Carrier-Grade Infrastructure Solution Brief: Deploying Huawei Switch BGP EVPN Configuration and Route Reflector Topology

MARKET POSITIONING

The convergence of data center virtualization and multi-tenant service delivery demands a control plane that scales without introducing protocol churn or operational fragility. Huawei CloudEngine series switches, when deployed with BGP EVPN as the VXLAN control plane and a Route Reflector (RR) topology, address this requirement directly. This solution brief positions the BGP EVPN RR architecture not merely as a configuration option, but as a strategic simplification of the overlay control plane. By centralizing EVPN route reflection, the architecture eliminates the full-mesh BGP peer requirement among VTEPs, reducing control-plane session count from $O(n^2)$ to $O(n)$ while preserving the policy granularity and multi-tenancy isolation that carrier and enterprise operators require. The target deployment is the leaf-spine data center fabric where spine nodes assume the RR role, or where a dedicated pair of RR-capable switches provides route reflection for a larger VTEP domain.



HIGH-AVAILABILITY REDUNDANCY

Route Reflector redundancy is non-negotiable in carrier-grade deployments. A single RR instance represents a control-plane single point of failure; Huawei’ s implementation supports RR client peering across two independent RR nodes, with BGP best-path selection and deterministic route reflection ensuring that VTEPs maintain EVPN route visibility if either RR becomes unavailable. The RR cluster-id mechanism prevents routing loops when multiple RRs serve the same client set, while BGP graceful restart and BFD for BGP sessions provide sub-second failure detection and session preservation across control-plane restarts. On the VTEP side, the anycast VTEP and M-LAG features complement the RR topology, ensuring that both the control plane (EVPN route distribution) and the data plane (VXLAN encapsulation and forwarding) maintain active-active or active-standby redundancy across the fabric. For EVPN

multi-homing scenarios, the ESI-based DF election operates independently of the RR topology, preserving multi-homing semantics even when route reflection is centralized.

PROTOCOL INTEROPERABILITY

The BGP EVPN implementation on Huawei switches conforms to RFC 7432 (BGP MPLS-Based Ethernet VPN) and RFC 8365 (Network Virtualization Overlay with EVPN), with full support for the EVPN route types required in VXLAN deployments: Type 2 (MAC/IP Advertisement), Type 3 (Inclusive Multicast Ethernet Tag), and Type 5 (IP Prefix). The RR function operates within the BGP EVPN address family (`l2vpn-family evpn`), and the `peer reflect-client` command designates clients without requiring client-side awareness of their RR status . For interworking with existing BGP domains, the EVPN address family can be enabled alongside IPv4/IPv6 unicast families, and route policies applied at the RR can manipulate next-hop attributes, extended communities, and route targets before reflection. The architecture also supports VXLAN-to-VXLAN and, where required, VXLAN-to-MPLS EVPN interworking through the same RR infrastructure, providing a migration path without control-plane redesign.

DETAILED PARAMETERS

The following parameters represent the validated configuration envelope for a typical RR-based BGP EVPN deployment on Huawei CloudEngine and S-series switches. Values are drawn from Huawei’ s official configuration guides and datasheet documentation for current-generation platforms.

Parameter	**Specification**
Form Factor	1RU / 2RU fixed chassis; modular chassis available in CloudEngine 16800/12800 series
Switching Capacity	Up to 12.8 Tbps (CE8850-64CQ-EI); 16 Tbps (CE8855H / CE8856); higher capacities in modular platforms
Power Supply	1+1 redundant, hot-swappable AC/DC/HVDC power modules; AC input 90–290 V
BGP EVPN Address Family	l2vpn-family evpn; enabled per peer or peer group
RR Configuration Command	peer <ipv4-address> reflect-client (within BGP-EVPN address family view)
RR Cluster ID	reflector cluster-id <id> (applies only to RRs; prevents loops in multi-RR

	topologies)
EVPN Route Types Supported	Type 2 (MAC/IP), Type 3 (Inclusive Multicast), Type 5 (IP Prefix)
VXLAN VNI Scale	Up to 16 million tenants (platform-dependent)
BGP EVPN Session Encryption	TCP MD5 authentication supported on BGP sessions
BFD for BGP	Hardware-based BFD supported for sub-second failure detection

LIFECYCLE ASSURANCE (MTBF)

Huawei CloudEngine data center switches are engineered for continuous operation, with hardware Mean Time Between Failures (MTBF) figures aligned to carrier-grade availability targets. The RR function itself is a software role executed on the switch’ s control-plane CPU; it does not alter the hardware reliability profile of the platform. In an RR topology, the control-plane load is concentrated on the RR nodes, which must be sized with sufficient CPU and memory headroom to handle the aggregate EVPN route table and BGP session churn from all clients. Huawei’ s CloudEngine platforms provide redundant supervisor modules (on modular chassis) and redundant power and fan

subsystems across the portfolio, ensuring that the failure of any single hardware component does not disrupt EVPN route reflection. Software maintenance activities — such as ISSU (In-Service Software Upgrade) on supported platforms—allow RR software updates without dropping BGP EVPN sessions to clients, preserving overlay reachability throughout the maintenance window.

TARGET NETWORK TOPOLOGIES

The canonical deployment topology places the Route Reflector function on the spine layer of a leaf-spine fabric. Each leaf switch acts as a VTEP and establishes a BGP EVPN session only with the spine RRs, rather than with every other leaf. In a two-spine fabric, each leaf maintains two BGP sessions (one per spine); a 64-leaf fabric with two RRs therefore requires 128 BGP EVPN sessions total, versus 4,032 sessions in a full-mesh design. This reduction directly translates to lower control-plane CPU utilization, faster convergence after topology changes, and simpler troubleshooting. For larger fabrics or multi-site deployments, a dedicated RR pair (or set) can serve multiple fabrics, with route targets and route policies enforcing tenant isolation across sites. The RR topology also integrates cleanly with Huawei’ s CloudFabric solution, where the RR nodes may be managed through the iMaster NCE controller for automated underlay and overlay provisioning.


SYS
FAN
PSU
SYS ALM
FAN
PSU
INIT
READY
V300R019
INIT
V300R019
MGMT
CON
MGMT
USB
RST
PSU
INIT
READY
READY
V300R019