Cisco Ise Design Guide

INTRODUCTION TO CISCO ISE DESIGN GUIDE

Cisco Identity Services Engine (ISE) is a cornerstone of modern network access control (NAC) and policy enforcement. Deploying ISE successfully requires more than just installing software; it demands a well-thought-out architecture that aligns with your organization's security, scalability, and operational requirements. This **Cisco ISE Design Guide** provides a comprehensive overview of the key considerations, best practices, and architectural patterns that will help you build a robust and future-proof ISE deployment. Whether you are planning a greenfield implementation or migrating from an older platform, understanding the design principles behind ISE is essential for achieving consistent policy enforcement, high availability, and seamless integration with your existing network infrastructure.

UNDERSTANDING ISE NODE ROLES AND PERSONAS

Before diving into design specifics, it is critical to understand the different node roles (also called personas) that an ISE appliance can assume. Each persona performs specific functions, and the way you combine them directly impacts performance and redundancy.

Primary Administration Node (PAN) and Secondary PAN

The PAN manages all administrative tasks, including configuration, policy creation, and reporting. For high availability, you can deploy a secondary PAN that synchronizes configuration data from the primary. In active/passive mode, the secondary takes over only if the primary fails. The **Cisco ISE design guide** recommends using the PAN for administration only — avoid assigning other personas to it unless your deployment is very small.

Policy Service Node (PSN)

PSNs are the workhorses that handle authentication, authorization, and accounting (AAA) requests, posture assessment, and guest services. They are the nodes that network devices (switches, wireless controllers, VPN concentrators) communicate with directly. For most designs, PSNs should be dedicated to policy services only. You can scale PSNs horizontally to handle increased traffic or to provide geographic distribution.

Monitoring and Troubleshooting Node (MnT)

The MnT collects logs, alarms, and reports from all other nodes. It can also serve as a syslog collector and provide live monitoring. For large deployments, dedicate separate MnT nodes to avoid performance impact. Redundancy is achieved by pairing primary and secondary MnT nodes that synchronize operational data.

pxGrid Node

The pxGrid (Platform Exchange Grid) node enables integration with third-party security tools such as SIEMs, firewalls, and threat detection systems. It shares contextual information (e.g., user identity, endpoint security status) with other platforms. In a production environment, deploy pxGrid on a dedicated node or co-locate it with the PAN only if the deployment is moderate.

SINGLE-NODE VS. DISTRIBUTED DEPLOYMENTS

One of the first decisions in any **Cisco ISE design guide** is whether to use a single-node or a distributed architecture. A single-node deployment (all personas on one appliance) is suitable for lab environments or very small branch offices with minimal authentication traffic. However, production networks almost always require a distributed design to ensure availability, performance, and disaster recovery.

Small-Scale Design (Up to 5000 Endpoints)

For small environments, a two-node pair (one primary PAN/MnT and one secondary PAN/MnT) plus one or two PSNs may suffice. However, the official **Cisco ISE design guide** advises against combining PAN and PSN on the same node in production because a failure affects both administration and policy enforcement. Instead, consider a dedicated PSN even for small offices.

Medium-to-Large-Scale Design

As the number of endpoints and authentication requests grows, you need separate nodes for each persona. A typical mid-sized design might include:

- Two PAN nodes (primary and secondary) – administration only.
- Two MnT nodes (primary and secondary) – monitoring only.
- Multiple PSN nodes (at least two for redundancy) distributed across data centers or campus cores.
- Optional pxGrid node for integration.

For large enterprises with thousands of concurrent authentications, the design becomes more complex, often involving regional PSN clusters, load balancers, and geographic redundancy. The **Cisco ISE design guide** provides sizing tables based on peak concurrent sessions, authentication rates, and logging volume.

POSITIONING ISE WITHIN YOUR NETWORK TOPOLOGY

Where you place ISE nodes in your network topology directly affects latency, failover behavior, and the user experience. The PSNs must be reachable by all network devices that enforce policies (e.g., switches, WLCs, firewalls). Latency between the network device and the PSN should be minimal—preferably within the same campus or metropolitan area.

Centralized vs. Distributed PSNs

In a centralized design, all PSNs reside in a primary data center. This simplifies management but introduces latency for remote sites. The **Cisco ISE design guide** typically recommends a hybrid approach: deploy a few PSNs in central data centers for headquarters and additional PSNs at large regional hubs to reduce authentication latency. For small remote sites, consider using Cisco ISE's "Proxy" mode (via network device AAA proxy) to forward requests to a central PSN, or deploy a local PSN if the site has more than a few hundred users.

Using Network Devices (NADs) as Proxies

Many Cisco switches and wireless controllers can act as RADIUS proxies, forwarding requests to an ISE PSN based on failover groups. This allows you to create logical tiers without additional hardware. The design guide stresses the importance of configuring RADIUS > prioritize local PSNs first, then fallback to a central PSN. This ensures that even if a local PSN is down, authentication still succeeds.

REDUNDANCY AND HIGH AVAILABILITY (HA) BEST PRACTICES

A well-designed ISE deployment must be resilient. The **Cisco ISE design guide** outlines several HA mechanisms:

PAN HA

Use a primary/secondary pair with automatic synchronization. The secondary PAN remains idle until a failure is detected. Ensure both PANs are the same hardware model and run the same software version. Avoid placing PANs in different geographic regions unless you have a high-latency link and are using manual failover.

PSN HA

PSNs operate in an active/active model. Use multiple PSNs and configure network devices with multiple RADIUS servers using primary and fallback groups. For critical authentication (e.g., 802.1X), ensure that the PSNs can handle the total load if one fails. The design guide recommends a N+1 redundancy model (e.g., two PSNs for peak load, plus an extra for failover).

MnT HA

Use a primary/secondary MnT pair. Logs are replicated continuously. To avoid losing data during a failover, ensure the network storage (if used) is reliable and low-latency. The **Cisco ISE design guide** also advises periodic backup of configuration and operational data.

INTEGRATION WITH CISCO DNA CENTER AND OTHER TOOLS

Modern network management increasingly involves automation and intent-based networking. ISE integrates with Cisco DNA Center (formerly APIC-EM) to provide a unified policy lifecycle. The design guide explains how to connect ISE as a policy repository and endpoint profiler. Similarly, integration with Cisco Stealthwatch, Firepower, and third-party SIEMs via pxGrid enriches security responses.

When designing for integration, consider the number of pxGrid sessions, the throughput of contextual data, and the security zones. The **Cisco ISE design guide** recommends isolating pxGrid traffic on a separate VLAN or using certificates for mutual authentication. For large-scale deployments, dedicate a separate vm or appliance solely for pxGrid to avoid competition for resources.

ENDPOINT PROFILING AND GUEST SERVICES

ISE's profiling capabilities (active and passive) help identify device types, operating systems, and users. The design should account for the profiling probe placement: active probes (e.g., nmap scans) should be used sparingly to avoid network impact, while passive probes (e.g., DHCP, HTTP, NetFlow) are less intrusive. The guide suggests using dedicated collectors or enabling passive probes on PSNs if the volume is manageable.

For guest services (sponsored, self-registration), ISE can host a captive portal. The portal pages can be customized, but the design must consider scalability: each concurrent guest portal session consumes PSN resources. For large guest networks, use dedicated PSNs or external web servers.

PERFORMANCE SIZING AND CAPACITY PLANNING

A central part of any **Cisco ISE design guide** is sizing. ISE performance depends on the number of concurrent authentications, session duration, logging levels, and enabled features (e.g., posture assessment, device registration). Cisco provides a **ISE Sizing Tool** that estimates required CPU, memory, and disk based on inputs like total endpoints, peak authentication rate, and logging. Always over-provision slightly — plan for at least 20% headroom for growth.

Key performance metrics to monitor in the design phase:

- Average authentication response time (target < 1 second).
- Maximum throughput per PSN (measured in authentications per second – APS).
- Disk I/O for MnT logs (high message rates can saturate storage).
- Network bandwidth between ISE nodes (replication traffic, RADIUS packets).

The design guide recommends using 10 GbE links for inter-node communication in large deployments and placing MnT nodes on separate compute resources with fast disks (SSD recommended).

COMMON DESIGN MISTAKES TO AVOID

Even experienced engineers can fall into traps. Here are frequent pitfalls highlighted in the **Cisco ISE design guide**:

1. **Oversizing or undersizing the deployment** – Relying on guesswork instead of using the sizing tool.
2. **Incorrect node persona assignment** – Combining PAN and MnT on the same node in a production environment,

which leads to degraded performance during peak loads.

3. **Ignoring certificate management** – ISE heavily uses certificates for EAP, portal, and node-to-node communication. Plan a certificate authority (internal or public) and lifecycle early.
4. **Not planning for failover testing** – Many designs look good on paper but fail during actual node outage because RADIUS server groups are misconfigured or timeouts are too short.
5. **Underestimating log volume** – Default logging can generate terabytes of data per day. Tune log levels based on necessity and archive older logs.

CONCLUSION

Designing a Cisco ISE deployment is a complex but manageable task when you follow a structured approach. This **Cisco ISE design guide** has covered the fundamental building blocks: node personas, topology placement, high availability strategies, integration points, and performance considerations. By carefully planning each element and leveraging Cisco's official documentation and sizing tools, you can create an ISE architecture that delivers secure, consistent, and scalable network access control. Always remember that a good design is not static — it should evolve with your network requirements and emerging threats. Regular reviews and capacity adjustments will ensure that your ISE deployment remains a reliable pillar of your security infrastructure.