

When organisations start planning network segmentation, the first question is often, “Which VLANs should we create?” It is a useful question, but not the first one to answer. Start with risk: what could an attacker reach if a device or account is compromised, and what would cause the most harm if it failed?

That makes segmentation priorities specific to each environment. A business with exposed legacy servers may need to isolate those before anything else. A manufacturer with operational technology (OT) may need to contain a vulnerable controller. An office with unmanaged guest access may have a simpler, more urgent first step: stop guest devices reaching corporate systems.

Good VLAN planning is not a substitute for security architecture. VLANs provide logical separation, but the design only works when switching, routing, firewalls and physical infrastructure are configured to enforce it. The goal is risk-based design: limit unnecessary access, preserve business operations and make it possible to see what is happening across the network.

Start with risk, not a standard VLAN list

There is no universal rule that says servers always come first, or that guest Wi-Fi is always the biggest concern. Rank the work by considering four things:

- **Impact:** What is the likely consequence if this device or network is compromised or unavailable?
- **Exposure:** How many users, devices or external connections can reach it?
- **Trust and condition:** Is it managed, patched and monitored, or is it old, poorly documented or running a default configuration?
- **Containment:** How much access could an attacker gain from this starting point?

Also check dependencies before changing anything. A server may rely on a particular workstation, a printer may communicate with a print server, and an OT controller may depend on a vendor’s maintenance path. Mapping those flows helps avoid turning a security change into an outage.

As a practical baseline, every environment should have a clear boundary for guest access, separation for critical services, and a plan for IoT and OT devices that cannot meet normal endpoint security requirements. The order and scope of those changes should follow the organisation’s actual risks.


A practical order for segmentation priorities

1. Isolate guest Wi-Fi from internal networks

Guest Wi-Fi is often the quickest, clearest place to establish a boundary. Visitors’ phones and laptops are outside the organisation’s control. They should not be able to browse file shares, management interfaces, printers or internal services just because they have been given internet access.

Put guest traffic in its own network and allow only the services it needs, typically internet access and the supporting services required to connect. Confirm that the wireless configuration, switching and firewall rules all enforce the boundary. A separate Wi-Fi name alone does not prove that guest clients are isolated.

Then test it. Connect a guest device and cyberdaily.au check whether it can reach internal addresses, discover devices or access administrative pages. If the network is provided by a third party, clarify who owns the configuration and who verifies it. The 2am question matters here too: if the guest network is misconfigured, can someone identify and disable it without disrupting business-critical wireless?



2. Protect high-value and exposed servers

Servers can contain valuable data, provide identity services, or support essential business processes. They are also common targets, particularly when exposed to the internet, reachable from broad user networks, or running services that no one has reviewed for years.



Group servers by function and sensitivity rather than putting every system into one large “server VLAN”. Apply access rules that permit only required connections. For example, users may need to reach a specific application, while administrative access should come from designated management systems and authorised accounts. Public-facing services should not have a broad, trusted route into internal infrastructure.

Use firewalls to enforce and log boundaries between important security zones, including where traffic must be inspected or tightly controlled. Pair that with patching, secure configuration, backups and endpoint detection platforms where the server operating system and workload support them. Segmentation limits paths; it does not make an unpatched server safe.

3. Contain IoT and OT devices

IoT is a broad category: cameras, access control, building systems, environmental sensors and specialist equipment may all have different needs. Many devices have limited security features, infrequent updates or communications that are not obvious from their labels. OT adds another constraint: a change that is acceptable for office IT could interrupt a physical process.

Inventory these devices before designing rules. Record their location, owner, purpose, vendor, required communications and whether they can be patched or monitored. Give them appropriately restricted network zones, then permit only the necessary flows to management systems, collectors or approved vendor access paths. Avoid allowing a camera or building controller to initiate arbitrary connections to user devices or servers.

Industrial suppliers such as MOXA provide networking equipment used in industrial environments, while the exact segmentation approach must still reflect the site's architecture, operating requirements and safety constraints. Do not assume that using industrial-grade equipment automatically creates a secure design. Validate the complete path, from the device port through switching and routing to the systems it communicates with.

Coordinate OT changes with operations and vendors. Use maintenance windows and test plans, and understand what a rollback looks like. In some environments, passive discovery and carefully observed traffic capture are safer first steps than active scanning.

4. Separate user devices and constrain administration

User endpoints are numerous and exposed to email, web content, removable media and account compromise. Separating user networks from servers reduces the chance that a compromised laptop can communicate freely with sensitive systems. Consider whether different user groups need different access, especially for contractors, shared devices and privileged administrators.

Administration deserves special attention. Management interfaces for switches, firewalls, servers and wireless systems should not be reachable from every ordinary user subnet. Use a dedicated management network or equivalent controlled access model, with strong authentication and logging. Keep switch and firewall administration accounts unique and remove default credentials—"it is only on the internal network" is not a security control.

Endpoint detection platforms can help identify suspicious activity on supported user devices and servers. Their alerts become more useful when network boundaries are known and monitored: responders can identify which zone a device belongs to and limit its access without guessing which port or patch lead it uses.

VLAN planning: the logical design must match the real network

VLANs are useful building blocks, but they are not a security policy by themselves. A VLAN separates Layer 2 traffic; routing between VLANs may still allow broad access unless controls are deliberately applied. Define which paths are permitted, where they are enforced, and how changes will be reviewed.

A sensible VLAN plan uses clear names, documented purposes and consistent addressing. Avoid creating a VLAN for every small variation without an operational reason, but do not combine systems with materially different trust levels just to keep the configuration tidy. The design should be understandable by the engineer who has to troubleshoot it after hours.

Zone Typical priority Design consideration Guest Wi-Fi Early, low-complexity boundary Internet access only; verify isolation from internal networks Servers High when sensitive or exposed Group by function and permit only required user, service and admin flows IoT and OT High where devices are fragile or high-impact Document

dependencies; restrict communications without disrupting operations User devices Essential baseline Limit access to servers and management systems; account for contractors and shared devices Network management Essential for control-plane security Restrict administrative access and log changes

Physical infrastructure is part of cybersecurity

Logical segmentation depends on physical infrastructure that behaves as expected. A port connected to the wrong patch-panel outlet, a mislabelled cable, or an unmanaged switch added under a desk can undermine the design. Mystery cables are not merely an inconvenience; they are a security smell because they make it harder to know which device is connected where and which controls apply.

Physical layer faults can create outages and security blind spots at the same time. A failed uplink may cause devices to move to an unexpected path. A patching error can place an endpoint in the wrong VLAN. An undocumented switch can bypass monitoring or expose ports that the approved design assumes are controlled. During incident response, uncertainty about the physical path slows containment and makes evidence harder to interpret.

That is why cabling is not “just plumbing”. Label both ends of cables and patch-panel ports, maintain accurate diagrams, secure network rooms and cabinets, and control access to switch ports. Use managed switches in production so that port configuration, VLAN membership, authentication and monitoring can be administered consistently. Vendors such as Cisco and Garland offer switching products used in different network environments, but product selection does not replace sound design, configuration or lifecycle management.

Before a segmentation rollout, verify switch capabilities and configuration. Check that required VLANs are supported along the full path, that trunks carry only the intended VLANs, and that unused ports are disabled or otherwise controlled. Confirm that the switching design can deliver the performance and resilience the site needs. If the hardware cannot enforce or report the intended controls, the architecture may exist only on paper.

Align with ACSC guidance—and the foundations it relies on

Australian Cyber Security Centre (ACSC) guidance promotes layered security, controlled access and sound network practices as part of reducing cyber risk. Those controls are most effective when the underlying network is understood and managed: assets are inventoried, systems have owners, configurations are maintained, and network paths can be verified.

In practice, that means treating segmentation as one part of a broader security program, not as a quick compliance exercise. Define the permitted flows, apply least privilege, review rules regularly and monitor activity at the boundaries. Keep diagrams and inventories current enough that responders can use them. Validate the design against the real site rather than relying on a spreadsheet that has not been checked since the last fit-out.

ACSC-aligned security objectives are difficult to achieve when ports are undocumented, equipment is unmanaged or teams do not know which systems depend on one another. Solid network foundations are not optional groundwork to do later; they are what make segmentation enforceable and incident response practical.

A safe way to begin

1. **Inventory:** Identify network devices, endpoints, servers, IoT and OT assets, their owners and their connections.

2. **Map flows:** Record which systems communicate, including vendor access, identity services, DNS, logging and management traffic.
3. **Rank risks:** Prioritise exposed or high-impact systems, uncontrolled access paths and devices that are difficult to patch or monitor.
4. **Design boundaries:** Define VLANs and firewall rules around trust and business function, not convenience alone.
5. **Test and phase:** Pilot changes, monitor for unexpected dependencies, document rollback steps and schedule disruptive work appropriately.
6. **Verify physically:** Trace the actual cabling and switch ports, update labels and diagrams, and remove unauthorised equipment.
7. **Review:** Check logs, rule sets and asset records regularly, and revisit the design when systems or business processes change.

So, what should you segment first?

For many organisations, guest Wi-Fi is a good early win because the boundary is clear and the required access is limited. Next, focus on high-value or exposed servers, then contain IoT and OT according to operational risk. Build strong separation for user and management networks as a foundational control—not an optional later improvement.

But let the evidence decide. The right segmentation priorities come from understanding what is connected, what it can reach and what a failure would mean. A well-labelled, managed network with carefully enforced boundaries gives security teams something reliable to defend. And when something breaks at 2am, it gives the on-call engineer a much better chance of finding the right cable, port and policy before the business feels the impact.