

IT/OT Convergence

The Demilitarized Zone

NETWORK SEPARATION AND GOVERNED EXCHANGE IN IT AND INDUSTRIAL ENVIRONMENTS

Martin van der Linden, 2026
The Governed Boundary

Abstract

A demilitarized zone (DMZ) provides an intermediate network in which exchanges between differently protected domains can be constrained, mediated and observed. Its industrial variant, the IDMZ, applies this principle between enterprise information technology (IT) and operational technology (OT) while preserving safe and reliable plant operation. This paper explains the historical metaphor, the two separation boundaries, policy-based traffic mediation, identity and access management, and protection of information in both directions. It argues that the boundary becomes effective through explicit ownership, risk-based design, controlled changes and verifiable evidence. The resulting architecture supports regulatory obligations and incident reporting, while retaining limits that must be addressed through wider security and operational governance.

Purpose and scope

The paper is intended for architects, industrial operators and governance decision makers who need a common understanding of the DMZ concept. It describes a pattern for controlled exchange rather than a product configuration or a claim that every organisation needs the same topology. “Governed boundary” means that each permitted exchange has an approved purpose, defined conditions, an accountable owner and evidence of its operation. The same principles apply to physical and virtual networks, provided that their separation is enforced. It concentrates on traffic mediation and on the checks applied to instructions at the boundary. The translation of business objects into operational commands, and of governance obligations between the IT and OT disciplines, is outside its scope.

The DMZ pattern

Historical origins and the computing metaphor

Military and diplomatic arrangements provide documented precedents for separating potentially hostile activities. Articles 42 and 43 of the Treaty of Versailles, signed on 28 June 1919, prohibited German fortifications and specified military activities on the left bank of the Rhine and within a defined area extending 50 kilometres east of the river. These provisions illustrate demilitarisation through geographical limits and agreed restrictions. They do not establish the first use of the English term. [2]

The Korean Armistice Agreement of 27 July 1953 explicitly established a demilitarized zone. Article I required both opposing sides to withdraw two kilometres from the military demarcation line, creating a buffer four kilometres across. The agreement identified northern and southern boundaries and arrangements for marking and supervising them. Its stated purpose was to prevent incidents that might restart hostilities. [3]

The terminology was in documented networking use by January 1997: RFC 2072 discusses a “demilitarized zone (DMZ)” alongside external and internal screening routers and a bastion host. This provides dated evidence of the computing usage, without assigning the invention to a particular person or year. The metaphor highlights separation and controlled interaction. In computing, the intervening area contains active services and enforcement mechanisms; its name does not imply that it is inherently safe or free of threats. [4]

In computer networking, the military metaphor describes a separate network segment between domains with different security requirements. A familiar example is the DMZ between the public internet and a private corporate network. Public services can be made accessible there while access to internal systems remains restricted. Its protection depends on configuration and operation. [5]

Two boundaries define the intervening zone

In the architecture developed here, a DMZ has two enforced separation boundaries: external network to DMZ, and DMZ to protected internal network. A dual-firewall design places the intermediate network between an outer and an inner firewall. Without separation on both sides, the intended governed buffer is not established. Naming a subnet a DMZ does not create that separation. Terminology varies across implementations; this paper makes the two-boundary requirement explicit for its controlled-exchange model. [6]

The two boundaries can be implemented by separate firewall appliances or by distinct interfaces and policies on a firewall with at least three interfaces. Separate appliances can reduce some shared failure dependencies, although shared administration or configuration can still create common weaknesses. A single appliance remains a possible implementation with its own failure implications. The requirement concerns enforced separation, not an assertion that every security zone universally requires two physical firewalls. [5, 6]

Underpinning network zones and conduits

The zone-and-conduit model gives this separation a systematic basis. In the ISA/IEC 62443 approach, a zone groups assets with common security requirements; a conduit groups communication channels with common security requirements connecting zones. Risk assessment determines the required protection for each. The model originated for industrial automation and control security, but also provides a useful way to explain the relationships around a DMZ. [11, 12]

The DMZ is a zone, or a set of separately protected zones, hosting exchange services. Conduits describe the controlled communications between it and neighbouring zones. A conduit is not simply a cable, and a zone is not necessarily identical to a subnet. Network segmentation provides the infrastructure; boundary controls give effect to the security requirements. For an internet-facing DMZ (internet-to-DMZ access and DMZ-to-enterprise network access) represent distinct communication relationships with distinct permissions. [11, 12]

Risk determines the boundary

Design begins with the system under consideration: its assets, information, dependencies and operational consequences. Relevant scenarios include compromised enterprise accounts, malicious or mistaken maintenance activity, unauthorised commands, data theft and service outages. Asset owners should determine the required protection and acceptable residual risk for each zone and conduit. IEC 62443-3-2 connects this assessment to target security levels and documented requirements; a diagram or firewall installation alone does not establish that those requirements are met. [1]

Cyber security

A DMZ supports defence in depth by separating exposed services from internal resources and limiting the opportunities for a compromised service to reach other systems. Web, mail and externally accessible DNS services are typical candidates. Protection also requires hardened hosts, restricted administrative access, monitoring and maintenance; placing a vulnerable server in a DMZ does not make it secure. [7]

Policy-based traffic mediation

The boundary implements an explicit communication policy: which source may contact which destination, using which protocol or service, and in which direction. A deny-by-default approach permits only justified flows. Firewalls enforce network rules, while application proxies can terminate a connection and make a separate, controlled request on its behalf. Stateful filtering admits return traffic for permitted sessions without granting unrestricted reverse access. Rules should be documented, tested and reviewed as requirements change. [5]

For example, permitting an internet client to reach a web service in the DMZ need not permit that client to reach an internal database. Any required connection from the web service to the database is a separate policy decision, restricted to the necessary endpoint and service. That connection should retrieve only the data needed to answer each request. Because it is opened from the DMZ inward, it is an accepted exception to initiating transfers from the more protected side. Service availability need not extend general network access. [5]

Bidirectional protection and confidentiality

Traffic mediation must address both directions. Inbound controls protect private networks against unauthorised access; outbound controls protect corporate confidentiality and the privacy of personal information by restricting what can leave, where it may go and who may release it. An authorised connection is not blanket permission to disclose its contents. The policy should therefore govern data flows as well as connectivity, including responses sent over connections initiated from outside. [5, 8]

For example, an approved reporting service may release selected production totals while withholding recipes, customer details or employee records. Destination restrictions, application authorisation, data filtering and monitoring support this distinction. Encryption protects information in transit, but does not make disclosure to an unauthorised recipient acceptable. Bidirectional protection means controlling permitted exchanges in both directions; it does not require identical permissions, or bidirectional communication where a one-way transfer is sufficient. [8]

Identity and access management

Identity and Access Management (IAM) governs which parties may use protected exchange services and what each of them may do. Individual user identities and managed service identities establish who or what may act; audit records link actions to those identities. Anonymous access and untraceable shared accounts should therefore be excluded from protected exchanges. IAM supports accountability, but its mere presence does not identify every packet or guarantee the real-world identity of an actor using stolen credentials. [9]

For protected services, the identity lifecycle includes approval, provisioning, privilege changes, credential renewal and prompt revocation. Human remote access should use multi-factor authentication; automated services need managed credentials and narrowly scoped permissions. IAM decisions must be enforced by the gateway or application handling the exchange. Public services may deliberately permit anonymous access to releasable information, so the prohibition on anonymity applies to protected exchanges rather than every service in every DMZ. [9, 10]

The Industrial Demilitarized Zone

A variant is the Industrial Demilitarized Zone (IDMZ), which mediates exchanges between enterprise Information Technology (IT) and industrial Operational Technology (OT), including industrial control systems (ICS). Its purpose is to support necessary business and operational information exchange while maintaining separation. Corporate ownership of both networks does not make enterprise access appropriate for every industrial asset. [13]

Separation and mediated exchange

Typical designs control both the enterprise-facing and OT-facing boundaries. Application exchanges terminate on an IDMZ service and continue through separately authorised connections. No industrial application or engineering connection traverses the IDMZ from end to end. A maintenance session may span both boundaries, but it is composed of separately authorised connections that terminate in the IDMZ. Data transfers should be initiated from the more protected side wherever the exchange permits it. Interactive access initiated from the enterprise side should end at an IDMZ broker. Common services include replicated

production data, application proxies, managed file transfer and update staging. Enterprise users can consult a replicated historian rather than querying control systems directly. Essential control functions and their supporting services remain within OT so that plant operation does not depend on IDMZ connectivity. Buffering can preserve data during a temporary disconnection. [13]

Some reference designs allow selected network services to communicate directly between the enterprise and industrial zones, where the organisation accepts the risk. In the model developed here, such an exception is never made for industrial application or engineering protocols. A named network service may cross only as an individually assessed and owned conduit. It may do so only when its risk is assessed, documented and accepted at the level of accountability assigned to it. [13]

These controls apply in both directions across the IDMZ. Access to production information should be attributable to an authorised user or service, and the release policy should protect confidential industrial information as well as the receiving OT systems. The two controlled boundaries and IAM enforcement together make the exchange accountable; neither network location nor corporate affiliation alone authorises access. [6, 8]


Many industrial architectures describe this placement with the Purdue reference model. Enterprise systems occupy Levels 4 and 5, site operations Level 3, and control and field equipment Levels 0 to 2. The IDMZ is then conventionally drawn as Level 3.5, between site operations and the enterprise. The model is a useful orientation, but it is a functional hierarchy rather than a security architecture. Cloud services, vendor platforms and direct device connectivity do not always fit its levels. Zones and conduits therefore remain the basis for protection decisions. [13]

Business intent entering operations

Some exchanges carry instructions rather than information. Production orders, set-points and replanning decisions move from enterprise systems towards equipment that acts on them. Their consequences are physical, so they require stricter conditions than data moving outward. Business approval alone is not sufficient for any of them to be executed. Figure 1 shows the resulting sequence.

The IDMZ checks what can be checked at the boundary: the identity of the originator, its permission to send the instruction, and the conformance of the instruction to the agreed message requirements. Where the operational side can confirm in advance that an instruction is currently executable, the IDMZ may obtain that confirmation before forwarding. That answer is a precondition, not a guarantee, because equipment state can change between the check and execution. When any boundary check is uncertain, the instruction is withheld rather than forwarded. Where the exchange permits it, the receiving OT application collects forwarded instructions from the IDMZ rather than accepting connections from it.

The final decision to execute remains within OT. The receiving OT application independently authorises the instruction and validates it against current conditions and safety constraints before acting. A partial failure of the boundary should result in non-actuation, not in unvalidated commands reaching the process. Withholding or rejecting an instruction must itself be safe, which the process safety assessment should confirm. Every forwarded, withheld or rejected instruction should be recorded with its originator, its authorisation and the checks on which the decision rested.



Boundary checks are a precondition, not a guarantee. The final decision remains within OT. The arrowhead indicates the direct ion of the instruction or status. The dot marks the side that opens the connection. Push: the opening side sends. Pull: the opening side retrieves.

Figure 1 – An instruction entering operations: boundary checks and the final execution decision

Safety and predictable operation

Industrial security decisions must account for human safety, environmental consequences, process integrity, availability and timing requirements. Confidentiality remains relevant, but security mechanisms must not introduce unacceptable operational disruption. Real-time control paths should remain outside the IDMZ. Redundancy and tested recovery reduce the effect of boundary failures; failure responses must follow the process safety assessment rather than assuming that either blocking all traffic or allowing all traffic is always safe. [10]

Least privilege and controlled remote access

Only identified users, devices and applications should reach authorised services. Remote maintenance should pass through a managed access point, such as an IDMZ jump host, with individual accounts, strong authentication, limited privileges and session logging. Access should be enabled for the required task and duration, with a means to revoke it. A remote connection must not create an alternative path around the enterprise–OT boundary. IT and OT teams need clear responsibilities for approving and administering these connections. [14]

Identity domains across the boundaries

The structure of identity domains is itself a design decision. Whatever the structure, identities should be attributable, trust between domains restricted, authorisation decided locally and revocation assured. The less protected domain should never administer the more protected one. Where identities cross, the IDMZ may accept enterprise identities for named IDMZ services only, and OT should not accept enterprise identities directly. A directory shared by the IDMZ and OT can create unwanted dependencies. Depending on directory placement, privileges and replication arrangements, OT authentication could come to rely on IDMZ availability. An IDMZ compromise could also extend into OT. The design must prevent both risks. Domain controllers located within OT, for example, can keep authentication available when connectivity is lost, provided their physical security can be assured. Essential OT authentication must in any case remain available when the IDMZ is lost. [9, 10, 24]

Vendor engineers can be admitted through federation with managed attributes or through individually assigned, locally managed accounts. Federation is one implementation choice among these. Shared accounts are excluded in both cases. [9]

Controlled changes and proportionate inspection

Long equipment lifecycles, legacy protocols and restricted maintenance windows require careful change control. Patches, firewall changes and inspection features should be tested for compatibility and operational impact before deployment. Passive monitoring is often appropriate where active probing could disturb equipment. Where the required exchange is strictly outward, a unidirectional gateway can strengthen separation; it cannot support an ordinary bidirectional session. Safety systems may need additional isolation, and an IDMZ does not replace segmentation within OT. [10]

Alternative paths and bypass prevention

The boundary must account for every connection capable of crossing its intended separation. Inventory vendor tunnels, cloud connectors, wireless links, maintenance laptops and devices attached to more than one network. Route necessary exchanges through approved controls or isolate them through an explicitly assessed alternative. Removable media cannot be forced through a firewall, so it requires a separate transfer and inspection procedure. A bypass can invalidate the protection assumed in the zone-and-conduit design. [10]

Some connectivity patterns challenge the IDMZ model directly. Cloud connectors installed in OT often open outbound sessions to a provider, creating a path that bypasses enterprise controls. Vendor edge platforms combine equipment functions with remote connectivity under the vendor's administration. Zero-trust network access products are sometimes presented as a replacement for the IDMZ. Each of these patterns should be treated as a conduit and assessed against the same requirements. Identity-based access can strengthen the boundary, but it does not remove the need to terminate sessions and govern what crosses. A cloud or vendor path is acceptable only when it is inventoried, owned, monitored and routed through controls that the asset owner can verify. [10, 14]

Failures, isolation and restoration

Define responses to loss or compromise of the firewalls, IAM, exchange services and log collection. Separate decisions about admitting new sessions from decisions about retaining existing ones. Deny new privileged remote access when authentication cannot be established; use a separately authorised, logged emergency-access procedure where necessary. Essential OT control should continue locally or move to a process-defined safe state. Specify how long data can be buffered, when services must stop, and who may isolate the boundary. [10]

Restoration should follow an approved sequence: establish trustworthy configurations and identities, recover monitoring, validate the exchange services, and reconnect the approved conduits. Failover must preserve filtering and authorisation rather than silently opening a route. Test loss of shared dependencies, including power, name resolution and time services, as well as individual appliance failures. These are design and acceptance decisions to resolve before an incident, with operational staff involved.

Worked example of governed exchange

This illustrative design uses two firewall boundaries and separate IDMZ services. Each arrow is a separately authorised connection, and a dot marks the side that initiates it. The outer zones show the functions that govern each exchange, not only its endpoints. The layout is a logical example.




Figure 2 – Two controlled exchanges through an industrial DMZ, with the functions that govern them

A production report moves outward

An OT export service publishes approved production totals to the IDMZ replica through Boundary B. Its managed identity permits only that publication. A reporting client separately authenticates through Boundary A to retrieve the report; it cannot query the OT historian directly. The export service excludes recipes and personal records before publication, so withheld data never enters the IDMZ. Recording the dataset version, publishing identity, release decision and retrieval event supports later investigation of what was disclosed.

A maintenance session proceeds inward

An engineer receives approval for a named OT endpoint and maintenance period, then authenticates with multi-factor authentication to the IDMZ broker through Boundary A. The broker then opens a separate, scoped connection through Boundary B. There the engineer authenticates again with an OT-local identity, and the OT endpoint authorises the session locally. Log approval, authentication, target access and relevant actions; revoke temporary rights when the task ends. Files use a separately approved inspection and transfer service.

If IAM fails, new maintenance access is withheld while local control continues. The designated operator decides whether an emergency procedure is necessary. Firewall, IAM and application events from all three zones are forwarded outward through a log relay in the IDMZ to security monitoring. Together they provide correlated evidence for review and reporting.

Governance, compliance and evidence

Ownership throughout the lifecycle

The process owner approves operational need and the consequences of interruption. The information owner authorises what may be released. IT and OT security owners define boundary controls and the division of

administration; the designated risk owner accepts residual risk and exceptions. These roles may be combined in a small organisation, but their decisions should remain explicit and recorded. Implementation and approval should be separated where the risk warrants independent review.

Each exchange should have a record of its purpose, endpoints, permitted actions, data classification, identity requirements, owner and review or expiry date. Commissioning verifies the intended permissions; change control covers both firewall rules and application behaviour. Reviews remove obsolete access, renew justified exceptions and confirm that service identities still have owners. Retirement revokes credentials, closes rules and disposes of retained data under the agreed policy. This proposed lifecycle makes responsibility persistent rather than a one-time design decision.

Monitoring and continued governance

The IDMZ itself requires protection because compromised exchange services can become a route into control systems. Asset owners should restrict its services, monitor boundary activity and investigate unexpected communication. Agreed ownership of rules, service maintenance and incident response helps keep the boundary effective as production requirements change. [15]

Adherence to EU, national and sector requirements

The DMZ and IDMZ also support adherence to applicable rules and regulations. Their communication policies should translate relevant obligations into controlled access, authorised data release, monitoring and retained evidence. The organisation must first establish which requirements apply to its activities, locations, information and sector. At EU level, examples include NIS2 cybersecurity risk-management and incident-reporting obligations as implemented nationally, and GDPR obligations where personal data are processed. The CER Directive adds resilience obligations for designated critical entities. The Cyber Resilience Act places its obligations mainly on manufacturers, importers and distributors of products with digital elements, including equipment deployed at the boundary. Its reporting obligations have applied since 11 September 2026, and its main obligations apply from 11 December 2027. An operator is affected chiefly through the products it procures. An operator's CRA obligations depend on whether it also acts as a manufacturer, importer or distributor, including through substantial modification and subsequent supply of a product. A DMZ supports these duties; its deployment alone does not establish compliance. [16, 17, 20, 21]

At national level, the design and operating procedures must reflect the applicable implementing legislation, additional domestic requirements and competent authorities. This includes identifying reporting recipients, thresholds, deadlines and required evidence for each jurisdiction. Sector requirements can further specify controls and reporting arrangements. For example, DORA establishes ICT risk-management and incident-reporting duties for financial entities within its scope; it is not a general requirement for all industrial operators. In aviation, Part-IS requires organisations within its scope to manage information security risks that could affect aviation safety. Operators within the scope of both NIS2 and Part-IS should align their controls rather than duplicate them. The competent authority verifies whether NIS measures are equivalent in effect to Part-IS requirements. NIS compliance does not by itself establish Part-IS compliance, and the Part-IS external reporting requirement (IS.OR.230) applies in any case. Applicable sector-specific EU rules may replace corresponding NIS2 provisions where the legal conditions are met. [16, 18, 22, 23]

A practical approach is to maintain a traceable mapping from each applicable obligation to the relevant boundary rule or service control, its owner, its verification method and its evidence record. Approved exceptions and changes should be recorded and reviewed. Technical standards such as ISA/IEC 62443 can support implementation and assurance, but should be distinguished from legislation and from standards made binding through a contract or regulatory requirement. For ISA/IEC 62443, the assessed requirements of IEC 62443-3-2 and IEC 62443-3-3 are best evidenced per zone and conduit. Each zone and conduit is assigned a target security level from the risk assessment. The implemented controls are then verified against the corresponding system requirements of IEC 62443-3-3. The gap between target and achieved security

levels is recorded, owned and either closed or accepted. Acceptance of residual risk does not, by itself, demonstrate achievement of the target security level or conformance with unmet requirements. A claim of conformance should therefore name the parts and requirements assessed. [1, 11]

Facilitating incident reporting

The controlled exchange points help detect and reconstruct incidents by correlating firewall, proxy, IAM and application events. These records support an account of when an event occurred and was detected, which identities and systems were involved, what information crossed the boundary, and which containment actions followed. They should feed an incident-handling process with assigned owners, escalation routes and evidence preservation. Boundary logs need to be combined with endpoint and process evidence; they do not provide a complete picture by themselves. ENISA's implementation guidance formally applies only to specified digital-infrastructure and ICT service providers. Its logging and monitoring practices are nevertheless a useful reference for other operators. [19]

Reporting procedures should distinguish a security alert from a legally reportable incident and apply the relevant significance or risk threshold. NIS2 and GDPR have different triggers, recipients and reporting duties, so one incident may require more than one notification. Under NIS2, as transposed nationally, a significant incident requires an early warning without undue delay and in any event within 24 hours of becoming aware of it. An incident notification follows without undue delay and in any event within 72 hours of becoming aware of the significant incident. That period runs from awareness, not from the early warning. A final report is due within one month of that notification. If the incident is still ongoing at that point, a progress report is due instead, followed by a final report within one month of handling the incident. Under GDPR, a personal data breach is notified to the supervisory authority without undue delay and, where feasible, within 72 hours, unless it is unlikely to result in a risk to individuals. Evidence gathering should support timely initial reports and subsequent updates rather than delay notification until the investigation is complete. Human responsibility remains essential for assessing impact and deciding which obligations are triggered. [16, 17]

Audit trails and evidence integrity

An audit trail should connect an event to its timestamp, authenticated user or service, source and destination, requested action, applicable policy and outcome. Retaining rule versions, approvals and administrative changes allows reviewers to establish which controls operated at the time. Synchronised clocks, protected central log collection, restricted administrative access and integrity checks help make the evidence reliable. Retention periods and retrieval procedures should be defined, and logging failures should be detected. These measures support incident investigation and internal or external audits. [19]

Audit records may themselves contain personal or confidential information. Their collection, use, access and retention must therefore be limited to justified purposes and applicable requirements. Unnecessary payload capture should be avoided, while relevant evidence is preserved for investigations and legal obligations. Accountability requires usable, protected records, not unrestricted surveillance or indefinite retention. [17]

For evidence used in a formal investigation, record who collected it, how it was exported, integrity checks and subsequent transfers or access. Protect evidential copies from alteration and separate their custody from routine administration. A sequence of authenticated events supports attribution, but should be interpreted alongside endpoint and operational evidence rather than treated as conclusive proof of personal responsibility.

Verification and residual limitations

Acceptance should test approved exchanges and attempts that must fail: unapproved direct IT-to-OT access, including all direct industrial application and engineering connections; unauthorised destinations; revoked identities; and prohibited data release. Periodic checks should cover rule drift, stale permissions, log

completeness and time correlation. Exercise isolation, recovery and reporting procedures without jeopardising production. Record the results, unresolved findings and the person accepting any remaining exposure. [19]

Permitted traffic can carry malicious instructions, and an authorised user can make a harmful change. Encryption can conceal content from network inspection unless an approved termination point provides visibility. Compromise of a proxy, management service or shared dependency can undermine several controls together. The IDMZ therefore complements secure applications, endpoint protection, physical security and segmentation within OT. Assurance comes from tested controls and understood limits, not from topology alone. [5, 10]

Principles

The following principles define the DMZ/IDMZ pattern developed in this paper.

1. Separation is enforced at two boundaries. Naming a subnet a DMZ does not create one.
2. No industrial application or engineering connection traverses the IDMZ from end to end. Every such exchange terminates on an IDMZ service and continues through a separately authorised connection. A network-service exception is permitted only as an owned conduit whose risk is assessed, documented and accepted.
3. Data transfers are initiated from the more protected side wherever the exchange permits it. Interactive access initiated from the less protected side ends at an IDMZ broker.
4. Every exchange is denied by default. It is permitted only with an approved purpose, defined conditions, an accountable owner and a review or expiry date.
5. Both directions are governed. Inbound controls protect the receiving systems. Outbound controls protect the information released.
6. Protected exchanges admit no anonymous access. Every action is attributable to an individual or a managed service identity.
7. Essential control never depends on the IDMZ. Real-time control paths remain within OT.
8. Instructions entering operations are withheld at the boundary when a check is uncertain. The final decision to execute them remains within OT.
9. Failure never silently opens a route. New privileged access is withheld when authentication cannot be established.
10. Every path across the boundary is known. An unassessed alternative path is a defect, not an exception.
11. Effectiveness is demonstrated by tests and retained evidence. Topology alone establishes nothing.

Conclusion

A DMZ becomes a governed boundary through enforced separation, explicit permissions, accountable identities and observable exchanges. Its two boundaries define the intermediate area; mediation determines what may cross. Industrial operation additionally requires safe local control, isolation and restoration.

Risk assessment, ownership and lifecycle management sustain the concept. Outgoing information needs protection alongside defence against incoming threats. Applicable regulations shape controls and evidence; protected audit trails support investigation and reporting. Effectiveness must be tested, and residual risks must remain visible.

References

Technical guidance and standards inform the architectural discussion; legal applicability depends on the organisation and jurisdiction. Illustrative exchanges and proposed governance procedures are the author's synthesis.

- [1] IEC (2020). IEC 62443-3-2, Security risk assessment for system design. Scope and publication information. [Source](#)
- [2] Treaty of Versailles (1919), Articles 42–44, reproduced in the Office of the Historian diplomatic records. [Source](#)
- [3] Korean Armistice Agreement (27 July 1953), Article I, paragraphs 1–4. Transcript held by the US National Archives. [Source](#)
- [4] Berkowitz, H. (1997). Router Renumbering Guide. RFC 2072, section 9.2, Special Firewall Considerations. [Source](#)
- [5] NIST (2009). Guidelines on Firewalls and Firewall Policy. SP 800-41 Rev. 1, sections 2–4. [Source](#)
- [6] NIST (n.d.). CSRC Glossary, demilitarized zone (DMZ). Definitions covering intermediate networks and firewall interfaces. [Source](#)
- [7] NIST (2008). Guide to General Server Security. SP 800-123, sections 2 and 4; server placement, hardening and administration. [Source](#)
- [8] ISO/IEC (2022). ISO/IEC 27001, Information security management systems – Requirements. Annex A controls 5.14, 8.12 and 8.20–8.22; information transfer, data leakage prevention and network segregation. [Source](#)
- [9] NIST (2018). Identity and Access Management for Electric Utilities. SP 1800-2, Volume B; identity lifecycle, authentication and access management. [Source](#)
- [10] NIST (2023). Guide to Operational Technology (OT) Security. SP 800-82 Rev. 3, sections 2.4, 5.2–5.4, 6.2.11 and 6.3.2. [Source](#)
- [11] IEC (2013). IEC 62443-3-3, System security requirements and security levels. SR 5.1 network segmentation and SR 5.2 zone boundary protection. [Source](#)
- [12] ISA Global Cybersecurity Alliance (2025). Quick Start Guide: An Overview of ISA/IEC 62443 Standards; zone and conduit definitions. [Source](#)
- [13] Cisco and Rockwell Automation (2022). Securely Traversing IACS Data across the IDMZ Using Cisco Firepower Threat Defense. Chapter 2, System Design Considerations. [Source](#)
- [14] DHS (2010). Configuring and Managing Remote Access for Industrial Control Systems. Recommended practice, hosted by CISA. [Source](#)
- [15] DHS (2016). Improving Industrial Control System Cybersecurity with Defense-in-Depth Strategies. Recommended practice, hosted by CISA. [Source](#)
- [16] European Union (2022). Directive (EU) 2022/2555 (NIS2), Articles 4, 21, 23 and 41; sector-specific acts, risk management, incident reporting and transposition. [Source](#)
- [17] European Union (2016). Regulation (EU) 2016/679 (GDPR), Articles 5, 6, 32, 33 and 34; processing principles, security and personal data breach obligations. [Source](#)
- [18] European Union (2022). Regulation (EU) 2022/2554 (DORA), Articles 1, 2, 6 and 17–20; financial-sector scope, ICT risk management and incident reporting. [Source](#)
- [19] ENISA (2025). NIS2 Technical Implementation Guidance, section 3; incident handling, monitoring and logging. Guidance scope covers specified digital infrastructure, ICT service management and digital providers. [Source](#)
- [20] European Union (2022). Directive (EU) 2022/2557 (CER Directive); resilience of critical entities. [Source](#)
- [21] European Union (2024). Regulation (EU) 2024/2847 (Cyber Resilience Act); horizontal cybersecurity requirements for products with digital elements. [Source](#)
- [22] European Union (2022, 2023). Delegated Regulation (EU) 2022/1645 and Implementing Regulation (EU) 2023/203 (Part-IS); information security management for aviation organisations. [Source](#)
- [23] EASA (n.d.). Frequently asked questions, Information Security (Part-IS), applicability; equivalence of NIS and Part-IS requirements. [Source](#)
- [24] Microsoft (2025). Planning Regional Domain Controller Placement; local authentication during WAN outages and physical security of domain controllers. [Source](#)