Network Topography & Routing Guide

Enterprise Campus & Data Center Fabric — Design Reference

Meridian Dynamics

Network Engineering & Infrastructure Group

Document Property	Value
Document Owner	Network Engineering — Core Infrastructure Team
Version	4.2
Status	Approved for Production Use
Effective Date	March 3, 2026
Review Cycle	Semiannual, or upon topology change
Classification	Internal — Distribution restricted

Prepared by Elena Cho, Principal Network Architect
Approved by Marcus Webb, Director of Infrastructure

This document supersedes Version 4.1 (2025-09-15).

Contents

1 Introduction & Design Principles

1.1 Purpose and Scope

This guide is the authoritative reference for the physical and logical network architecture of Meridian Dynamics' primary campus and its paired data center. It documents the routed topology, IP addressing plan, interior and exterior routing policy, VLAN segmentation, and high-availability mechanisms currently in production. It is intended for network engineers, security operations staff, and infrastructure vendors performing authorized change work.

The scope covers the headquarters campus (Meridian Tower, Singapore), the primary data center (DC-East), the disaster-recovery site (DC-West), and three regional branch offices. It does not cover application-layer configuration, endpoint management, or the public-facing CDN, which are documented separately by the Platform Engineering group.

1.2 Design Principles

The fabric is built around four principles that every change request is evaluated against:

- **No single point of failure.** Every layer from transit circuit to access switch is deployed in an active/standby or active/active pair.
- **Deterministic addressing.** Every subnet is derived from a documented supernet so that any address can be traced to a site and function without querying a live system.
- **Defense in depth.** Traffic between trust zones always crosses at least one stateful inspection point; no zone trusts another implicitly.
- **Change is reversible.** Every routing and VLAN change ships with a pre-verified rollback step before it is scheduled.

At a Glance

Sites in scope: **5** (HQ, DC-East, DC-West, 3 branches shown as one class)
Core address space: **10.0.0.0/8**
Interior routing: **OSPFv2**, 4 areas
Exterior routing: **eBGP**, dual transit, AS 65001
Core switching: redundant pair, MLAG
Firewall posture: active/standby, stateful failover

2 Network Topology

2.1 Topology Overview

The fabric follows a classic three-tier design — core, distribution, access — with a dedicated security tier sitting between the Internet edge and the core. Figure 2.1 shows the logical layout for the HQ campus; DC-East mirrors the same pattern with the distribution tier replaced by top-of-rack leaf switches.

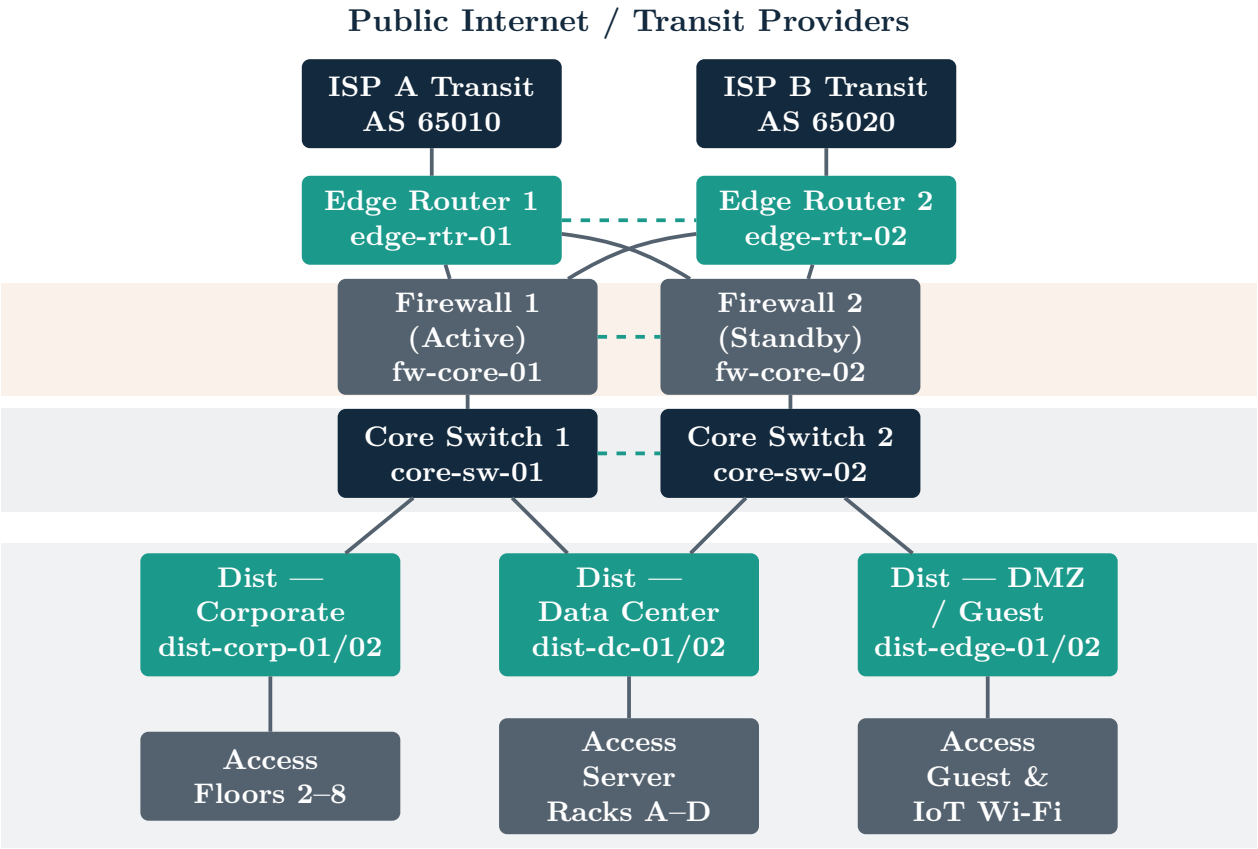

Figure 2.1: Logical topology of the HQ campus fabric, including the Internet edge and security tier.

2.2 Site Inventory

Table 2.1 lists every routed site in the fabric and its role in the overall design. All inter-site links are encrypted point-to-point circuits terminated on the edge routers shown in Figure 2.1.

Site	Role	Primary Link	Backup Link
HQ — Meridian Tower	Campus, user access	2 × 10 Gbps dark fiber	1 Gbps LTE failover
DC-East	Primary data center	40 Gbps DWDM to HQ	10 Gbps DWDM to DC-West
DC-West	DR / warm standby	10 Gbps DWDM to DC-East	1 Gbps Internet VPN
Branch — Austin	Regional office	500 Mbps MPLS	200 Mbps SD-WAN broadband
Branch — Lisbon	Regional office	500 Mbps MPLS	200 Mbps SD-WAN broadband
Branch — Nairobi	Regional office	200 Mbps MPLS	100 Mbps SD-WAN broadband

Table 2.1: Site inventory and circuit redundancy.

3 IP Addressing Plan

3.1 Address Space Summary

All internal address space is carved from the `10.0.0.0/8` private block. Each site receives a dedicated `/16`, and each functional zone within a site receives a `/20` or smaller reserved range, leaving headroom for growth without renumbering. Table 3.1 is the top-level allocation; Section 3.2 breaks each site down to individual VLANs.

Site / Zone	Supernet	Reserved For
HQ Campus	10.10.0.0/16	User, voice, and management VLANs across 8 floors
DC-East	10.20.0.0/16	Server, storage, and out-of-band management subnets
DC-West (DR)	10.21.0.0/16	Mirrors DC-East allocation for failover consistency
Branch — Austin	10.30.0.0/16	Local user and voice VLANs
Branch — Lisbon	10.31.0.0/16	Local user and voice VLANs
Branch — Nairobi	10.32.0.0/16	Local user and voice VLANs
DMZ / Guest (shared)	172.16.0.0/16	Internet-facing services, guest Wi-Fi, IoT
Point-to-point links	192.168.0.0/22	/31 transit links between routed devices
Future data center	10.40.0.0/16	Reserved, unallocated

Table 3.1: Top-level supernet allocation.

⚠ Operational Warning

The `10.40.0.0/16` reservation for the planned DC-Central build-out must not be assigned to any interim project. Two prior incidents trace back to VLANs borrowed from an "unused" future block that later conflicted with a live migration.

3.2 VLAN and Subnet Allocation

Table 3.2 enumerates every production VLAN at HQ and DC-East. Branch offices follow the same VLAN ID scheme (100s for user, 200s for voice, 900s for management) against their own site supernet.

VLAN	Name	Subnet	Gateway	DHCP Range
110	HQ — User Data (Floors 2–4)	10.10.0.0/22	10.10.0.1	.0.50–.3.250
120	HQ — User Data (Floors 5–8)	10.10.4.0/22	10.10.4.1	.4.50–.7.250
210	HQ — Voice / SIP	10.10.16.0/23	10.10.16.1	.16.20–.17.250
310	HQ — Guest Wi-Fi (in DMZ block)	172.16.10.0/24	172.16.10.1	.10.20–.10.250
320	HQ — IoT / Building Systems	172.16.20.0/24	172.16.20.1	.20.20–.20.250
910	HQ — Network Management	10.10.250.0/24	10.10.250.1	Static only
920	HQ — Security Cameras & Badging	10.10.251.0/24	10.10.251.1	Static only
501	DC-East — Application Tier	10.20.1.0/24	10.20.1.1	Static only
502	DC-East — Database Tier	10.20.2.0/24	10.20.2.1	Static only
503	DC-East — Storage / iSCSI	10.20.3.0/25	10.20.3.1	Static only
510	DC-East — Load Balancer VIPs	172.16.30.0/25	172.16.30.1	Static only
930	DC-East — Out-of-Band Mgmt (IPMI)	10.20.250.0/24	10.20.250.1	Static only

Table 3.2: VLAN and subnet allocation, HQ and DC-East (abbreviated addresses share the row's /16).

4 Routing Architecture

4.1 Interior Routing — OSPFv2

The campus and data center backbone runs OSPFv2 in four areas. Area 0 is the transit backbone linking core and distribution devices; each site's user- and server-facing subnets are stubbed into their own totally-stubby area to keep the link-state database small on branch routers with limited memory.

Area	Scope	Type
0.0.0.0	Core-to-distribution backbone, core-sw-01/02, dist-*-01/02	Backbone
0.0.0.10	HQ campus user and voice VLANs	Totally stubby
0.0.0.20	DC-East and DC-West server, storage, and OOB VLANs	Totally stubby
0.0.0.30	Branch offices (Austin, Lisbon, Nairobi)	Totally stubby

Table 4.1: OSPF area design.

4.2 Exterior Routing — BGP

Internet transit is dual-homed with route-based failover and no default static route — the edge routers rely entirely on learned BGP routes so that failover behavior matches production during a planned circuit test.

Peer	Peer ASN	Local ASN	Policy
ISP A (primary)	65010	65001	Prefer via local preference 200; full table
ISP B (secondary)	65020	65001	Local preference 150; full table, AS-path prepend outbound
DC-West DR link	65001	65001	iBGP over encrypted VPN; used only on DC-East failure

Table 4.2: External BGP peering summary.

- Only the summarized `10.0.0.0/8` aggregate and the DMZ `172.16.0.0/16` block are advertised outbound; component subnets are filtered by an outbound prefix list.
- Inbound, only the default route and two backup provider routes are accepted; a maximum-prefix limit of 25 stops an accidental full-table leak from either peer.
- Redistribution between BGP and OSPF is one-directional (BGP → OSPF) and tagged with community `65001:900` so it can never be redistributed back out.

Operational Warning

Never remove the `65001:900` no-re-export community from the BGP-to-OSPF redistribution route-map. Doing so during the 2024 edge upgrade briefly leaked internal OSPF routes toward ISP B and triggered a maximum-prefix session reset.

5 VLAN Segmentation & Security Zones

5.1 Trust Zone Model

Every VLAN is assigned to exactly one of four trust zones. Traffic crossing a zone boundary always transits a firewall context; there is no routed path between zones that bypasses inspection.

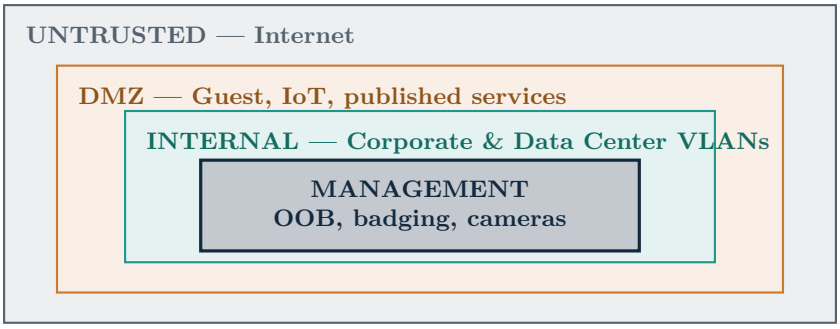

Figure 5.1: Nested trust zones: traffic only moves inward or outward through firewall policy, never laterally without inspection.

5.2 Firewall Policy Summary

Table 5.1 summarizes the default posture between zones. All rules are logged; any traffic not explicitly permitted is denied and alerts on the first match.

Source Zone	Destination Zone	Allowed Services	Action
Untrusted	DMZ	HTTPS, DNS to published resolvers only	Permit, logged
DMZ	Internal	None by default	Deny
Internal	DMZ	HTTPS to published services only	Permit, logged
Internal	Management	SSH, HTTPS, SNMPv3 from jump hosts only	Permit, logged
Management	Internal	None outbound except patch mirrors	Permit, logged
Untrusted	Management	—	Deny, alert on match

Table 5.1: Default inter-zone firewall policy.

6 High Availability & Maintenance

6.1 Redundancy Protocols

First-hop redundancy uses HSRP on all campus and data center distribution pairs, with priority tuned so the primary distribution switch owns the virtual gateway under normal conditions and preemption is disabled to avoid unnecessary flapping during brief link blips.

Group	Virtual IP	Protocol	Members
HSRP 110	10.10.0.1	HSRPv2	dist-corp-01 (active), dist-corp-02 (standby)
HSRP 120	10.10.4.1	HSRPv2	dist-corp-01 (active), dist-corp-02 (standby)
HSRP 501	10.20.1.1	HSRPv2	dist-dc-01 (active), dist-dc-02 (standby)
VRRP 310	172.16.10.1	VRRPv3	dist-edge-01 (master), dist-edge-02 (backup)

Table 6.1: First-hop redundancy group assignments.

6.2 Maintenance Window Policy

Design Principle

Standard maintenance windows run Sundays 01:00–05:00 SGT. Any change touching the core, security, or edge tiers requires a peer-reviewed rollback plan and a pre-staged configuration diff attached to the change ticket before the window opens.

6.3 Monitoring & Alerting

- BGP session state, OSPF neighbor adjacency, and HSRP/VRRP role changes page the on-call engineer within 60 seconds of a state transition.
- Interface utilization above 80% sustained for 5 minutes on any core or edge link opens a capacity-planning ticket automatically.
- Synthetic transactions verify end-to-end reachability between HQ, DC-East, and DC-West every 30 seconds; three consecutive failures trigger failover validation.

7 Appendix

7.1 Glossary

OSPF
: Open Shortest Path First — interior gateway routing protocol used within the fabric.

BGP Border Gateway Protocol — exterior routing protocol used with ISP transit and the DR site.

HSRP / VRRP
: Hot Standby / Virtual Router Redundancy Protocol — first-hop gateway failover mechanisms.

DMZ Demilitarized Zone — a segregated network exposing services to less-trusted networks.

MLAG
: Multi-Chassis Link Aggregation — allows a pair of switches to present a single logical link-aggregation partner to downstream devices.

7.2 Revision History

Version	Date	Author	Summary
4.2	2026-03-03	Elena Cho	Added DC-West DR peering, updated BGP max-prefix limits
4.1	2025-09-15	Priya Nair	Reworked VLAN scheme for guest/IoT separation
4.0	2025-02-20	Elena Cho	Migrated core to MLAG pair, retired legacy STP root
3.3	2024-07-08	Marcus Webb	Added Nairobi branch, corrected supernet table

Table 7.1: Document revision history.

7.3 Approval & Contacts

Role	Name	Contact
Principal Network Architect	Elena Cho	ext. 4021
Director of Infrastructure	Marcus Webb	ext. 4002
Security Engineering Lead	Priya Nair	ext. 4187
Network Operations Center	On-call rotation	NOC bridge, 24/7

Table 7.2: Document approval chain and operational contacts.