

Network infrastructure projects are often treated as facilities work: install the cabling, fit the switches, connect the racks and hand the keys to IT. That approach misses a critical point. Physical infrastructure is part of cybersecurity. The condition of the cabling, the capability of the switches and the way the network is designed all affect whether an organisation can apply security controls effectively.

For Australian organisations, aligning infrastructure work with **ACSC guidance** means considering security from the first site survey through to commissioning and ongoing maintenance. The Australian Cyber Security Centre's Information Security Manual (ISM) provides guidance and controls for protecting systems and information. Those controls do not replace sound network engineering. They rely on it. If the physical and logical foundations are unreliable, security measures may be difficult to implement, hard to verify or easy to bypass.

This article explains how to translate ACSC guidance into practical decisions about cabling, switching, network segmentation, monitoring and operational resilience. The aim is not to turn every fit-out into a compliance exercise. It is to build networks whose security can be designed, tested and maintained.

Start with the security outcome, not the shopping list

A project brief that says "replace the switches" is not enough to guide a secure design. Start by identifying what the network must protect, which systems need to communicate and what failure would mean to the organisation. Include corporate IT, operational technology (OT), wireless networks, management interfaces, building systems and any third-party connections in scope.

Then map those needs to relevant ACSC guidance and the organisation's security obligations. The ISM is maintained over time, so teams should consult the current publication and the controls applicable to their environment rather than rely on an old checklist or a vendor's interpretation. Record the control intent, the design decision that supports it and how the decision will be verified.



Here's what kills me: for example, a requirement to restrict access to sensitive systems needs more than a firewall order. It may depend on correctly identifying assets, placing them in appropriate network zones, controlling routes between zones and limiting administrative access to network devices. (my cat just knocked over my water). Last month, I was working with a client who learned this lesson the hard way.. A useful project question is: what evidence will show this control is working after handover?

Physical infrastructure is part of cybersecurity

Cabling, patch panels, racks and switch locations affect who can connect to a network and how easily teams can identify unauthorised or unexpected connections. A live socket with no label, an undocumented patch lead or an unmanaged switch added under a desk can create a path that was never considered in the security design.

Physical security and logical security should therefore be planned together. Consider where network equipment is located, who can access it, how changes are authorised and how ports are managed. Protecting a server room while leaving accessible network outlets in public or shared areas may not match the threat model. Equally, an access-controlled cabinet is not a substitute for appropriate network controls.

Good documentation is a security measure, not administration for its own sake. Maintain accurate rack elevations, patch-panel schedules, cable labels, switch port descriptions and network diagrams. Label both ends of cabling where practical, and keep records aligned with the installed environment. A label maker in the laptop bag is a small investment compared with hours spent tracing a mystery cable during an incident.

Design for change and incident response

Infrastructure changes should have an approval path, an implementation record and a rollback plan. Keep configuration backups for managed network [network labelling](#) devices, protect those backups appropriately and document the expected state. Where feasible, compare the live configuration with the approved baseline.

During an incident, responders may need to isolate a device, trace a connection or confirm which systems share a path. If the patching and diagrams do not reflect reality, containment takes longer and investigators may miss important traffic. A tidy, labelled physical layer helps teams answer practical questions quickly: which port is this, what is connected to it, and what happens if we disable it?

Physical faults create outages and security blind spots

A damaged fibre, a loose patch lead, an overloaded switch or an incorrectly configured trunk can cause obvious service disruption. Less obviously, the same fault can create a security blind spot. A monitoring feed might stop reaching its destination. A sensor could be connected to the wrong segment. A failover path might bypass expected inspection, or a network team might restore service by introducing an undocumented temporary connection.

During an outage, pressure to “get it back up” is real. That is precisely when unauthorised workarounds tend to appear. Build resilience into the design, document recovery steps and define who can approve emergency changes. After the incident, review temporary configurations and remove them promptly. Otherwise, the workaround can become a permanent, poorly understood part of the network.

For each critical link or device, ask what happens at 2am when it fails. Is there a tested spare or redundant path? Can the on-call engineer identify the connection? Is there a supported escalation route? Does failover preserve the intended security controls? Resilience is not simply uptime; it is the ability to recover without quietly weakening protections.

Network segmentation depends on capable switching and correct design


Network segmentation limits unnecessary connectivity and helps contain compromise. It can separate business functions, user groups, management traffic, guest access, servers and OT environments. But segmentation is not

achieved merely by creating VLANs. It depends on a design that controls where traffic can flow, on equipment that supports the required functions, and on configuration that is consistent and maintained.

Begin with communication requirements. Identify which systems need to talk, in which direction, using which services, and why. Avoid broad rules such as “allow the whole corporate subnet” where a narrower policy is practical. Define trust boundaries, document permitted paths and review them as systems and business needs change.

Switching provides important building blocks: VLANs, access control features, management-plane protections and, in some designs, routing between segments. The exact capabilities vary by product and model. Select equipment against the security and operational requirements, not just port count or purchase price. Vendors such as Cisco and MOXA offer a range of switching products for different environments, including enterprise and industrial settings. The right choice depends on the site, support requirements, environmental conditions and the features needed for the design.

Where traffic crosses trust boundaries, use controls appropriate to the risk. These may include firewalls, access control lists or other enforcement points. A firewall rulebase should reflect documented business flows, have an accountable owner and be reviewed. Avoid treating the firewall as a magic boundary: mispatched ports, alternate routes and poorly controlled device management can undermine the intended architecture.



Do not confuse segmentation with isolation

Segmentation reduces exposure, but it does not guarantee that compromise cannot cross between zones. Review inter-zone rules, management access, shared services and remote support paths. Pay particular attention to temporary links, vendor access and connections between corporate IT and OT. A design that looks separated on a diagram may still have broad reachability in practice.

Test the implemented design. Confirm that authorised communications work and that prohibited paths are blocked. Validate configurations on the installed devices, not only in design documents. Where the environment changes frequently, establish a repeatable review process so that new ports, routes and firewall rules do not gradually erode the boundaries.

Connect infrastructure controls to ISM controls

ISM controls are more useful when teams turn them into clear implementation and assurance activities. Infrastructure work commonly supports broader security outcomes such as controlled access, secure configuration, protection of network communications, logging and incident response. The specific controls that apply depend on the system, its classification, the operating context and the organisation's obligations.

Build a traceability table that links the project requirement to the design, implementation evidence and owner. This makes it easier to see where a control depends on multiple layers—for example, physical access, switch configuration and firewall policy—and where a gap remains.

Security objective	Infrastructure contribution	Example verification evidence
Restrict unauthorised connections	Controlled equipment rooms, managed switch ports, documented outlets and approved change processes	Access records, port configuration review and current patching documentation
Limit lateral movement	Defined network zones, controlled inter-zone paths and appropriately configured switching	Approved network diagram, rule review and segmentation test results
Support detection and investigation	Reliable connectivity for logs, monitoring and security tools, with time and traffic sources understood	Successful log delivery, monitoring checks and documented incident procedures
Recover securely from failure	Tested redundancy, configuration backups, labelled cabling and controlled emergency changes	Recovery test records, backup checks and post-change review

The table is a starting point, not a substitute for assessing the applicable ISM controls. Use it to make dependencies visible and to assign responsibility. Some actions belong to facilities, some to network operations and some to security. The project must connect those teams rather than assume one group will cover everything.

Make monitoring part of the network design

Security tools need dependable data and connectivity. Firewalls enforce traffic policies at defined points. Endpoint detection platforms help identify suspicious activity on supported endpoints. Neither can compensate for an unknown network path, missing assets or a monitoring connection that silently stopped working.

Include logging and monitoring requirements in the design. Decide which network devices and security controls will generate logs, where those logs will be sent, who will review alerts and how connectivity failures will be detected. Protect management interfaces and monitoring paths, and make sure network changes do not interrupt them without an approved reason.

For OT and other specialised environments, monitoring and active testing must be planned carefully. Some systems have strict availability or compatibility requirements. Work with system owners and vendors to understand operational constraints, then choose an approach that provides useful visibility without creating avoidable risk.

Commission the network as a security deliverable

Handover should include evidence that the network matches the approved design. Do not accept “it pings” as the commissioning test. Check the physical installation, device inventory, software and configuration baselines, segmentation, management access, monitoring and resilience.

A practical commissioning sequence could include:

1. Inspect racks, patch panels and cable management; confirm labels and physical access arrangements.
2. Reconcile the installed equipment and port connections against the inventory and diagrams.
3. Review device configurations, management access, authentication and approved services.

4. Test required and prohibited traffic paths between network zones.
5. Confirm that logs and alerts reach their intended monitoring destinations.
6. Test failover and recovery procedures where the design requires them.
7. Record deviations, assign owners and close high-risk items before operational handover.

Keep test results with the project records, and make sure the operational team receives usable diagrams, configuration backups, support details and change procedures. A document pack that nobody can understand is not an effective handover.

Keep alignment after the project ends

Security alignment is not a one-off sign-off. Networks change as people move, equipment is replaced, applications are added and suppliers request access. Establish periodic reviews of segmentation rules, device configurations, asset records and physical access. Update diagrams when changes are made, not months later during an audit or incident.

Include network infrastructure in vulnerability and risk management processes. Track vendor support status, apply approved updates and assess the effect of changes on both security and availability. Ensure unmanaged switches are not introduced into production as a quick fix. If a temporary device is genuinely required, control it, document it and set a removal date.

Finally, make ownership explicit. Facilities teams, network engineers, security teams and system owners each hold part of the picture. A short review before a fit-out can prevent expensive redesign later; a clear escalation path can prevent an emergency workaround from becoming permanent.

Conclusion

Aligning network infrastructure work with ACSC guidance starts with recognising that physical and logical security are connected. Cabling faults and undocumented changes can cause outages, weaken segmentation and create blind spots. Strong network segmentation relies on capable switching, deliberate traffic flows and tested enforcement—not just VLAN names on a diagram.

Use the current ISM and relevant ACSC guidance to shape requirements, then translate those requirements into designs, implementation evidence and operational responsibilities. Build in labelling, monitoring, resilience and change control from the beginning. When the network is clear, supportable and testable, security controls have a solid foundation—and the on-call engineer has a better chance of understanding what is happening when something fails at 2am.