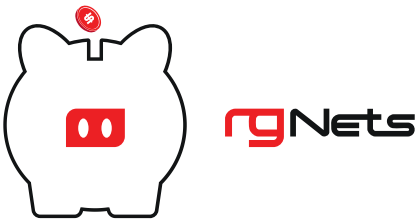

VXLAN vs GRE for Software-Defined Access Networks

A Topology-First Decision Framework for SDAN Deployments

Document ID: RGN-WP-VXLAN-GRE
Version: 5.0
Date: April 2, 2026

Contents

1 Executive Summary

The choice between Virtual eXtensible Local Area Network (VXLAN) and Generic Routing Encapsulation (GRE) in a Software-Defined Access Network (SDAN) is not a binary decision. It is a multi-dimensional engineering evaluation shaped by the IP path between initiator and concentrator, the capabilities of edge hardware, and the operational requirements of the deployment. The following factors are any one of which can independently determine the choice:

- **Network Address Translation (NAT) / Carrier-Grade NAT (CGNAT) on the path?** Use VXLAN. Its UDP encapsulation allows NAT devices to create session mappings—but the overlay still requires control-plane coordination to discover NAT-translated endpoints and maintain bindings. The rXg platform provides this coordination natively.
- **East-west traffic between campus edge devices?** Use VXLAN. In the standard SDAN model, initiator traffic hairpins through the concentrator; in campus extensions, managed switches with VXLAN Tunnel Endpoint (VTEP) capability can forward locally without hairpinning.
- **Equal-Cost Multi-Path (ECMP) across spine paths?** Use VXLAN. Its flow-hashed UDP source port provides per-flow entropy for load balancing.
- **Fleet-wide segmentation via control plane?** Use VXLAN. Border Gateway Protocol (BGP) Ethernet VPN (EVPN) propagates VXLAN Network Identifier (VNI) definitions without per-device static configuration.
- **Broad hardware acceleration?** Favor VXLAN. First-class Data Processing Unit (DPU) and switch Application-Specific Integrated Circuit (ASIC) offload are broadly available; GRE offload varies by vendor and may lack advanced operations.
- **Constrained firmware (Optical Network Terminals, cable modems)?** Use GRE. Simpler encapsulation is easier to implement and validate.
- **Passive Optical Network (PON) / Data Over Cable Service Interface Specification (DOCSIS) point-to-point underlay?** Use GRE. Natural fit for direct Layer 2 paths with no ECMP or east-west requirements.
- **Cost-sensitive edge, CPU-only, no boundary crossings?** Use GRE. Lower overhead, fewer control-plane dependencies, and ∼38% lower CPU utilization (5% vs 8% at 1 Gbps) in software-only deployments without DPU offload.

Both protocols are valid, production-proven choices. The rXg platform supports both natively and can deploy them simultaneously across a mixed portfolio—GRE on private-path sites, VXLAN where NAT traversal, ECMP, or operational scale require it—managed from a single control plane.

This whitepaper explains the Book-Ends architecture that frames the decision, details how each protocol works, and presents the multi-factor decision framework with quantified trade-offs. Five deployment examples—on-premises concentrators, central-office hosting, distributed remote sites, hospitality with Direct Internet Access (DIA), and mixed-protocol portfolios—illustrate the frame-

work in practice.

2 The Book-Ends Architecture

SDAN deployments share the same fundamental structure: two purpose-built endpoints—the **initiator** and the **concentrator**—with a distribution underlay between them that serves as transparent transport. This is the Book-Ends architecture. Understanding it is essential before evaluating encapsulation protocols, because the protocol choice is a property of the path between the two book-ends, not a property of the venue type or the number of access points.

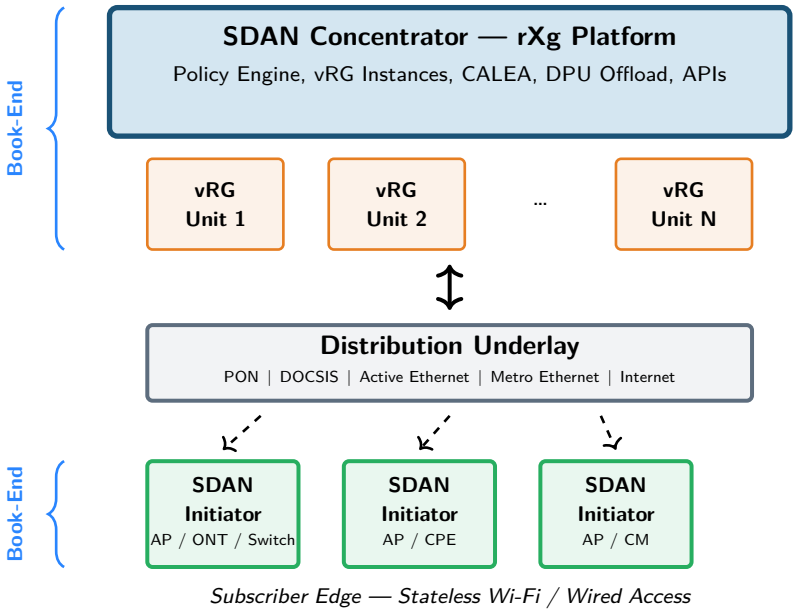

Figure 1: SDAN Book-Ends Architecture: Concentrator, Underlay, and Initiators

2.1 The Concentrator: rXg as Intelligence Anchor

The SDAN concentrator is the headend of every deployment. On the RG Nets platform, the rXg *is* the concentrator. It hosts all per-subscriber virtual residential gateway (vRG) instances, implements the policy engine that governs network behavior, enforces Communications Assistance for Law Enforcement Act (CALEA)-compliant subscriber attribution, and exposes northbound APIs for fleet orchestration. It runs in the operator's infrastructure—most commonly in a Communications Service Provider (CSP) central office, a regional data center, or on-premises at the property headend.

The concentrator is also the overlay termination point. Every subscriber overlay—whether GRE or VXLAN—terminates at the concentrator. The rXg decapsulates inbound traffic, routes it through the appropriate vRG instance (one per subscriber), applies policy via the integrated Broadband Network Gateway (BNG), and re-encapsulates outbound traffic for delivery to the initiator. On BlueField-3 DPU-equipped SKUs, encapsulation and decapsulation are hardware-accelerated for

both protocols, though VXLAN receives broader offload support (including header modification and field copy operations that GRE lacks in the NVIDIA DOCA Flow API).

Because the rXg is simultaneously the concentrator *and* the control plane, the encapsulation decision is not a vendor ecosystem choice. There is no separate decision about BGP EVPN vs. Locator/ID Separation Protocol (LISP) vs. static flooding. The rXg control plane manages overlay state for both GRE and VXLAN natively. BGP EVPN is used as the overlay signaling protocol for VXLAN deployments, but it is an implementation detail of the platform, not a competitive architectural decision the operator must evaluate independently.

2.2 The Initiator: Stateless Subscriber Edge

The SDAN initiator is the device at the subscriber edge—a Wi-Fi access point, an ONT with integrated wireless, a cable modem, a lightweight Customer Premises Equipment (CPE) device, or a high-density Power over Ethernet (PoE) switch at an Intermediate Distribution Frame (IDF), as in active Ethernet deployments for hospitality, large public venues, or Multi-Dwelling Unit (MDU) / Multi-Tenant Unit properties. Its responsibilities are deliberately narrow: 802.11 association and encryption, overlay encapsulation of subscriber traffic toward the concentrator, and basic link-layer operations. All stateful functions—policy enforcement, subscriber identity, Quality of Service (QoS), CALEA compliance, service segmentation—are implemented in the vRG on the concentrator, not in the initiator.

This statelessness is architecturally significant for the encapsulation decision. The initiator's firmware must implement the encapsulation protocol. GRE requires implementing IP Protocol 47 encapsulation with an 8-byte header (base GRE plus key)—a minimal addition to any IP stack. VXLAN requires implementing UDP encapsulation on port 4789, a VXLAN header with VNI lookup, and a source-port hashing function for ECMP entropy. On general-purpose Linux-based APs, both are straightforward. On constrained embedded platforms—low-cost ONTs, DOCSIS cable modems, Internet of Things (IoT) gateways—GRE's simpler encapsulation path is substantially simpler to implement and validate in firmware.

2.3 The Underlay: Transparent Transport

The distribution underlay connects initiators to the concentrator. It may be a PON tree, a DOCSIS Hybrid Fiber-Coaxial (HFC) plant, an active Ethernet campus, a metro Ethernet ring, a carrier Multiprotocol Label Switching (MPLS) network, or the public internet. The Book-Ends architecture treats the underlay as transparent IP transport. The overlay carries subscriber traffic end-to-end; the underlay merely forwards encapsulated packets.

This transparency means the underlay's characteristics—its maximum transmission unit (MTU), its NAT behavior, its routing topology—directly determine which overlay protocol is appropriate. A PON underlay with direct L2 adjacency between ONT and Optical Line Terminal (OLT) provides a

private IP path with no NAT: GRE works naturally. A hospitality deployment where a remote property connects to a cloud-hosted concentrator through the public internet crosses NAT boundaries: VXLAN is the right choice. The underlay does not care which protocol it carries, but its properties constrain which protocol *should* be carried.

3 How GRE and VXLAN Work

Both protocols create overlays across an IP underlay, but they encapsulate at different layers and produce meaningfully different overhead profiles. Understanding these mechanics is essential before applying the topology-first decision framework.

Note

Layer Distinction: VXLAN carries complete inner Ethernet frames (Layer 2), adding 50 bytes of overhead including the inner Ethernet header. Although base GRE (RFC 2784) can carry Layer 3 payloads directly, the rXg platform uses L2oGRE (GRE with Transparent Ethernet Bridging, protocol 0x6558), which carries full Ethernet frames—the same Layer 2 capability as VXLAN. The key architectural difference is segmentation: L2oGRE uses inner 802.1Q VLAN tags within a single GRE overlay, while VXLAN uses VNIs natively.

3.1 GRE: Minimal Overhead, Point-to-Point

The rXg platform uses L2oGRE—GRE (RFC 2784/2890) configured to carry Ethernet frames (Transparent Ethernet Bridging, EtherType 0x6558). The encapsulation adds an outer IP header (20 bytes), a GRE header with key (8 bytes), and an inner 802.1Q VLAN tag (4 bytes) for service segmentation, producing 32 bytes of encapsulation overhead (46 bytes total wire overhead including the inner Ethernet header). The protocol field (IP Protocol 47) identifies GRE traffic, and all overlay endpoints must be explicitly provisioned. The outer header contains only source IP, destination IP, and protocol number—no UDP/TCP ports—providing minimal entropy for underlay ECMP hashing.

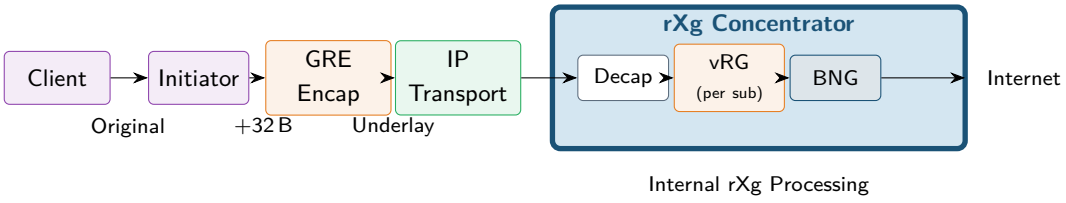

Figure 2: GRE Encapsulation Flow: Initiator to Concentrator

Strengths for SDAN. GRE's 32-byte encapsulation (outer IP + GRE header + key + inner VLAN tag) preserves slightly more payload headroom than VXLAN on 1500-byte links (inner IP MTU = 1454, TCP Maximum Segment Size (MSS) ≈ 1414—a ∼3.2% reduction vs. VXLAN's ∼3.4%). Its minimal overlay definition (source IP, destination IP, protocol) is implemented across major IP stacks—Linux, Juniper, Cisco, embedded real-time operating system (RTOS)—with no control plane dependencies. On single-core or resource-constrained initiators, GRE achieves ∼95%

line-rate efficiency at 10 Gbps in software and ∼38% lower CPU utilization than VXLAN at 1 Gbps throughput. The simpler outer header makes packet captures easier to interpret with standard tools. Most significantly for SDAN, GRE's encapsulation path is simple enough to implement reliably in constrained embedded firmware on ONTs, cable modems, and low-cost APs—devices where code complexity directly correlates with defect rates.

Trade-offs. GRE requires a separate overlay for each initiator-concentrator pair; overlay count grows linearly with initiator count (though not quadratically in the SDAN model, since all overlays terminate at the concentrator). Each new initiator requires explicit overlay configuration. GRE's static outer headers prevent ECMP load balancing across multiple spine paths. Most critically, GRE uses IP Protocol 47 with no TCP/UDP ports, making it operationally impractical through NAT or CGNAT—many NAT devices can track only one GRE session per public IP per destination.

Note on GRE-in-UDP (RFC 8086). Wrapping GRE inside a UDP header (GRE-in-UDP, RFC 8086) can theoretically address NAT traversal for GRE. However, GRE-in-UDP lacks widely available hardware acceleration—it is not recognized as a standard tunnel type in the NVIDIA DOCA Flow API or Open vSwitch (OVS) hardware offload paths as of this writing. Its overhead (36–40 bytes with key) equals or exceeds VXLAN's 36 bytes, eliminating GRE's overhead advantage. In practice, GRE-in-UDP offers no benefit over VXLAN for NAT-crossing scenarios: both require control-plane coordination, both have similar overhead, but only VXLAN has full hardware acceleration. **For any path that crosses a NAT boundary, VXLAN with rXg control-plane coordination is the recommended encapsulation.**

3.2 VXLAN: Layer 2 Overlay with UDP Transport

VXLAN (RFC 7348, August 2014) encapsulates entire Ethernet frames inside a UDP overlay. The encapsulation adds an outer IP header (20 bytes), an outer UDP header (8 bytes), and the VXLAN header (8 bytes), producing 36 bytes of encapsulation overhead—or 50 bytes of total wire overhead when the inner Ethernet header (14 bytes) is included. A 24-bit VNI field supports up to 16,777,216 distinct overlay segments, and the UDP source port is flow-hashed per RFC 7348 Section 5, enabling ECMP across spine-layer paths.

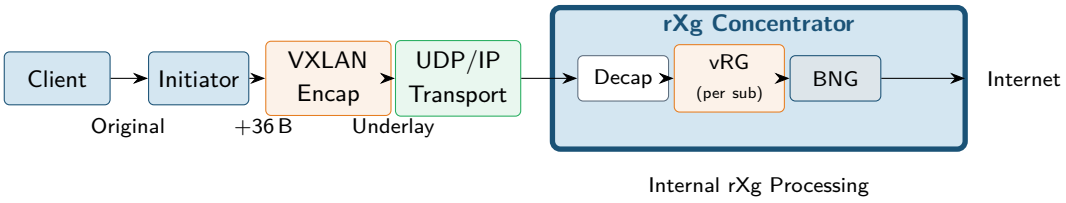

Figure 3: VXLAN Encapsulation Flow: Initiator to Concentrator

Strengths for SDAN. VXLAN's UDP encapsulation on destination port 4789 allows NAT devices to create session mappings—multiple overlays can share a single public IP via port translation. Reliable NAT traversal requires control-plane coordination to discover translated endpoints and

maintain bindings; the rXg platform provides this natively. This capability is critical for deployments crossing carrier or internet boundaries. The flow-hashed UDP source port enables Receive Side Scaling (RSS) on multi-core concentrators, distributing overlay traffic across CPU cores and scaling past GRE's single-core ceiling. The 24-bit VNI field provides native multi-service segmentation: a single VTEP simultaneously carries resident broadband, guest networks, building systems, and staff operations via independent VNIs, each with isolated Media Access Control (MAC) learning. BGP EVPN (RFC 7432, RFC 8365) provides dynamic endpoint learning, MAC mobility tracking, and centralized provisioning as the overlay signaling protocol.

Trade-offs. VXLAN's 50-byte overhead reduces the overlay MTU to 1450 bytes on a 1500-byte underlay, implying a TCP MSS of $\sim$1410 bytes—a $\sim$3.4% payload capacity reduction per segment. VXLAN's encapsulation path requires more firmware complexity in the initiator: UDP header construction, source port hash calculation, VXLAN header construction with VNI lookup, and outer Ethernet framing. On single-core software-only initiators, this produces measurably higher CPU utilization than GRE. East-west traffic between VTEPs on the underlay requires VTEP-capable switches at each participating node, adding hardware cost if the underlay does not already include such switches. Broadcast, Unknown-unicast, and Multicast (BUM) traffic is replicated to all VTEPs in the same VNI, which can become significant in large fabrics until MAC routes are learned.

3.3 Side-by-Side Comparison

Aspect	L2oGRE (RFC 2784/2890)	VXLAN (RFC 7348)
Layer	Layer 2 (Ethernet frames)	Layer 2 (Ethernet frames)
Encap Overhead	32 bytes (IP + GRE + key + inner VLAN)	36 bytes (IP + UDP + VXLAN header)
Wire Overhead	46 bytes (encap + inner Ethernet)	50 bytes (encap + inner Ethernet)
Topology	Point-to-point (explicit)	Multipoint (dynamic)
MTU Impact	1500 → 1546 needed	1500 → 1550 needed
TCP MSS	∼1414 (∼3.2% reduction)	∼1410 (∼3.4% reduction)
Segmentation	Single GRE overlay with inner 802.1Q VLAN multiplexing	VNI per resident/service (no inner VLAN)
NAT Traversal	IP Protocol 47; NAT-hostile	UDP 4789; NAT-traversable with rXg control-plane coordination
ECMP	No entropy fields	UDP source port entropy
Firmware Complexity	Low-Moderate (IP + GRE + key + inner VLAN tag)	Moderate-High (UDP + 5-tuple entropy hash + VNI)
SW Compute (no offload)	Lower (∼38% less CPU; better single-core)	Higher (extra UDP/VNI/entropy per packet; better multi-core via RSS)
HW Acceleration	Available (Network Interface Cards (NICs), DPUs; limited switch-ASIC)	Broad (NICs, DPUs, switch-ASICs; de facto overlay target)
Control Plane	Static provisioning + VLAN mapping	BGP EVPN (on rXg)
Maturity	RFC 2784/2890 (2000); 25+ years	RFC 7348 (2014); broad production

Table 1: GRE vs VXLAN Protocol Comparison

4 The Decision: Topology Drives Protocol Choice

The encapsulation choice is shaped by the topology of the path between initiator and concentrator, the capabilities of edge hardware, and the operational requirements of the deployment. The following questions identify the primary decision drivers.

Important

Primary Decision Rule:

1. Does the path between initiator and concentrator cross a NAT, CGNAT, or firewall boundary that blocks IP Protocol 47?
2. If yes: **use VXLAN**. UDP port 4789 allows NAT session creation; the rXg control plane provides the endpoint discovery and binding maintenance needed for reliable operation.
3. If no (private, directly routable IP path): **use GRE**. Lower overhead, simpler firmware, suits PON/DOCSIS underlays.

4.1 The Topology-First Decision Flowchart

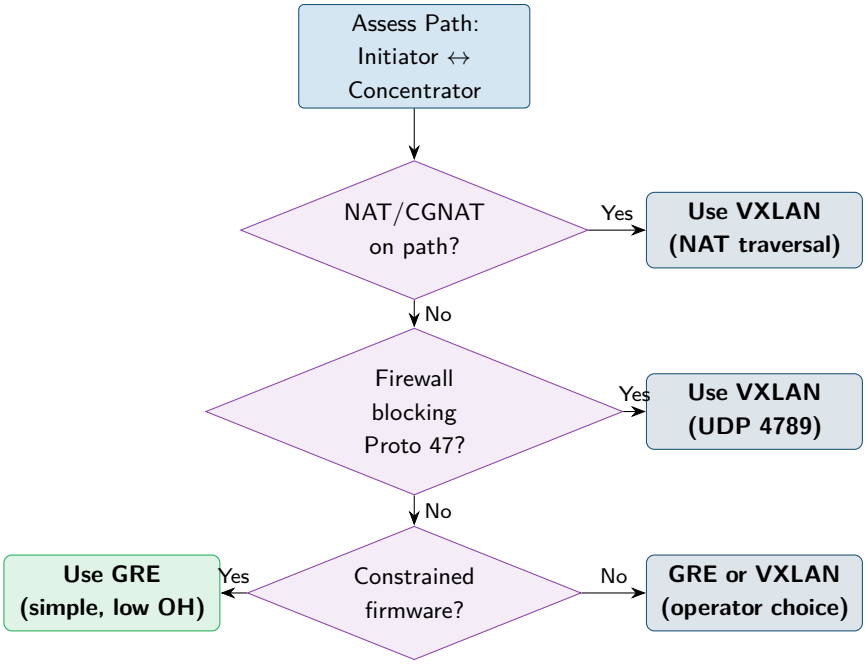

Figure 4: Topology-First Protocol Decision Flowchart

The flowchart captures the primary decision logic. NAT presence is the first and most decisive filter—it effectively eliminates GRE as a practical option for that path. Firewall policy is the second filter, since many enterprise and carrier firewalls block IP Protocol 47 by default but permit UDP 4789. Only when the path is private, directly routable, and firewall-permissive does the secondary question of firmware complexity arise.

4.2 GRE Strengths: When the Path Is Private

When the initiator-to-concentrator path provides direct IP reachability with no NAT, GRE delivers five concrete advantages.

Slightly lower overhead. L2oGRE's 32-byte encapsulation (outer IP + GRE base + key + inner VLAN tag) yields an inner IP MTU of 1454 bytes and TCP MSS of ∼1414 bytes on a 1500-byte underlay—a ∼3.2% payload reduction, only 4 bytes better than VXLAN's ∼3.4%. On a mini-jumbo underlay of 1546 bytes or above, the L2oGRE penalty disappears entirely; a 1550-byte underlay eliminates it for VXLAN as well.

Firmware simplicity. L2oGRE's encapsulation path—insert an inner VLAN tag, prepend a GRE header (with key), and wrap in an outer IP header—is straightforward to implement. For embedded platforms with limited code space and no operating system (bare-metal ONT firmware, cable modem DOCSIS stacks), this simplicity translates directly to faster development, smaller attack surface, and fewer firmware defects. The rXg concentrator handles all complex functions; the initiator's firmware burden is deliberately minimal.

PON and DOCSIS suitability. PON and DOCSIS underlays are inherently point-to-point or point-to- multipoint with direct L2/L3 adjacency between subscriber equipment and headend. There is no NAT in the path, no intermediate firewall, and no need for ECMP (the transport is a single fiber or coax path). GRE maps naturally to this topology. The underlay is commodity L2-oriented transport; GRE adds the minimum necessary encapsulation to carry subscriber traffic to the concentrator.

Debugging simplicity. GRE's outer header contains only source IP, destination IP, and protocol number. A packet capture on the underlay shows the encapsulation structure immediately. There is no UDP layer, no VXLAN header, and no VNI to decode. For operators troubleshooting overlay failures on carrier networks where they may not control the underlay infrastructure, this transparency is operationally valuable.

Lower software-only compute overhead. On CPU-only concentrators or initiators without DPU hardware offload, GRE's simpler encapsulation path consumes ∼38% less CPU than VXLAN at equivalent throughput. For private-path deployments where hardware acceleration is not available or not yet deployed, this translates directly to higher single-core headroom and lower per-packet processing cost.

4.3 VXLAN Strengths: When the Path Crosses Boundaries

When the path between initiator and concentrator crosses a NAT boundary, traverses the internet, or spans distributed sites, VXLAN delivers five advantages that GRE cannot match.

NAT and firewall traversal. VXLAN's UDP encapsulation on destination port 4789 is handled by standard UDP session tracking in mainstream NAT implementations. Multiple VXLAN overlays from different initiators can share a single public IP via port translation. Firewalls treat UDP 4789 as a standard application port. By contrast, GRE's IP Protocol 47 has no port numbers for session multiplexing—NAT devices typically track at most one GRE session per public IP per destination, and many enterprise firewalls block Protocol 47 by default.

ECMP and multi-core scaling. VXLAN's flow-hashed UDP source port provides per-flow entropy that enables ECMP load balancing across spine paths and RSS distribution across CPU cores on the concentrator. On multi-core rXg platforms, this allows VXLAN to scale aggregate throughput past GRE's single-core ceiling. For deployments with multiple spine paths between initiators and concentrator, VXLAN distributes load; GRE sends all traffic from a given overlay over an identical path.

Multi-service segmentation. VXLAN's 24-bit VNI field allows a single VTEP to carry up to 16 million distinct overlay segments. In an SDAN deployment, this means one VTEP simultaneously carries resident broadband (VNI 4100), guest Wi-Fi (VNI 4101), cameras (VNI 4102), access control (VNI 4103), and staff operations (VNI 4104)—five services through one overlay, each with hardware-enforced isolation. Under L2oGRE, services are multiplexed with inner 802.1Q VLAN tags on the per-initiator overlay, but each initiator still requires explicit per-service VLAN mapping configuration.

East-west traffic on the underlay. In the standard SDAN model, stateless initiators hairpin all inter-initiator traffic through the concentrator regardless of encapsulation protocol. However, in campus deployments where managed switches with VTEP capability are present, VXLAN enables east-west forwarding between those VTEPs without routing through the concentrator—reducing latency for local traffic at the cost of additional switch hardware. GRE's point-to-point model has no equivalent mechanism; inter-initiator traffic traverses the concentrator.

DPU-accelerated throughput. With NVIDIA BlueField DPUs in eSwitch mode, the rXg platform achieves 100 Gbps aggregate VXLAN throughput with full hardware offload—encap, decap, header modification, and field copy operations all handled in silicon. GRE also receives DPU offload for basic encap/decap, so raw tunnel throughput is comparable for both protocols. The gap appears when the full processing pipeline—header rewrite, field copy, policy marking—is required: GRE lacks hardware-accelerated modify and copy operations, so those stages fall back to the CPU data path. VXLAN remains the clear choice for deployments that depend on the full DPU pipeline at scale.

4.4 NAT/CGNAT Traversal: The Decisive Factor

NAT traversal capability is the single most decisive factor for most SDAN deployments. The following table summarizes the practical differences.

Factor	GRE (Protocol 47)	VXLAN (UDP 4789)
NAT compatibility	No port numbers for session tracking; NAT-hostile	UDP port translation allows NAT session creation; return-path discovery and binding maintenance require rXg control-plane coordination
Multiple overlays/IP	Typically one session per IP per destination	Multiple sessions via UDP port translation
CGNAT	Requires carrier cooperation (no practical workaround)	UDP session tracking with rXg keepalives to maintain bindings
Firewall rules	"Allow Protocol 47"—uncommon default	"Allow UDP 4789"—standard
Carrier support	Varies; some carriers filter Proto 47	Widely supported across standard IP transport

Table 2: NAT/CGNAT Traversal Comparison

Important

NAT Decision Rule: Any deployment where the overlay path crosses a NAT'd carrier boundary strongly favors VXLAN. GRE should only be considered where the carrier path provides direct IP reachability with no NAT/CGNAT. GRE has no practical NAT traversal option with hardware acceleration support.

4.5 Decision Matrix

The following matrix maps topology characteristics to protocol recommendations. The topology drives the choice; the venue type is incidental.

Path Characteristic	NAT?	Protocol	Confidence
Private IP, PON/DOCSIS underlay	No	GRE	High
Private IP, campus Ethernet	No	GRE or VXLAN	High
Private IP, carrier metro Ethernet	No	GRE	High
CGNAT on carrier path	Yes	VXLAN	Very High
Internet/cloud path	Yes	VXLAN	Very High
Firewall blocking Proto 47	N/A	VXLAN	Very High
Mixed portfolio (some NAT, some not)	Both	Both	High

Table 3: Topology-to-Protocol Decision Matrix

5 Deployment Examples

The following examples illustrate how the topology-first framework applies to real-world SDAN deployments. The venue types are *examples*, not drivers—the protocol choice follows from the path topology in each case, not from the fact that the venue is a stadium, an MDU, or a hotel.

5.1 On-Premises Concentrator with Private Path

Typical venues: large public venue (LPV) stadiums, campus MDU, convention centers

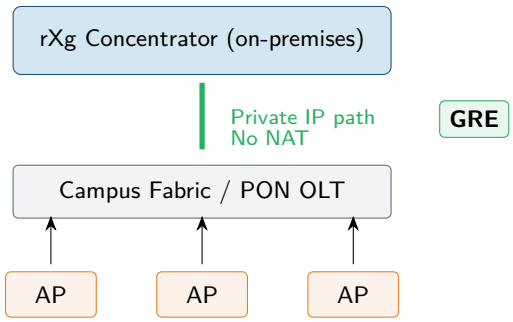

Figure 5: On-Premises Concentrator: Private Path Favors GRE

When the rXg concentrator sits on-premises—in the venue's main distribution frame (MDF), a property headend, or a campus data center—and connects to initiators through a private campus fabric or PON infrastructure, the path is direct and NAT-free. GRE is the natural choice. The initiators (APs or ONTs) encapsulate subscriber traffic with 32 bytes of L2oGRE overhead, the campus fabric or PON transport delivers encapsulated packets to the concentrator, and the rXg decapsulates and routes traffic through the appropriate vRG.

For venues with spine-leaf fabrics where ECMP across spine links is operationally important (e.g., large stadiums with thousands of APs), VXLAN's UDP source port entropy provides better load distribution. In these cases, the topology still permits GRE, but VXLAN's ECMP advantage may justify the additional overhead. The decision is situational, not categorical.

5.2 Central Office (CO)-Hosted Concentrator with Private Carrier Path

Typical venues: MDU portfolios, City Hot Zone (CHZ) districts, DOCSIS plants

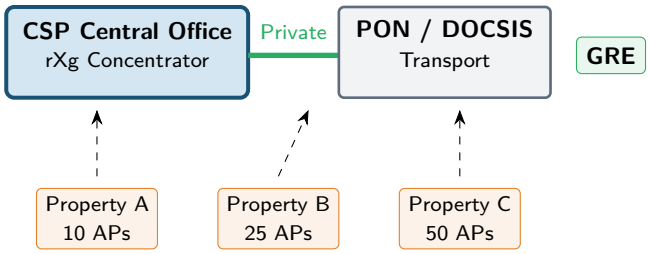

Figure 6: Central Office (CO)-Hosted Concentrator: Private Carrier Path Favors GRE

When the rXg concentrator is hosted in the CSP's central office and connects to property initiators through the CSP's own PON or DOCSIS infrastructure, the path remains private. The CSP controls both ends and the transport between them. There is no NAT, no firewall boundary, and no internet transit. GRE is the right choice.

This is the most common topology for MDU portfolio deployments where the CSP owns the fiber or coax plant. Each property's APs encapsulate subscriber traffic via GRE to the CO-hosted rXg, which hosts vRG instances for every unit across the portfolio. The configuration is minimal: each AP is provisioned with the concentrator's IP address and a GRE key. Fleet Manager orchestrates provisioning across all properties from a single pane of glass.

For CHZ deployments spanning dozens of buildings across an urban district on the same carrier infrastructure, the same logic applies. If the carrier provides direct IP reachability between every building's APs and the CO-hosted concentrator with no NAT, GRE's lower overhead and simpler firmware are preferred. If the carrier introduces CGNAT on any segment, that segment must use VXLAN.

5.3 Distributed/Remote Sites with Internet Path

Typical venues: out-of-region CHZ, remote campuses, franchise sites

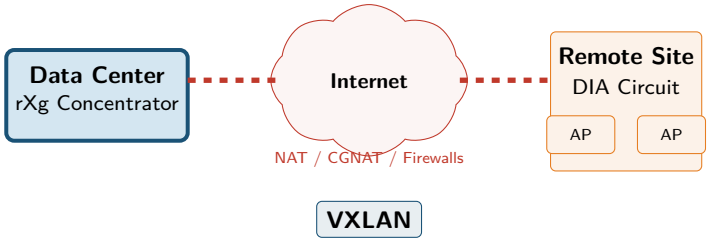

Figure 7: Distributed Remote Sites: Internet Path Requires VXLAN

When remote sites connect to the concentrator through the public internet, the path typically crosses NAT or CGNAT boundaries. The remote site's DIA circuit terminates behind a NAT device; the concentrator sits behind a data center firewall. GRE's Protocol 47 cannot reliably traverse this path. VXLAN's UDP 4789 can traverse this path with rXg control-plane coordination.

This topology is common for operators managing out-of-region properties, franchise networks, or campus satellite buildings that cannot be reached through the operator's own transport infrastructure. The rXg concentrator in the operator's data center terminates VXLAN overlays from all remote sites, each site's APs identified by their VTEP source address.

5.4 Hospitality with DIA and Remote Concentrator

Typical venues: hotel chains, resort properties, cruise ships

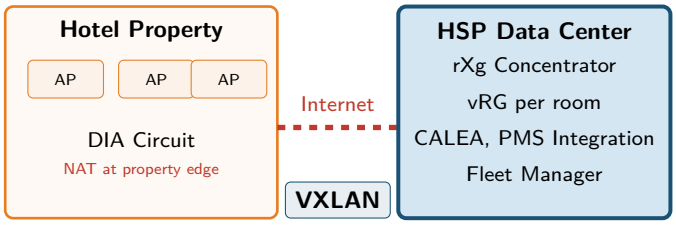

Figure 8: Hospitality with DIA: Remote Concentrator Requires VXLAN

Hospitality Service Providers (HSPs) typically operate a centralized data center that hosts rXg concentrators for dozens or hundreds of hotel properties. Each property connects to the HSP's data center through a DIA circuit and the public internet. The property's edge router performs NAT. The HSP's data center sits behind its own firewall. The path is NAT-traversing in practice.

VXLAN is the clear practical choice. Each hotel property's APs establish VXLAN overlays to the HSP's rXg concentrator. The concentrator hosts per-room vRG instances, integrates with the property management system (PMS) for guest identity, enforces CALEA compliance, and provides property-wide service segmentation (guest Wi-Fi, staff network, building systems, Point of Sale) through independent VNIs.

The DPU-accelerated rXg concentrator terminates VXLAN overlays from all properties simultaneously. Adding a new hotel to the portfolio requires provisioning the property's APs with the concentrator's public IP and deploying VXLAN overlay configuration via Fleet Manager—no site visit to the data center, no carrier coordination, no GRE-in-UDP workarounds.

5.5 Mixed-Protocol Portfolio

Typical operators: regional CSPs, national HSPs, university systems

Most operators manage a portfolio that includes both private-path sites and NAT-traversing sites. The topology-first framework produces a natural mixed-protocol deployment:

Site Type	Path	Protocol	Rationale
MDU on operator's PON	Private	GRE	Direct L2 path, no NAT
MDU on third-party fiber	CGNAT	VXLAN	Carrier NAT in path
Campus building on-prem	Private	GRE	Local fabric, no NAT
Remote satellite campus	Internet	VXLAN	NAT at both ends
Hotel on HSP DIA	Internet	VXLAN	NAT at property edge
CHZ on operator metro-E	Private	GRE	Carrier provides L2
CHZ on leased circuit with NAT	CGNAT	VXLAN	Carrier NAT in path

Table 4: Mixed-Protocol Portfolio: Topology Determines Each Site's Protocol

The rXg platform supports this natively. A single Fleet Manager instance orchestrates GRE-connected sites and VXLAN-connected sites simultaneously. The concentrator terminates both overlay types and routes traffic through the same vRG instances regardless of encapsulation. The protocol is transparent to the subscriber experience and to the operator's policy engine.

6 Quantified Trade-Offs

The topology-first framework determines *which* protocol to use. The quantified trade-offs in this section inform *how much* each choice costs in overhead, throughput, and operational complexity.

6.1 MTU and MSS Impact

Protocol	Wire OH	Tunnel MTU	TCP MSS	Payload Reduction
L2oGRE	46 B	1454 B	∼1414 B	∼3.2%
VXLAN	50 B	1450 B	∼1410 B	∼3.4%

Table 5: MTU and MSS Impact on 1500-Byte Underlay

The ∼0.3% difference in payload capacity between the two protocols (4 bytes per packet, or ∼3 Mbps at 1 Gbps) is operationally negligible at residential broadband speeds—well within normal throughput variance. The difference becomes relevant only in high-throughput, latency-sensitive applications where every byte of MSS matters, or in environments where the 1500-byte MTU ceiling is immovable and MSS clamping is undesirable.

6.2 Mini-Jumbo Frames: Eliminating the MTU Penalty

When the operator controls the underlay infrastructure, increasing the underlay MTU eliminates the overhead penalty entirely.

Protocol	Overhead	Min. Underlay MTU for 1500-Byte Inner	Result
L2oGRE (RFC 2784)	46 bytes	**1546 bytes**	Full 1500-byte inner MTU
VXLAN (RFC 7348)	50 bytes	**1550 bytes**	Full 1500-byte inner MTU

Table 6: Mini-Jumbo MTU Requirements

Frames slightly larger than the IEEE 802.3 maximum (1518 bytes) but smaller than full jumbo frames (9000+ bytes) are known as **baby giant** or **mini-jumbo** frames. This capability is widely available across commodity switching silicon—including Cisco Catalyst/Nexus 9000 and Juniper QFX/EX, among other major enterprise and carrier platforms. Most modern Gigabit PON (GPON) / 10-Gigabit Symmetric PON (XGS-PON) OLTs also support MTU values in the 1546–1600 range.

Feasibility depends on underlay ownership. Where the CSP controls building switches and fiber (on-premises MDU), feasibility is high. Where the CSP owns PON infrastructure (CO-hosted), feasibility is medium-high but requires OLT/ONT MTU verification. For carrier leased lines, it requires Service Level Agreement (SLA) renegotiation and is typically low.

6.3 Software-Only Performance

On pure software systems without hardware offload (Linux kernel, single core), GRE and VXLAN exhibit measurable performance differences:

Metric	GRE	VXLAN	Delta
Throughput (10 Gbps, single core)	9.5 Gbps (∼95%)	8.9 Gbps (∼89%)	GRE +6.7%
CPU utilization at 1 Gbps	∼5%	∼8%	GRE ∼38% lower
Added latency	+0.1 ms	+0.3 ms	GRE 0.2 ms lower

Table 7: Software-Only Performance Comparison (single core, AMD EPYC, Linux 6.1)

The difference arises because VXLAN requires additional outer-header processing: UDP header construction, source port hash calculation, VXLAN header construction/parsing, and outer Ethernet framing. On multi-core systems, however, VXLAN's UDP source port entropy enables RSS, distributing traffic across cores and potentially exceeding GRE's aggregate throughput.

For software-only edge initiators without DPU offload, GRE provides measurably better single-stream performance. For centralized concentrators with multi-core hardware, the difference diminishes or reverses.

6.4 DPU Hardware Acceleration

Modern Data Processing Units (DPUs) offload overlay encapsulation to hardware, fundamentally changing the throughput equation. VXLAN receives first-class hardware offload across all major DPU and switch-ASIC platforms. GRE offload is available for basic encap/decap operations but varies in feature depth by vendor—notably, the NVIDIA DOCA Flow API supports match, encap, decap, modify, and copy for VXLAN but only match, encap, and decap for GRE (no modify or copy in hardware). AMD Pensando lists GRE among supported flexible encapsulations, but VXLAN is the primary built-in pipeline feature with deeper integration.

Platform	VXLAN	GRE	Geneve	Notes
NVIDIA BlueField	Full (match, encap, decap, modify, copy)	Encap/decap only (no modify/copy)	Full	GRE lacks advanced DOCA Flow operations
AMD Pensando	Built-in pipeline	Supported via flexible encap pipeline	Built-in	GRE listed in Salina product brief
Intel IPU (E2100)	P4-programmable	P4-programmable	P4-programmable	All three require P4 implementation
Broadcom SmartNIC	Full stateless offload	Stateless offload	Full	Near-parity at NIC level

Table 8: DPU Overlay Protocol Support Matrix

Geneve (Generic Network Virtualization Encapsulation, RFC 8926) appears in this comparison because DPU vendors increasingly support it alongside VXLAN. Its extensible Type-Length-Value (TLV) header model can carry richer metadata than either VXLAN or GRE. The rXg platform currently deploys VXLAN and GRE for SDAN overlays; Geneve is tracked as a future option as ecosystem support broadens.

When encapsulation is handled in hardware, the software throughput gap diminishes to near zero. Published benchmarks include NVIDIA BlueField-2 achieving 96 Gbps software-defined networking (SDN) throughput with near-zero CPU utilization, and BlueField DPU delivering ∼54x higher packet-switching throughput per server (18.7M packets per second (PPS) vs. 350K PPS) compared to software-only OVS.

Note

RG Nets rXg Platform: The rXg platform uses NVIDIA BlueField-3 DPUs with FD.io Vector Packet Processing (VPP), supporting hardware- accelerated VXLAN and GRE on BlueField-3 SKUs (100/200/400 Gigabit Ethernet configurations). With DPU offload, basic encap/decap throughput is comparable for both protocols; the protocol choice is driven primarily by topology, NAT requirements, and whether the deployment needs full-pipeline hardware offload for header modification and field copy operations (VXLAN has broader DPU feature coverage than GRE).

6.5 Configuration Efficiency at Scale

In multi-service deployments, the configuration model differs significantly between the two protocols.

L2oGRE: single overlay with per-service VLAN mapping. L2oGRE uses one overlay per initiator-concentrator pair, multiplexing services with inner 802.1Q VLAN tags. Each property requires overlay endpoint configuration plus per-service VLAN mapping:

```
interface Tunnel1
  tunnel source 203.0.113.50
```

```
  tunnel destination 203.0.113.200
  tunnel key 1001

service_vlan_map: [
  {vlan: 100, name: "resident"},
  {vlan: 101, name: "guest"},
  {vlan: 102, name: "cameras"},
  {vlan: 103, name: "locks"},
  {vlan: 104, name: "staff"}
]
```

Listing 1: L2oGRE Single-Overlay Configuration (per property; illustrative)

VXLAN: VNI array on single overlay. VXLAN uses a single VTEP with a VNI array, where each VNI natively represents a service segment without inner VLAN tags:

```
interface vxlan0
  tunnel source 203.0.113.50
  tunnel destination 203.0.113.200

vni_list: [
  {vlan: 100, vni: 4100, name: "resident"},
  {vlan: 101, vni: 4101, name: "guest"},
  {vlan: 102, vni: 4102, name: "cameras"},
  {vlan: 103, vni: 4103, name: "locks"},
  {vlan: 104, vni: 4104, name: "staff"}
]
```

Listing 2: VXLAN Multi-Service Configuration (per property; illustrative)

Both models require per-service definitions at each property. The operational difference is in fleet-wide management: VXLAN's VNI definitions are propagated through the BGP EVPN control plane, while L2oGRE VLAN mappings are statically provisioned per device. Adding a sixth service under L2oGRE requires a new VLAN mapping entry per property; under VXLAN, a single VNI definition propagates fleet-wide through the control plane.

7 Implementation Guidance

Successful deployment requires three pre-design steps regardless of which protocol is selected.

7.1 Step 1: Underlay MTU Survey

Before selecting a protocol, survey every carrier path in the deployment portfolio for MTU support:

- **Controlled underlay** (operator-owned fiber/coax): verify switch and OLT/ONT MTU support. Most modern PON equipment supports 1546–1600 byte MTU. If confirmed, mini-jumbo frames eliminate the overhead penalty for either protocol.
- **Carrier leased lines:** request MTU confirmation from the carrier. If the carrier guarantees only 1500 bytes, plan for MSS clamping (L2oGRE: 1414, VXLAN: 1410).
- **Internet paths:** assume 1500-byte MTU. MSS clamping is required. Path MTU Discovery (PMTUD) may be unreliable across the internet.

7.2 Step 2: NAT/CGNAT Audit

For every initiator-to-concentrator path, determine whether NAT or CGNAT is present:

- **Operator-owned transport:** typically no NAT. Verify with traceroute and protocol-47 reachability tests.
- **Carrier transport:** request explicit confirmation that IP Protocol 47 is passed. Many carriers silently filter it. If the carrier introduces CGNAT on any segment, that segment must use VXLAN.
- **Internet/DIA paths:** assume NAT is present. Use VXLAN.

> **Tip**
>
> **Practical Test:** Test IP Protocol 47 reachability from the initiator's location to the concentrator's IP address. A successful result confirms GRE is permitted on that path; a failure indicates the path should be treated as VXLAN-only unless further carrier validation proves otherwise.

7.3 Step 3: Phased Rollout

Phase 1: Pilot (Months 1–3). Select one private-path site and one NAT-traversing site for pilot. Deploy GRE on the private-path site, VXLAN on the NAT-traversing site. Validate MTU

assumptions, overlay stability, DPU offload performance, and Fleet Manager orchestration across both protocols. Benchmark throughput on target hardware.

Phase 2: Portfolio Classification (Months 2–4). Complete the MTU survey and NAT/CGNAT audit across the entire portfolio. Classify each site as GRE-eligible or VXLAN-required based on path topology. Identify any sites where the path is ambiguous (e.g., carrier may introduce CGNAT in the future) and default them to VXLAN.

Phase 3: Fleet Rollout (Months 4+). Deploy across the portfolio site-by-site. GRE sites and VXLAN sites are configured through the same Fleet Manager workflow. The rXg concentrator terminates both overlay types simultaneously. Monitor for carrier path changes (e.g., a carrier introducing CGNAT) that would require migrating a site from GRE to VXLAN.

8 Conclusion

The choice between VXLAN and GRE for SDAN deployments is not a question of which protocol is architecturally superior. It is a question of topology.

When the path between the SDAN initiator and the SDAN concentrator is private and directly routable, L2oGRE is the preferred choice. Its 32-byte encapsulation overhead (only 4 bytes less than VXLAN) preserves slightly more payload capacity. It is simpler to implement in constrained embedded firmware and maps naturally to the point-to-point topology of PON and DOCSIS underlays. On these paths, GRE's simplicity is not a limitation—it is a strength.

When the path crosses a NAT boundary, traverses the public internet, or spans geographically distributed sites through firewalls, VXLAN is the right choice. Its UDP-based transport allows NAT session creation (with rXg control-plane coordination for endpoint discovery and binding maintenance), its flow-hashed source port enables multi-core scaling and ECMP load balancing, and its VNI-based segmentation provides multi-service isolation without per-service overlay provisioning. On these paths, VXLAN's additional overhead is the cost of traversing boundaries that GRE cannot cross.

The rXg platform supports both protocols natively. A single concentrator terminates GRE overlays from private-path sites and VXLAN overlays from NAT-traversing sites simultaneously, routing all traffic through the same vRG instances and applying the same policies regardless of encapsulation. With DPU hardware acceleration on BlueField-3 SKUs, both protocols achieve near-line-rate encap/decap throughput—though VXLAN benefits from deeper offload support—and the decision is driven by topology, NAT requirements, and operational scale rather than raw throughput.

Three pre-design steps ensure a sound deployment regardless of protocol: an underlay MTU survey to determine whether mini-jumbo frames can eliminate the overhead penalty, a NAT/CGNAT audit to classify each site's path topology, and a phased rollout that validates both protocols at pilot scale before fleet deployment. With these in place, operators can match each site in their portfolio to the right encapsulation with confidence—deploying GRE where the path is private and simple, VXLAN where the path crosses boundaries, and both where the portfolio demands it.

9 References

IETF Standards

- RFC 2784: Generic Routing Encapsulation (GRE), IETF, March 2000
- RFC 2890: Key and Sequence Number Extensions to GRE, IETF, September 2000
- RFC 7348: Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks, IETF, August 2014
- RFC 7432: BGP MPLS-Based Ethernet VPN, IETF, February 2015
- RFC 8086: GRE-in-UDP Encapsulation, IETF, March 2017
- RFC 8365: A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN), IETF, March 2018
- RFC 8926: Geneve: Generic Network Virtualization Encapsulation, IETF, November 2020

Vendor Documentation

- Cisco Systems: Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 10.6(x), "MTU Size in the Transport Network." cisco.com
- Cisco Systems: Document ID 25885, "Resolve IPv4 Fragmentation, MTU, MSS, and PMTUD Issues with GRE and IPsec." cisco.com
- Cisco Systems: "Configure Jumbo/Giant Frame Support on Catalyst Switches," Document ID 24048. cisco.com
- Cisco Systems: VXLAN EVPN Multi-Site Design and Deployment White Paper (Nexus 9000). cisco.com
- Cisco Systems: VXLAN Network with MP-BGP EVPN Control Plane Design Guide. cisco.com
- Juniper Networks: "Understanding VXLANs," Junos OS Documentation. juniper.net
- Juniper Networks: "Configuring GRE Tunnel Interfaces," Junos OS Documentation. juniper.net
- NVIDIA: BlueField-3 Networking Platform User Guide—ASAP2 overlay offload (VXLAN, GRE, Geneve). docs.nvidia.com
- NVIDIA: "OVS Offload Using ASAP2 Direct." docs.nvidia.com
- NVIDIA Developer Blog: "Accelerate Enterprise Apps with Microsoft Azure Stack HCI and NVIDIA BlueField DPUs," November 2022
- NVIDIA Developer Blog: "Accelerating Cloud-Ready Infrastructure with Red Hat OpenShift and NVIDIA BlueField DPU," April 2022
- NVIDIA: DOCA Flow Programming Guide, v2.x. docs.nvidia.com/doca/sdk/doca-flow
- AMD: Pensando Salina DPU Product Brief. amd.com/en/products/accelerators/pensando

- Intel: Infrastructure Processing Unit (IPU) E2100 Product Brief. intel.com
- Broadcom: NetXtreme-E SmartNIC Product Brief—Virtualization and Tunnel Offload. broadcom.com

Third-Party Benchmarks

- Virtuasys (VIRTUA SYSTEMS SAS): Overlay Throughput Benchmarks—GRE, VXLAN, IPsec, WireGuard. AMD EPYC 7543, Linux 6.1, 10G fiber. Published at virtuasys.eu

Industry References

- IEEE: "Software-Defined Access Network (SDAN)," IEEE Conference Publication, 2014. DOI: 10.1109/CCNC.2014.6994134
- IEEE: "Software-defined access networks," IEEE Communications Magazine, 2014. DOI: 10.1109/MCOM.2014.6894466
- Kerpez, K., Ginis, G., et al.: "Software-Defined Access Networks." ResearchGate, 2014

RG Nets

- RG Nets SDAN Architecture Documentation
- RG Nets BlueField-3 DPU GA Announcement, PR Newswire, March 7, 2025
- RG Nets Fleet Manager Documentation

Whitepaper Prepared By: RG Nets Engineering
Date: April 2, 2026
Classification: Technical Architecture Guide
Version: v5.0