


Figure 4. A fused stack vs separated concerns

Left: each ring is a separate product, configuration and failure point. An attacker past a ring reaches a core that does not check who is asking. Right: trust is its own layer, checked on every request.

Communication engineering already works this way. The OSI reference model (ISO/IEC 7498-1) gives each layer one job and a defined interface to its neighbors, so one layer can change without touching the others. Parnas (1972) gave the rule behind it: draw boundaries around the decisions most likely to change. Substation implementations did not carry that rule through. Naming the concerns tells us what to separate, but not where to draw the lines. For that, we look at how fast each concern changes.

Principle 2: Cut where the speeds differ

Each of the four concerns draws on a different field, and each field changes on its own timeline.

Where two parts change at different speeds, put an explicit boundary between them. Otherwise the fast part is held back by the slow one. Assets last decades. Data models change slowly and stay backward compatible (IEC 61850 Edition 1 around 2003, Edition 2 in 2010, Edition 2.1 in 2020). Transports turn over about once a decade, crypto every few years (TLS 1.0 and 1.1 were deprecated in 2021), and attack tools in months.

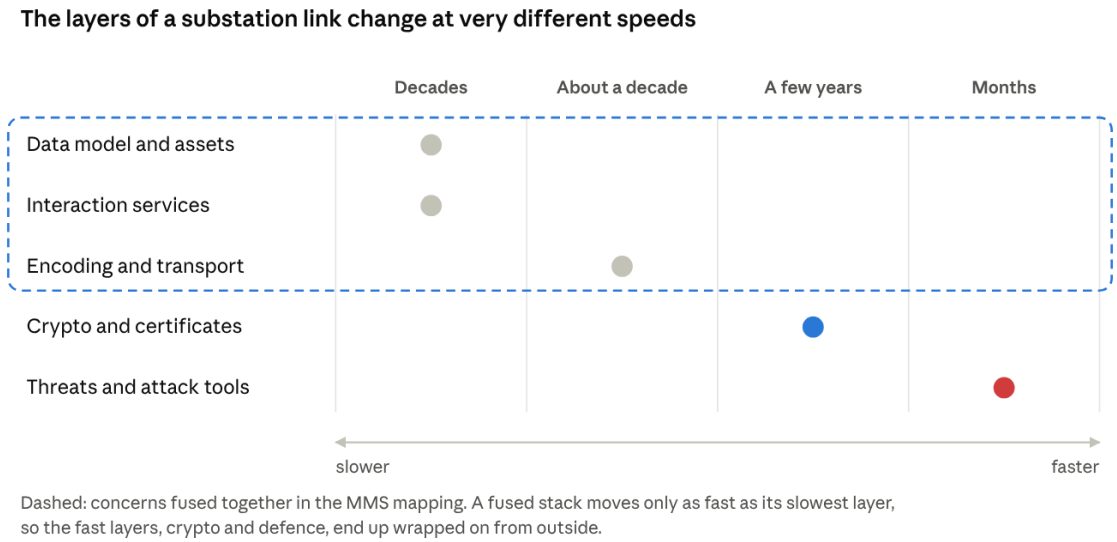


Figure 5. How fast each concern changes

TLS (IEC 62351-3) wraps MMS so cipher suites can change. But certificates and security support still live inside each IED's MMS stack, so they change only when the device changes, on a decades-long clock.

The most important boundary is between interaction (ACSI) and carriage (encoding and transport).

Principles 1 and 2 tell us what to separate and where to cut.

One question remains: where should trust be checked?

Principle 3: Enforce trust where intent is visible

A security check is only as precise as what we can see.

This sounds obvious, and it is. It is also the basis for observability: we can only log what we can see.

Table 1. What each guard can see, and the best decision it can make

Guard	What it can see	Decision it can make
Perimeter firewall	Addresses and ports	"This host may talk to that relay"
TLS wrapper	A certificate and an encrypted stream	"The pipe is authentic and private"
A check that reads each request	The caller, the service and the object	"Engineers may read Bay 3; only operators may operate its breaker"

A status read and a trip command use the same port and the same TLS session. Firewalls and TLS cannot tell them apart, so a stolen session can do both.

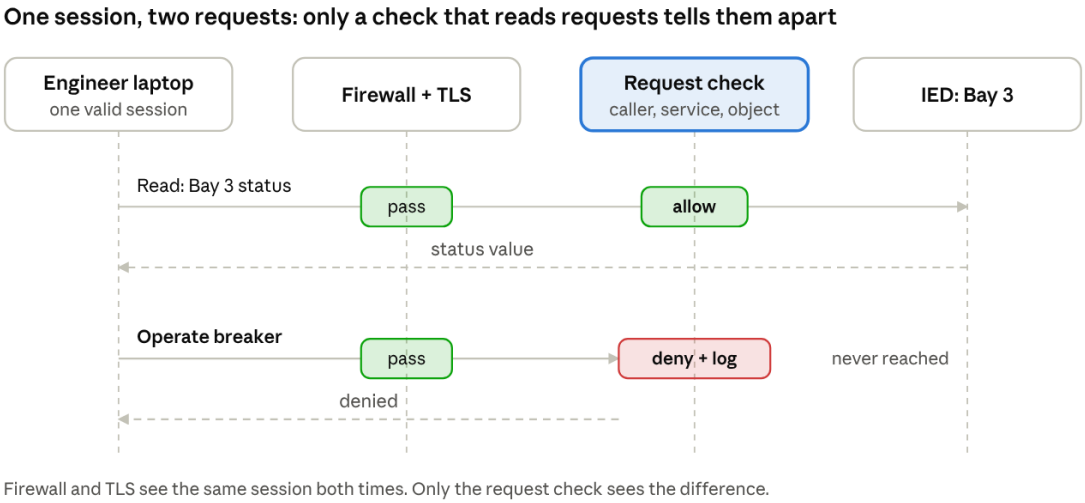


Figure 6. Sequence: one session, a read and an operate

Authorize each request, not each connection (zero trust, NIST SP 800-207). This is exactly how the 2015 attack worked: valid sessions, legitimate-looking commands. IEC 62351-8 already defines the roles; what is missing is a check that understands each service call and enforces them on every request.

With all three principles stated, we can now put the design together.

Deriving the ideal communication transport

Applied in order, the three principles narrow the options to one shape.

Where can the check sit? Not at the perimeter or in TLS: both see the pipe, not the request (Principle 3). Inside each IED, IEC 62351-4 and -8 can see every request, but only on devices that support them, and

the installed base mostly does not (Principle 2). What remains is a check placed in front of the devices that understands ACSI requests and does not depend on the transport. Step by step:

1. **Keep the data model.** The IEC 61850 data model does not change.
2. **Separate the services from the wire.** ACSI requests are understood on their own, whatever transport carries them.
3. **Check trust where requests are understood.** Every call is authenticated, authorized by role, service and object, and logged.
4. **Make the transport a plug-in.** MMS through an adapter for existing devices, a web-native carrier for new ones. Ciphers are updated in the carrier.
5. **Log in the grid's terms.** Logs record who did what to which object.

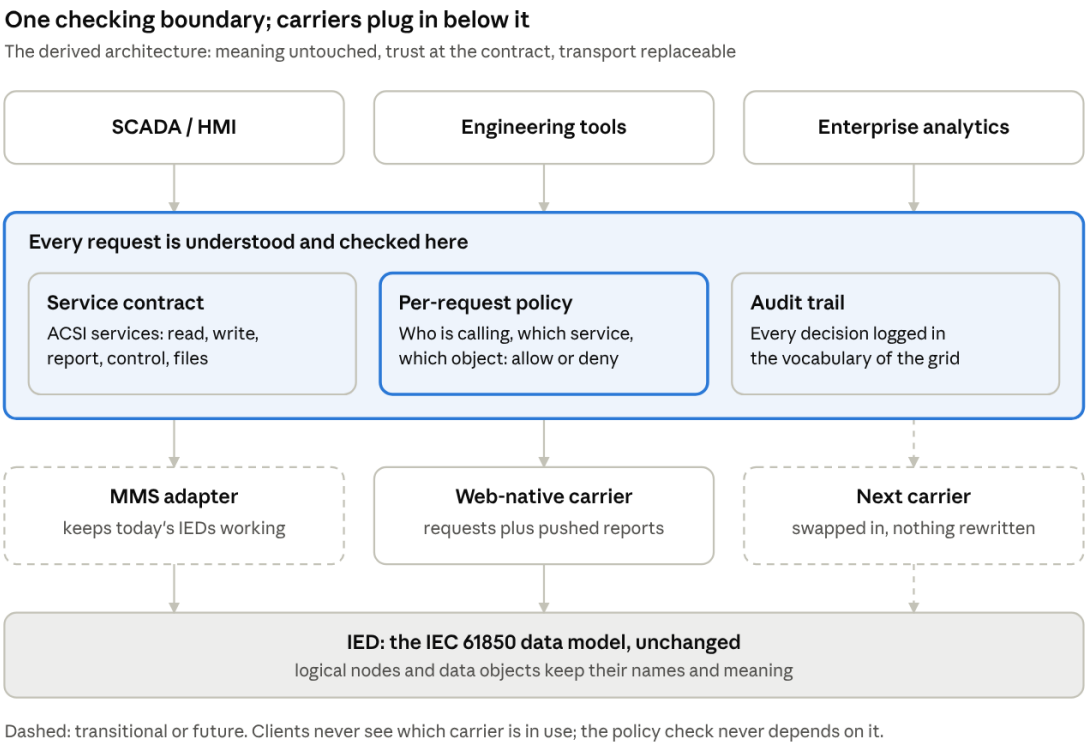


Figure 7. The derived design: one checking boundary, carriers below it

The design holds only under these conditions:

- Every request reaches the IED through the check. The network must block any other path, not just discourage it.
- The allow/deny decision depends on caller, service and object, never on the transport.
- The link between the check and the IED is protected too, or isolated so no one else can reach it.
- If the check fails, local protection keeps working; only remote requests stop.
- Control commands still meet their timing limits (IEC 61850-5), so the added check must be measured against them.

What defenders gain

What does this change for a defender? Each MMS shortcoming traces to a missing boundary, and each boundary derived above closes one.

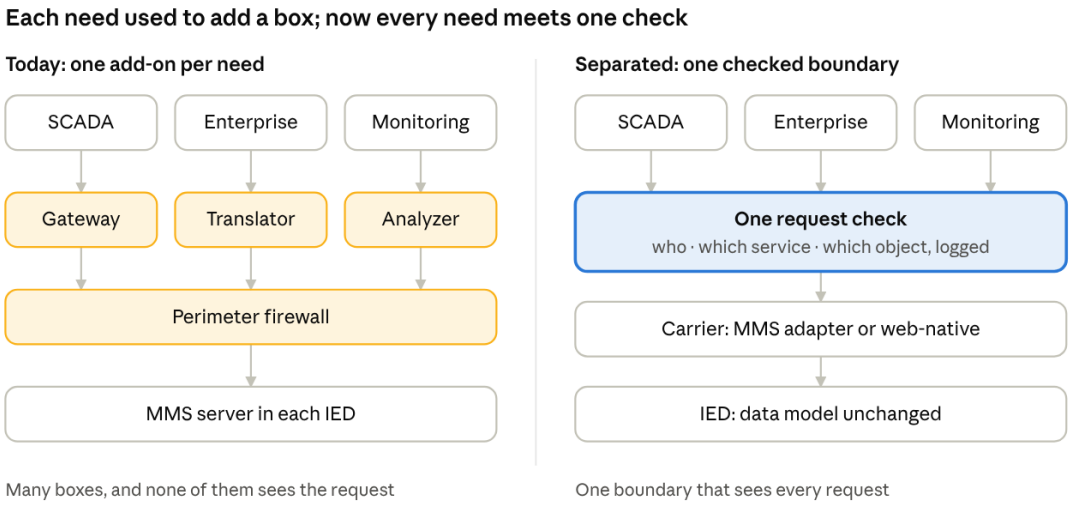


Figure 8. One add-on per need vs one checked boundary

Table 2. Each MMS shortcoming, the boundary it lacks, and what separation gives a defender

Shortcoming in MMS	Missing boundary	What separation gives a defender
Security is an overlay	Trust had no layer of its own	Every request is checked for who is calling, which service and which object. An attacker who gets past the firewall still cannot operate a breaker without the right role. The firewall becomes a second line of defense, not the only one.
Crypto upgrades are painful	Crypto support tied to each IED's MMS stack	Cipher suites and certificates are handled in the carrier. TLS already lets ciphers change; with the check separated, a crypto upgrade no longer waits for every IED's own security support.
Integration means translation	No contract above the wire	Enterprise tools call the ACSI service contract directly. One adapter per carrier replaces a chain of point-to-point protocol translators, and every request through it is checked.
Observability is limited	Logs speak bytes, not intent	Every log entry records who made the request, which service, which object, and whether it was allowed or denied, for example "operator X operated breaker Bay 3: allowed". A security analyst can read it without decoding MMS

Shortcoming in MMS	Missing boundary	What separation gives a defender
		packets.
Expertise concentrates	Knowledge locked in one encoding	Carriers use common web technologies (HTTP, TLS, Websockets) that many engineers already know how to secure and audit. Checking the link no longer depends on a few MMS specialists. The trade-off: web stacks are large, so they must be kept minimal and patched.
Innovation avoids the core	Core cannot change without breaking all	The data model and the ACSI services stay fixed. New transports, ciphers and tools are added at the carrier or client side, without changing the core.

Nothing else changes: the data model, the ACSI services and the engineering files stay as they are. That is why any existing link can be tested against the questions below today.

Key takeaways

Secure the conversation, not just the cable. Ask these five questions about any substation link:

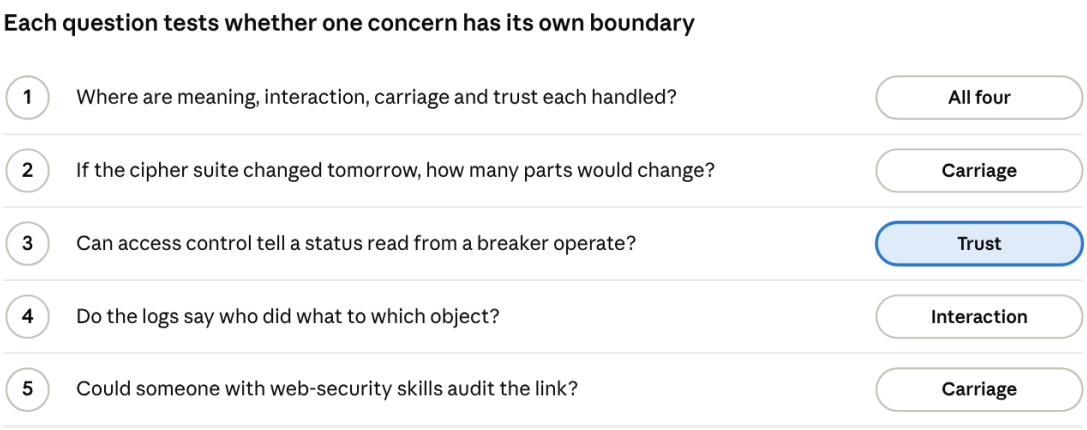


Figure 9. Which concern each question tests

1. Can you name, separately, where meaning, interaction, carriage and trust are handled?
2. If you had to replace the cipher suite tomorrow, how many components would change?
3. Can your access control tell a status read from a breaker operation from the same client?
4. Do your logs say *who did what to which object*, or only *which IP talked to which port*?
5. If MMS expertise left your team next year, could someone with web-security skills still audit the link?

If you cannot answer one of them, that concern has no boundary of its own.

Abbreviations

Term	Meaning
2FA	Two-factor authentication
ACSI	Abstract Communication Service Interface (IEC 61850-7-2)
HMI	Human–machine interface
ICS	Industrial control system
IED	Intelligent electronic device, such as a protection relay
MMS	Manufacturing Message Specification (ISO 9506)
OSI	Open Systems Interconnection
OT	Operational technology
SCADA	Supervisory control and data acquisition
TLS	Transport Layer Security
VPN	Virtual private network

References

1. PGCoE (2026). *Manufacturing Message Specification (MMS): A Protocol That Refuses to Retire.*
2. ISO/IEC 7498-1 (1994). *Information technology - Open Systems Interconnection - Basic Reference Model*.
3. Parnas, D. L. (1972). On the criteria to be used in decomposing systems into modules. *Communications of the ACM*, 15(12), 1053–1058.
4. IEC 61850-5: Communication requirements for functions and device models.
5. IEC 61850-7-2: Abstract communication service interface (ACSI).
6. IEC 61850-8-1: Specific communication service mapping (SCSM) to MMS.
7. IEC 62351-3: Profiles including TCP/IP (TLS).
8. IEC 62351-4:2018: Profiles including MMS and derivatives.
9. IEC 62351-8: Role-based access control for power systems management.
10. NIST SP 800-207 (2020). *Zero Trust Architecture*.
11. RFC 8996 (2021). *Deprecating TLS 1.0 and TLS 1.1*.
12. ESET (2017). Industroyer: Biggest threat to industrial control systems since Stuxnet.
13. Dark Reading (2016). Lessons from the Ukraine electric grid hack.
14. POWER (2016). How attackers caused the Ukraine blackouts.
15. CCDCOE. Power grid cyberattack in Ukraine (2015).
16. Dragos (2026). OT/ICS Cybersecurity Year in Review.