

Network infrastructure projects are often treated as a cabling, switching or connectivity exercise, with cybersecurity brought in later to review firewalls and endpoint detection platforms. That sequence creates avoidable risk. Physical infrastructure is part of cybersecurity: cables, patch panels, switches, ports and equipment rooms all affect whether security controls work as intended.

Aligning infrastructure work with Australian Cyber Security Centre (ACSC) guidance means turning security objectives into practical design and operational decisions. The Information Security Manual (ISM) provides a framework for managing information security risks, while other ACSC guidance, including the Essential Eight, addresses important [patch panels](#) mitigation strategies. Applying that guidance depends on the organisation's risk, systems and operating environment. It also depends on foundations that are easy to overlook: accurate diagrams, capable switches, controlled physical access and a network that can be understood when something fails at 2am.


This article sets out a practical way to connect infrastructure projects with ACSC guidance, with particular attention to network segmentation, physical-layer reliability and ISM controls.

Start with the security outcome, not the equipment list

A project brief that says "replace the core switches" describes an activity, not the security outcome. Before selecting hardware or arranging a cutover, establish what the organisation needs the network to protect and support.

For example, a project might need to limit communication between corporate users and operational technology (OT), reduce the impact of a compromised workstation, or improve visibility of traffic between sites. Those objectives can then be mapped to relevant risks and ISM controls. The mapping should be documented and checked against the current ISM and applicable ACSC guidance; requirements and recommended practices may change over time.

Ask questions such as:



- Which systems and information are in scope, and who owns them?
- What communications are genuinely required between users, servers, cloud services and OT?
- Where are the trust boundaries, and which devices enforce them?
- How will administrators identify and investigate unusual traffic?

- What is the safe operating state if a switch, link or security device fails?

These questions keep the design connected to risk rather than turning it into a shopping list. They also help surface dependencies early. A firewall rule cannot provide meaningful separation if traffic can bypass the firewall through an unmanaged switch or an incorrectly configured trunk.

Physical infrastructure is part of the security boundary

Security discussions often focus on logical controls: firewall policies, identity, encryption and endpoint detection platforms. Those controls matter, but they rely on a physical environment that behaves predictably. A person with unauthorised access to an exposed network port, an unlocked communications cabinet or an undocumented patch can undermine assumptions about who and what is connected.

Physical security does not mean every network cabinet needs the same protection. It means assessing access according to the systems and data it supports. Consider who can enter equipment rooms, who can access floor cabinets, how visitor access is supervised and whether contractors can connect devices. Where the risk warrants it, use appropriate access controls, monitoring and procedures for handling keys and credentials.

Physical documentation is security documentation. Label patch panels, switch ports and fibre links so authorised staff can trace a connection without guesswork. Keep diagrams and port schedules current, protect them from inappropriate access and record changes. I keep a label maker in my laptop bag because an accurate label today can save an incident responder a long night later.

Unlabelled patch panels and mystery cables are not merely untidy. They can obscure rogue connections, make isolation slower and increase the chance that a technician disconnects the wrong service. Treat unknown cabling as something to investigate, not as harmless background noise.

Physical-layer faults can create outages and security blind spots

A damaged fibre, loose patch lead or failing power supply can look like a routine availability problem. Yet physical faults can also interfere with security monitoring and enforcement. A link that repeatedly flaps may disrupt logs or monitoring feeds. A poorly understood redundant path may carry traffic around a control that the design assumes is in-line. An unplanned temporary switch may reconnect networks that were meant to be separate.

These failures create two related risks: services may be unavailable, and the organisation may have an incomplete picture of what is happening on the network. During troubleshooting, emergency workarounds can make that second risk worse if they are not recorded and removed afterwards.

Build fault handling into the design and operating procedures. Identify critical links and power dependencies, document redundancy, monitor interface errors and link state, and define how temporary changes are authorised. Ensure that network and security teams know which alerts indicate a possible loss of visibility, not just a service interruption.

Ask the practical question: what happens at 2am when this fails? The answer should include who gets called, what information they need, how to isolate the fault and how to restore the intended security posture. If the only available answer is “find someone who remembers how it was patched”, the design is not operationally ready.

Make network segmentation a design property

Network segmentation is a way to limit unnecessary communication and contain the impact of compromise. It can support separation between user groups, servers, management networks, guest access and OT. However, creating

VLANs alone does not guarantee useful security boundaries. Segmentation depends on capable switching, a correct design and enforcement points that control traffic between segments.

Begin with communication requirements. Identify which systems need to talk, what protocols and services they require, and why. Then design segments and enforcement so that permitted traffic is explicit and unnecessary paths are blocked. Depending on the environment, enforcement may involve firewalls, routing policy, access control lists or other mechanisms. The design must account for where traffic actually flows, including wireless networks, remote access, shared services and failover routes.

Switching capability matters. The chosen platform needs to support the required features, scale and management model, and it must be configured consistently. Products from vendors such as Cisco may be used in enterprise networks; industrial environments may also use specialist equipment from providers such as MOXA. Garland is another name organisations may encounter in network infrastructure. The appropriate choice depends on the operating environment and requirements—not on a brand name alone. Confirm the platform's security capabilities, support lifecycle, management options and fit for the proposed design.

Be wary of unmanaged switches in production. They can hide device connections, complicate traffic tracing and make the intended segmentation difficult to verify. Where an unmanaged device is genuinely necessary, document its purpose, physical location, attached equipment and risk treatment. In a controlled design, it should not become an invisible bridge between trust zones.

Test the segmentation, don't just document it

After implementation, verify both allowed and denied communications. Check that each segment can reach only the services it needs, that management interfaces are restricted, and that alternative paths do not bypass enforcement. Review switch configurations and firewall rules against the approved design. A network diagram showing separate VLANs is not proof that the boundaries are effective.

Testing should also account for operational requirements. In OT, for instance, a broad block rule introduced without application knowledge may interrupt a process. Work with system owners to understand necessary flows and plan changes in a controlled way. The objective is not maximum isolation at any cost; it is deliberate, risk-based communication with controls that can be maintained.

Connect infrastructure decisions to ISM controls

ACSC guidance is most useful when it is translated into controls people can implement, operate and evidence. The ISM is not a substitute for understanding the environment. Use it to inform risk treatment and control selection, then show how the network design supports the organisation's security objectives.

Infrastructure area Security contribution Evidence to retain Asset and connection records Helps identify what is connected and who is responsible for it Current inventory, diagrams, port schedules and ownership records Segmentation and boundary enforcement Limits unnecessary communications between systems and trust zones Approved diagrams, rule reviews and connectivity test results Secure device management Reduces the risk of unauthorised changes to network infrastructure Configuration standards, access records and change approvals Monitoring and incident response Supports detection, investigation and timely containment Log source records, alert tests and incident procedures Physical access and cabling Helps prevent or detect unauthorised connections and tampering Access procedures, cabinet records, labels and inspection results

This is not a one-to-one control checklist: a single infrastructure measure may support several security outcomes, and a control may require more than one measure. Keep the mapping specific to the organisation. Record the rationale for design choices, any accepted residual risks, and the person accountable for reviewing them.

Secure the network devices themselves

Switches, routers, wireless controllers and firewalls are security-relevant assets. Apply a consistent device-hardening process that includes supported software, approved configuration baselines, restricted administrative access and secure management paths. Remove default credentials and unused services, use strong authentication appropriate to the environment, and protect configuration backups. Changes should be authorised, logged and recoverable.

Management traffic should be separated from ordinary user traffic where the design allows, with access limited to authorised administrators and systems. Use centralised logging and time synchronisation so events can be correlated. Where a device or network segment cannot provide the required telemetry, document the limitation and consider compensating measures.



Firewalls should enforce deliberate policy, not become a dumping ground for exceptions that nobody can explain. Review rules for ownership, purpose and continued need. Endpoint detection platforms provide valuable visibility on supported endpoints, but they do not replace network controls or physical safeguards. A strong approach combines endpoint, network and operational evidence rather than relying on a single tool to reveal every threat.

Build security into the project lifecycle

Infrastructure work aligns more reliably with ACSC guidance when security is included from discovery through operations. Use a repeatable lifecycle:

1. **Discover:** Inventory devices, links, services, owners, dependencies and physical locations. Validate records on site where possible.
2. **Assess:** Identify risks, required communications, trust boundaries, availability needs and applicable ISM controls.
3. **Design:** Produce logical and physical diagrams, segmentation policy, management design, monitoring plan and failure behaviour.
4. **Implement:** Use approved configurations, controlled change windows, labelled cabling and documented rollback steps.
5. **Verify:** Test connectivity, blocked paths, failover, alerting, logging and access restrictions against agreed acceptance criteria.

6. **Operate:** Maintain records, review rules and access, monitor faults, and remove temporary connections after the work is complete.

Include network, security, facilities and system owners in the right stages. The team responsible for a production line may know about a required protocol that is absent from the diagrams. Facilities staff may know which cabinet shares power with another service. That knowledge is part of the security picture.

Measure whether the foundations are holding

Compliance evidence should be supported by operational checks. Useful measures can include the percentage of network devices with approved configurations, the age of unresolved diagram discrepancies, the number of unreviewed firewall rules, and the time required to identify the owner of a connected device. Track link faults and monitoring gaps alongside security events so recurring physical problems are not dismissed as routine noise.

Use findings to improve the design and maintenance process. If technicians regularly discover undocumented switches, improve change control and site records. If segmentation exceptions keep multiplying, revisit the underlying application requirements. If a cutover repeatedly loses monitoring, make telemetry validation a formal acceptance test.

Conclusion

Aligning network infrastructure with ACSC guidance is not simply a matter of buying security products or mapping a project to a control list. It means making the physical and logical network support the security outcomes the organisation needs. Reliable cabling, controlled access, accurate records, capable switching and tested network segmentation help firewalls, endpoint detection platforms and other controls work as intended.

The ISM and broader ACSC guidance provide a valuable basis for risk-informed decisions, but those decisions need sound network foundations beneath them. Label the panels, remove the mystery cables, validate the paths and plan for failure. When the network is understandable and its boundaries are real, security teams can see more clearly, operators can recover faster, and the organisation is better prepared for the 2am call.