Enterprise Wireless Network Transformation with Cisco SD-Access

2026-09-11

© 2026 Cisco and/or its affiliates. All rights reserved.

Contents

Executive summary 2
Enterprise manufacturing network: scale, challenges, and context 2
Network scale, legacy challenges, and manufacturing context 2
Wireless best practices: fabric wireless and central switching architecture 3
Introduction 3
SD-Access fabric roles 3
The role of overlays in SD-Access 3
Security through segmentation 4
Architectural shift: centralized versus distributed 4
Configuration-based differences for migration 5
Telemetry 5
Guest and BYOD access 5
Quality of service (QoS) 5
DHCP enforcement 5
mDNS (Bonjour) gateway 6
Peer-to-peer blocking 6
Application QoS scale 6
DHCP required on SSID 6
Default-handled attributes with Catalyst Center 6
Strategic migration insights 6
Conclusion 7

Executive summary

Enterprise fabrics must deliver seamless wired and wireless mobility, robust segmentation, and programmability at scale. Cisco SD-Access meets these goals with an overlay control plane based on LISP, integrated TrustSec segmentation, and a distributed data plane for fabric-enabled wireless. This document explains how fabric-enabled wireless differs from traditional centrally switched designs, identifies which network services operators must rethink and possibly redesign (for example, DHCP, mDNS, QoS trust, and telemetry), and describes what Catalyst Center automates by default. It concludes with practical migration guidance for organizations moving WLAN service to the fabric.

Enterprise manufacturing network: scale, challenges, and context

Network scale, legacy challenges, and manufacturing context

Large enterprise manufacturing organizations often operate heterogeneous campus and plant-floor networks spanning thousands of switches, access points (APs), and connected endpoints. Environments at this scale—typically thousands of switches, tens of thousands of wireless APs, and tens of thousands of wired and wireless clients—amplify any inconsistency in design, software revision, or operational practice.

© 2026 Cisco and/or its affiliates. All rights reserved.

Before adopting a fabric-based architecture, many organizations face similar baseline conditions that make transformation both necessary and complex.

Historically, these organizations face limited visibility into network-connected devices, difficulty standardizing device profiles and policies, and fragmented operational sources of truth that complicate end-to-end assurance and reliable alerting. Switch configurations vary by site, engineer, or integrator preference, and a broad sprawl of Catalyst software versions often exists across the installed base. Automation is frequently limited, so changes remain manual or dependent on ad hoc scripting. On the plant floor, Industry 4.0 and factory automation workloads stress legacy Wi-Fi designs: a shared wireless medium cannot guarantee deterministic latency, and safety-related use cases require a deliberately engineered wireless architecture (for example, higher-density AP placement where automated guided vehicles operate, with APs remounted for consistent RF coverage).

To address manufacturing IT/OT convergence, leading enterprises are evolving toward IT-managed, production-engineering-managed, and OT-bridged or gatewayed models—culminating in SD-Access manufacturing network initiatives that deliver segmentation, zero-trust alignment, analytics consumption, and automation at enterprise scale.

Wireless best practices: fabric wireless and central switching architecture

Introduction

Enterprise fabrics must deliver seamless mobility, segmentation, and programmability. This section compares fabric-enabled wireless with traditional centrally switched designs, including control-plane and data-plane roles, policy placement, and operational attributes that change when WLAN traffic is integrated with Cisco SD-Access.

SD-Access fabric roles

Before comparing centralized and fabric-enabled wireless, it is helpful to align on common SD-Access terms:

- Underlay: The physical IP network formed by routed campus switches. The underlay provides connectivity between fabric nodes using a Layer 3 routed-access design.
- Fabric/overlay: The collection of virtual networks (overlays) running on the underlay. An overlay is a logical topology that connects endpoints to the network and provides isolation.
- Fabric edge node: The access switch where endpoints and APs attach. In EVPN designs, this role is analogous to a leaf switch; in LISP-based SD-Access, it is the edge node where VNIDs, anycast gateways, and SGACL enforcement are applied.
- Extended node: A simplified access switch (typically Layer 2) that extends fabric connectivity to ports or areas that cannot support full edge-node capabilities.
- Border and control plane node: The control plane provides reachability between the fabric and external networks. Border nodes act as gateways for traffic entering and exiting the fabric.
- Wireless LAN controller (WLC): Retains the wireless control plane (association, roaming, RF management) but is not in the client data path for fabric-enabled SSIDs.

The role of overlays in SD-Access

The SD-Access architecture uses fabric technology to provide wired and wireless campus networks with programmable overlays and network virtualization. A physical underlay—formed by routed campus switches and routers—hosts one or more logical overlay networks. The underlay establishes IP connectivity using a

© 2026 Cisco and/or its affiliates. All rights reserved.

Layer 3 routed-access design and is engineered for reliability, scalability, and resiliency; end-user subnets and endpoints exist in the overlay, not the underlay. While the underlay is configured and tuned once, the overlay manages users, devices, mobility, and segmentation.

A fabric is an overlay network created through encapsulation—a process that adds headers to the original frame—forming a logical topology of tunnels that connects devices over the physical underlay. Each overlay is a virtual network (VN) provisioned as a VRF instance that isolates routing tables; data-plane traffic and control-plane signaling remain contained within each VN. Multiple overlay networks can run across the same underlay, enabling role-based segmentation with VRFs and Security Group Tags (SGTs), consistent wired and wireless policy, and seamless mobility without spanning large Layer 2 broadcast domains across the campus.

An overlay network in SD-Access relies on a sophisticated control plane to maintain real-time mappings of endpoints to their network locations. Cisco SD-Access uses the Locator/ID Separation Protocol (LISP) as its primary control plane. Functionally, LISP operates as a demand-based protocol similar to DNS, maintaining a mapping database that is queried when a connection is required and subsequently cached. This model keeps the overlay lightweight while supporting the services expected of a modern enterprise network.

Security through segmentation

As security threats increasingly move laterally, segmentation has evolved from an optional control to a baseline requirement. In the Cisco SD-Access architecture, segmentation is integrated into the fabric using the Instance ID (IID), which "colors" both control-plane exchanges and data-plane traffic. This scopes traffic into specific VRFs or bridge domains so that segmentation is enforced consistently without undermining fundamental forwarding behavior.

Software-defined segmentation integrates with Cisco TrustSec technology, providing micro-segmentation for groups inside a virtual network using SGTs. Automating virtual networks with integrated security via Cisco Catalyst Center reduces operational expense and risk.

The Cisco SD-Access fabric represents the next generation of switching infrastructure. While early demand-based fabrics such as Cisco Application Centric Infrastructure (Cisco ACI) used conversational learning to track endpoint locations, standards-based implementations have matured. By adopting LISP—a protocol with substantial deployment history and IETF formalization—SD-Access provides an optimized, standards-aligned foundation for the enterprise.

Architectural shift: centralized versus distributed

Organizations that operate a centralized wireless architecture should understand the architectural and operational differences before moving to fabric-enabled wireless inside SD-Access.

In a traditional centralized Cisco unified wireless architecture, control and data traffic are typically tunneled to a centralized controller for policy enforcement and forwarding decisions. Fabric-enabled wireless uses a distributed data plane: traffic is switched locally at the edge while the SD-Access fabric enforces policy and segmentation.

From a configuration perspective, SSID attributes such as security policies and RF parameters remain largely consistent and can often be retained during migration, which preserves the client experience and limits large-scale WLAN rework.

The data forwarding model, however, changes materially. In a centralized (non-fabric) architecture, client traffic is tunneled to the WLC, which applies policy and places traffic on the correct egress VLAN. A single SSID can map to different VLANs or policies per user or client group because enforcement occurs on the

© 2026 Cisco and/or its affiliates. All rights reserved.

WLC data path. With a fabric-enabled SSID, traffic is forwarded from the AP to the upstream fabric edge or extended node. Traditional VLAN semantics give way to Virtual Network Identifiers (VNIDs) that define logical segmentation and forwarding.

Two primary VNID types appear in fabric designs: Layer 2 VNID and Layer 3 VNID. For wireless deployments, client traffic maps to a Layer 2 VNID, which plays a role analogous to a VLAN in classic designs and defines broadcast domain behavior for wireless clients.

This evolution delivers practical benefits:

- Simplified policy: Cisco ISE serves as the single policy authority, pushing SGACLs to fabric edge nodes for consistent enforcement across wired and wireless.
- Simplified segmentation: Endpoints authenticate to Cisco ISE and receive SGTs via 802.1X, MAB, or static assignment. The same SGT model applies whether the device connects over Wi-Fi or a wall-port switch.
- Scale without WLC bottlenecks: The distributed data plane ensures aggregate throughput is not capped by WLC capacity. For example, a Catalyst 9800-M at ~50 Gbps can be saturated by a handful of high-throughput APs, yet supports thousands of APs for control. Each switch port/uplink becomes the local bottleneck.
- Fast roaming without centralized switching: WLCs manage the wireless control plane for seamless roam (as in centralized mode), while anycast gateways on each fabric edge keep default gateways local, delivering centralized-style mobility with distributed data-plane scale. This allows for high-scale seamless fast roaming across the fabric.

Configuration-based differences for migration

When migrating to fabric-enabled wireless, administrators move from a WLC-centric model to a distributed model in which the fabric edge plays a larger role. Key differences include the following.

Telemetry

In centrally switched deployments, NBAR2 visibility is managed by the WLC. In fabric-enabled environments, visibility is achieved by explicitly enabling wireless telemetry in Catalyst Center.

Guest and BYOD access

Both modes can use preauthentication ACLs, but the enforcement location differs. Centrally switched designs enforce at the controller; fabric-enabled wireless pushes policy from the WLC to the AP, where enforcement occurs.

Quality of service (QoS)

APs enforce per-user and per-radio QoS in both modes. Fabric edge nodes are expected to honor DSCP markings; administrators should verify that the edge is configured to trust those markings during migration.

DHCP enforcement

Centrally switched networks often configure DHCP on the Catalyst 9800 series WLC. Fabric-enabled wireless typically requires DHCP snooping and IP source guard on the fabric edge rather than relying on the WLC data path.

© 2026 Cisco and/or its affiliates. All rights reserved.

mDNS (Bonjour) gateway

In centrally switched designs, the WLC may act as the Bonjour gateway. In a fabric deployment, this function typically resides on the fabric edge node or the Layer 2 border, so mDNS configuration on the 9800 is not the primary placement.

Peer-to-peer blocking

Centrally switched networks may rely on explicit controller configuration to block unwanted peer traffic. Fabric-enabled designs often use SGT-based policy on the edge to permit only desired peer-to-peer flows.

Application QoS scale

Fabric-enabled wireless supports a bounded QoS schema (for example, eleven classes with three applications per class in typical deployments). Customers should audit existing QoS policies before migration because legacy centrally switched designs sometimes exceed those limits.

DHCP required on SSID

In SD-Access fabric deployments, enabling the DHCP Required option can introduce additional processing overhead on the WLC. When this feature is enabled, the WLC maintains IP-to-MAC binding information and distributes these bindings to the associated APs.

In large-scale fabric deployments, the replication and synchronization of IP-to-MAC bindings across APs increases the controller's processing overhead. As the number of clients and APs grows, this additional overhead impacts the performance and scalability of the WLC.

For this reason, enabling the DHCP Required option is not recommended for SD-Access fabric wireless deployments.

Default-handled attributes with Catalyst Center

A major benefit of fabric-enabled wireless under Catalyst Center is automation of manual tasks. The following are representative defaults when wireless is configured through Catalyst Center.

- QoS trust boundary: shifts from the WLC to the fabric edge.
- Network access control: VLAN assignment remains RADIUS-driven while gateway functions move to the fabric edge.
- IP theft protection: handled with Switch Integrated Security Features (SISF) on the fabric edge rather than on the 9800.
- ARP handling: LISP control-plane awareness reduces dependence on WLC ARP proxy or flood suppression for the same scenarios.
- DHCP handling: DHCP relay is handled by the fabric edge, removing the WLC from the data path.
- Multicast: enabled through Catalyst Center, with APs replicating traffic to clients according to the defined policy.

Strategic migration insights

The transition to fabric-enabled wireless offers distinct advantages, especially for large-scale industrial environments.

- Seamless roaming: An anycast gateway on each fabric edge helps roaming devices (for example, automated guided vehicles) avoid hairpinning traffic back to an anchor controller.

© 2026 Cisco and/or its affiliates. All rights reserved.

- Traffic optimization: local forwarding at the switch reduces unnecessary tromboning when wireless and wired endpoints in the same policy zone communicate, saving bandwidth and latency.
- Symmetric policy: a unified approach to wired and wireless operational-technology endpoints using a single subnet style with macro- and micro-segmentation (SGTs) enforces consistent security without subnet sprawl.

Conclusion

Migrating to fabric-enabled wireless is a shift from a WLC-centric model to a distributed, policy-driven fabric. By moving enforcement points for DHCP, mDNS, QoS, and telemetry to the fabric edge, organizations achieve better scalability, more efficient traffic paths, and consistent security.

Organizations should audit existing QoS classes, application mappings, and protocol controls before cutover so that fabric capabilities align with earlier centrally switched behavior.

© 2026 Cisco and/or its affiliates. All rights reserved.