

Quantum IPoE BNG — Community Edition

Alpha Release 1.2.0 — Release Notes

The first customer-visible release of the Quantum IPoE BNG — a self-contained Broadband Network Gateway for small and mid-size ISPs. One tarball, one installer, and roughly three minutes to a working control plane, data plane, and web UI on a fresh Ubuntu 24.04 box.

Document Type	BlueDot Release Notes
Product	Quantum IPoE BNG — Community Edition
Release	Alpha 1.2.0
Bundle	<code>quantumtouch-installer-1.2.0-01.tar.gz</code>

Status	Alpha — lab-ready; not yet ISP-scale production
Release Date	July 24, 2026
Classification	BlueDot Proprietary

Contents

What this is 4

What ships in this release 5

System requirements 7

NIC compatibility 9

Known limitations 12

Upgrade / downgrade 14

Feedback 15

Appendix A: Quantum Controller v1.2.0 release notes 16

20 — onboarding performance under storm (identity/xid dedup) 19

21 — duplicate sessions on re-DORA 19

22 — DHCP reply-path counters / session-cap observability 19

What this is

The first customer-visible release of the **Quantum IPoE BNG** — a self-contained Broadband Network Gateway (BNG) for small and mid-size ISPs. One tarball, one `sudo ./install.sh`, and roughly three minutes later you have a working control plane, data plane, and web UI on a fresh Ubuntu 24.04 box.

This is the most feature-rich, turnkey IPoE solution in the Community Edition segment. It is ready for lab environments; the data plane is line rate. It is not yet suitable for ISP-scale (>100 subscribers/sec churn) production — that is targeted for the Beta release.

The three planes:

Plane	Component	What it does
Management	QuantumTouch	Web UI on port 8443 — subscribers, plans, portal config, health, audit log
Control	Quantum Controller	Source of truth. DHCP, AAA, session state. Runs in a container on port 3000 (HTTPS, internal).
Data	VPP + tuned DPDK	Line-rate packet forwarding. Per-core worker threads, isolated CPUs, hugepages, NUMA-pinned.

The controller team ships their own release notes for the Quantum Controller v1.2.0; those notes are embedded verbatim as [Appendix A](#) at the end of this document.

What ships in this release

Setup Wizard (end-to-end first-run)

Fresh install → first-run URL in the install banner → 7-step wizard:

1. **Platform Settings** — hostname, timezone, admin email, NTP
2. **DPDK Bind** — drag NICs from Kernel to VPP column. Mgmt-port lock refuses to bind any PCI function on the same slot as the management interface (protects your SSH session).
3. **CPU Pinning** — auto-computed on mount using the NUMA node of the NICs you just bound; the data plane never lands on a different socket than its NICs.
4. **Data Plane** — pick uplink and access interfaces from the VPP-bound set.
5. **Plans / IP Pools / Authentication / Captive Portal / TLS**
6. **Review & Apply** — one click, all sections applied. A clean run lands 7/7 green.

Management plane (QuantumTouch)

- Dashboard, Topology, Sessions, Health, Audit (with rev_id chips, deep-link from events to point-in-time timeline, and "return to live"), Hardware
- Subscribers, AAA (Local), Pools & DNS, Subscriber Plans, Portal Config, Vouchers, Edge Routers, Dataplane
- API Keys management (mint, revoke, roles, scoping)
- Snapshots (settings + state, hash-verified, MERGE + REPLACE restore modes)
- Two-layer security:
- **Machine-to-machine** — the bootstrap admin API key, plus per-service keys the shim and captive portal use to reach the controller over authenticated HTTPS
- **User login** — QuantumTouch's own auth (password policy, optional MFA, session management, audit log, LDAP + passkey/WebAuthn + SAML supported)

Control plane (Quantum Controller)

See [Appendix A](#) for the full controller team notes. Headline:

- **Local AAA** — no RADIUS server required. Two modes: `open` (accept-all) and `gated` (per-line policy keyed on DHCP Option 82 identity).
- **1,000 concurrent-subscriber cap** in the VPP plugin (tamper-resistant, plugin-side).
- **DHCP lease-expiry timers** independent of session/idle timeouts.
- **Plan-edit propagation** — `PUT /bng/plans/{name}` re-rates every live session on that plan in place.
- **Auth required by default** — Bearer API keys, hashed at rest, role-based (`read_only` / `read_write` / `admin_only`).

Data plane (VPP)

- Line-rate forwarding on tested NICs (see [NIC compatibility matrix](#) below)

- **NUMA-aware CPU pinning** — auto-fill picks main + workers from the same socket as the bound NICs. Cross-NUMA data-plane pinning is a well-known throughput killer; the wizard now steers you away from it by default.
- Per-core worker threads on isolated CPUs (GRUB `isolcpus` set from wizard's selections)
- 1 GB hugepages (default $8 \times 1\text{G}$)

Captive portal

- Local mode — `bng-portal.service` on `:8088/`, ships enabled and Bearer-authenticated to the shim
- Voucher-based redemption supported
- External-URL redirect mode requires you to host your own portal (see [Known limitations](#))

Observability

- Prometheus TSDB (14-day default retention) scraping `/metrics` every 5s
- Dashboards for sessions, DHCP reply-path, session-cap headroom
- Sticky event history with RESOLVED badge, start/end/duration
- CPE fleet simulator via `quantumtest` (harness for load and churn testing; separate repo)

System requirements

We break these into **Minimum** (a working lab / small pilot / dev box) and **Recommended** (production pilot).

Minimum

CPU	4 physical cores, x86_64
RAM	8 GB
Storage	20 GB free on / after install
NICs	1 supported adapter (see matrix) + a separate management NIC
Hugepages	4 × 1G (allocated at boot via GRUB)
OS	Ubuntu 24.04 LTS (Noble), kernel 6.8+
BIOS	Intel VT-d / AMD IOMMU enabled

What Minimum gets you: enough to serve a small subscriber base (a few dozen active sessions), evaluate the UI, exercise the setup wizard, and test config flows. Not tuned for throughput.

Recommended

CPU	16+ physical cores, dual-socket, hyperthreading off
RAM	32 GB+
Storage	100 GB SSD on / (Prometheus TSDB + audit log grow over time)
NICs	2× supported adapters, same NUMA socket as the VPP worker cores
Hugepages	8 × 1G
OS	Ubuntu 24.04 LTS (Noble), kernel 6.8+
BIOS	Intel VT-d + Intel Turbo Boost off (for consistent worker latency)

Why hyperthreading off: VPP treats logical CPUs as physical for pinning; sibling hyperthreads on the same physical core fight for the same execution units and throughput drops. Turning HT off makes CPU numbering match physical cores 1:1 and lets the wizard's NUMA math be correct.

NIC compatibility

Tier 1 — tested, known-good

Cards in this list have been driven end-to-end through the wizard's DPDK Bind step, brought to line-rate, and exercised under CPE/traffic-generator load.

Vendor	Family	Speed	Kernel driver	DPDK PMD	Notes
Intel	X710	10 GbE	<code>i40e</code>	<code>net_i40e</code>	Widely deployed, mature
Intel	XL710	40 GbE (4×10G or 1×40G)	<code>i40e</code>	<code>net_i40e</code>	Primary dev target; 40G confirmed line-rate
Intel	XXV710	25 GbE	<code>i40e</code>	<code>net_i40e</code>	Same driver family as XL710
Intel	E810-C QSFP	100 GbE	<code>ice</code>	<code>net_ice</code>	Requires DDP package
Intel	E810-C SFP	25/100 GbE	<code>ice</code>	<code>net_ice</code>	Requires DDP package
Intel	E830 QSFP	200 GbE	<code>ice</code>	<code>net_ice</code>	Requires OEM-signed NVM 1.20+ (Lanner); DDP rev3 verified

Tier 2 — should work, not currently tested

Cards in this tier use the same driver families as Tier 1 cards, or ship supported DPDK PMDs, but we haven't personally validated them end-to-end. Report your experience to alpha@bluedot-insight.com.

Vendor	Family	Speed	Kernel driver	Notes
Intel	82599 (X520 / X540 / X550)	10 GbE	<code>ixgbe</code>	Mature DPDK driver; earlier generation than X710
Nvidia / Mellanox	ConnectX-5	25 / 40 / 100 GbE	<code>mlx5_core</code>	Requires MLNX_OFED + firmware you provide — see below
Nvidia / Mellanox	ConnectX-6 Dx / Lx	25 / 100 GbE	<code>mlx5_core</code>	Same as above
Nvidia / Mellanox	ConnectX-7	200 / 400 GbE	<code>mlx5_core</code>	Same as above
Broadcom	NetXtreme-E	25 / 50 / 100 GbE	<code>bnxt_en</code>	DPDK <code>net_bnxt</code> PMD available
Chelsio	T5 / T6	40 / 100 GbE	<code>cxgb4</code>	DPDK <code>net_cxgbe</code> PMD available
Marvell	OCTEON	25 / 100 GbE	vendor	Vendor DPDK PMD

Loading your own drivers / firmware (Nvidia ConnectX)

The installer does not bundle MLNX_OFED (Mellanox's driver + firmware stack) or Nvidia's firmware images. To use a ConnectX card:

1. **Install MLNX_OFED** for your kernel from [Nvidia's download center](#). Match your Ubuntu 24.04 + kernel 6.8 version exactly.
2. **Flash matching firmware** if your card is older than the OFED package expects. `mlxfwmanager --query` will tell you the current level; `mlxfwmanager --online -u -d <pci>` will update from Nvidia's catalog.
3. **Reboot** so the new driver + firmware load cleanly.
4. **Re-run the setup wizard's DPDK Bind step** — the ConnectX ports should now appear in the Kernel column, draggable to VPP. Mellanox NICs work best with **PCI passthrough** rather than `vfiopci` — VPP's DPDK PMD handles the `mlx5` bifurcated driver correctly, but you may need to leave the card on `mlx5_core` and let VPP grab it via `dpdk { dev <pci> }` in `startup.conf` rather than through the wizard's rebind flow. This is a real rough edge in this release — see [Known limitations](#).

If a card doesn't appear in either column after driver load, please open an issue with:

- `sudo lspci -nnk -s <pci>`
- `sudo dpdk-devbind.py --status`
- `sudo journalctl -u vpp -n 200`

Not currently supported

Card	Reason
Intel E825-C (NVM 3.81)	The DDP package required for E825-C-specific features fails signature verification on every public build we've tried; requires OEM-signed DDP or NVM update. Card is functional in Safe Mode (kernel <code>ice</code> driver, no DPDK).
1 GbE NICs (I350, DH8900CC series, etc)	Kernel / management-plane use only. Not routed through DPDK.

Known limitations

Being upfront about the shape of this alpha:

Control plane scale ceiling

The Quantum Controller currently processes **roughly 5–10 subscriber state changes per second** end-to-end (DHCP DORA + local-AAA lookup + session persist + interim tick). That's fine for:

- Small pilots (up to ~500 concurrent subscribers with normal churn)
- UI-driven ops (plan edits, pool changes, admin actions)
- Bench / lab evaluation

It's **not** fine for ISP-scale churn scenarios (bursts of >100 sessions/sec, mass CPE reboots, aggressive re-DORA storms). Under sustained churn above that rate, the controller backlogs; we've seen orphaned sessions and over-drain under stress. The controller team is actively working the throughput ceiling; expect meaningful improvement in the next release.

Data plane performance is not the bottleneck — VPP hits line-rate at 40G+/port on Tier 1 NICs. The scale limit is entirely in the control-path RPCs.

RADIUS not shipped

Standalone deployments use local-AAA (per-line records + Option-82 identity). If you need real RADIUS auth/accounting/CoA against an external server, that's a follow-on release — the `/bng/radius` config API is inert (returns 503) in this build.

DSCP remarking of out-of-contract traffic (issue #19)

The 3-color meter classifies traffic into green (in-contract), yellow (over CIR, under PIR), and red (dropped). Counters increment correctly per color. **Marking** — rewriting DSCP on yellow packets — is not yet implemented. Forwarding decisions are correct; only the DSCP bit-flip is missing.

External captive portal — you host it

Local captive-portal mode works. External-URL redirect mode (where subscribers get bounced to a portal on your own infrastructure) requires you to serve the portal; QuantumBNG doesn't ship a hosted one. RFC 7710 / RFC 8910 (DHCP Option 114 / IPv6 RA captive-portal URL emission) is targeted for a follow-on.

Mellanox ConnectX drag-to-VPP rough edge

As noted in the NIC section: the wizard's DPDK Bind step assumes `vfio-pci` rebinding, which works for Intel PMDs. Mellanox `mlx5` uses a bifurcated driver — the NIC stays on `mlx5_core` and the DPDK PMD grabs it in-place. If you drag a ConnectX in the wizard and it doesn't bind cleanly, you'll need to edit `/etc/vpp/startup.conf` by hand and add `dpdk { dev 0000:XX:00.0 }`, then restart VPP. We'll smooth this in a follow-on.

Setup wizard's DPDK Bind requires VPP briefly restart

Rebinding a NIC to VPP requires VPP to release its current port set. As of this release, dragging a NIC in the wizard will auto-stop + auto-restart VPP behind the scenes (the wizard is designed for fresh-install context where no subscribers are yet connected). On a live production box, the Settings-page ports panel prompts for confirmation first — active subscriber sessions will drop briefly during the VPP restart.

Deep-wipe / uninstall behavior

The bundled `uninstall.sh` now rebinds `vfio-pci` -owned NICs back to their kernel driver before purging VPP (fixed in this release). Prior to this fix, an `uninstall.sh` left NICs invisible to the kernel until manual `dpdk-devbind --bind=<drv>`. If you're upgrading from an earlier private build, run:

```
sudo dpdk-devbind.py --status          # find vfio-pci NICs
sudo dpdk-devbind.py --bind=i40e 0000:01:00.0  # rebind XL710
```

before reinstalling.

Upgrade / downgrade

None in this alpha. Uninstall (`sudo /home/<user>/quantumtouch-installer-1.2.0-01/uninstall.sh`), then reinstall with the new bundle. State under `/opt/bng-deploy/data` , `/var/lib/quantum-controller` , and `/opt/quantumtouch/data` is preserved by `uninstall.sh` for the settings + Prometheus TSDB paths (see the script header for the preserve/purge matrix), but for alpha we recommend a clean install per release to avoid unexpected state carryover.

Feedback

This is an alpha — we want to hear what works, what doesn't, and what surprised you.

Please include:

- Your hardware (`lscpu` , `lspci -nn | grep -i ether`)
- What you tried and what happened
- `sudo journalctl -u vpp -u vpp-bridge -u quantum-controller -u bng-deploy -n 300` if something crashed
- Which section of this doc you were following when it broke

Email: alpha@bluedot-insight.com

Appendix A: Quantum Controller v1.2.0 release notes

The following is the Quantum Controller team's own release notes, reproduced verbatim from github.com/Bluedot-Insight/bng/releases/tag/v1.2.0. This is the controller version bundled inside this QuantumBNG alpha.

Release v1.2.0

Overview

Standalone / free-tier release for Quantum BNG. The headline change is that the BNG can now run **without a RADIUS server**: a built-in **local-AAA** provider handles subscriber admission and tiering by line identity, making the product a self-contained BNG for small ISPs. This release also introduces a **free-tier session cap (1,000 concurrent subscribers)**, **DHCP lease-expiry timers**, **plan-edit propagation** to live sessions, and a broader **observability** surface (DHCP reply-path counters, session-cap visibility, write-path timing).

It is also the first **release-hardened** build: the management API now requires authentication by default, the test-only endpoints ship disabled, and the binary distribution carries its own end-user license (distinct from the source license).

Built against **quantum_core v1.4.0** and **quantum_test v1.2.0** (v1.1.0 shipped against quantum_core v1.3.0 / quantum_test v1.1.0).

Licensing is **not enabled** in this product build (`LICENSING_ENABLED := false`); the free tier requires no activation license. The core licensing machinery is present but compiled out / stubbed.

New Features

Local-AAA Server — Standalone Operation (no RADIUS required)

A built-in AAA provider lets the BNG admit and tier subscribers with **no external RADIUS server**. Identity is the **subscriber line** (DHCP Option 82 circuit-id / remote-id), and per-line records drive the onboarding outcome:

- `active` → accept (onboard at the record's plan)
- `suspended` → captive (dunning / policy hold — walled garden, plan retained)
- `terminated` → reject (dropped)

Two modes: `open` (flat-rate accept-all — the default) and `gated` (per-line policy). Captive-portal activation re-authenticates against the record and restores the subscriber's plan.

- `GET` / `PUT /bng/aaa/local` — provider config (mode, default plan, unknown-line policy).

- `GET / PUT / DELETE /bng/aaa/local/subscribers/{line}` — per-line records (status, plan, optional `session_timeout` / `idle_timeout` / `interim_interval`).

Free-Tier Session Cap

The product enforces a compiled **1,000 concurrent-subscriber** limit in the VPP plugin dataplane (tamper-resistant, independent of the JS controller). Onboards past the cap are silently dropped with a rate-limited log — no NAK storm.

- `GET /bng/health` reports `sessions.onboard_refused` (monotonic, climbs at the cap), plus `sessions.max` and `sessions.at_cap` sourced from the plugin.

DHCP Lease-Expiry Timers

A per-session lease timer reaps a subscriber whose CPE stops renewing (returns the address to the pool), closing a gap where an abandoned CAPTIVE or a CONNECTED-past-lease session could linger. The DHCP lease is now **pool-driven** and independent of any RADIUS Session-Timeout, so a Session-Timeout correctly drops an active subscriber to CAPTIVE while the lease stays valid.

Plan-Edit Propagation & Live Re-Tier

- `PUT /bng/plans/{name}` now **re-rates every live session on that plan** in place (response reports `re_rated` / `failed`) — an edit reaches existing subscribers, not just new ones.
- `POST /bng/sessions/{ip}/plan` — move a single live session to a different plan.

Expanded Observability

- `GET /bng/counters` gains a `dhcp` block — reply-path counters (`discovers_received` , `offers_emitted` , `requests_received` , `acks_emitted` , `naks_emitted`), monotonic since controller start.
- `GET /bng/health` gains `persistence` — interim-tick write-path timing (`save_ms` , `write_lock_wait_ms`) for spotting the DB-bound ceiling.

A Prometheus `/bng/metrics` endpoint is intentionally **not** provided — the JSON counters are designed to be scraped by the monitoring layer.

Security & Release Hardening

- **Authentication on by default** (`AUTH_REQUIRED_DEFAULT := true`). The management API requires a Bearer API key; provision keys with `create_api_key.sh` . Keys are hashed at rest, role-based (`read_only` / `read_write` / `admin_only`), and hot-reloaded.
- **Test endpoints off by default** (`TEST_ENDPOINTS_DEFAULT := false`). The `/bng/_test/*` and `/quantum/_test/*` routes are not mounted in the shipped product.

- **Binary distribution license.** The `.deb` packages and install bundle ship an end-user license (`LICENSE.txt`) distinct from the source `LICENSE` .

API Changes

Quick index; detail is in the sections above.

Added

- `GET / PUT /bng/aaa/local`
- `GET / PUT / DELETE /bng/aaa/local/subscribers/{line}`
- `POST /bng/sessions/{ip}/plan` — live re-tier

Changed

- `PUT /bng/plans/{name}` — re-rates live sessions on the plan (`re_rated` / `failed`)
- `GET /bng/counters` — adds the `dhcp` reply-path counter block
- `GET /bng/health` — adds `sessions.max` / `sessions.at_cap` / `sessions.onboard_refused` and a `persistence` timing block
- `GET / PUT /bng/radius` — the config surface is now **inert** (local-AAA is the backend): `GET` returns `{active:false, ...}` and `PUT` returns **503** rather than storing config nothing consumes. To be re-enabled when the real RADIUS client ships.

Removed from the shipped surface (test endpoints off by default — see Breaking Changes)

- `PUT / GET / DELETE /bng/_test/radius/stub` — RADIUS stub (dev / integration only)
- `POST /bng/_test/sessions/synthetic` — synthetic session injection (scale / cap testing)
- `POST /quantum/_test/vpp/{switch}/force_resync` and the other `/quantum/_test/*` fault-injection routes
- Not deleted from the code — simply not mounted unless `testEndpoints` is re-enabled via `/data/runtime_overrides.json` .

Bug Fixes

- **Onboarding / churn robustness (issues #20, #21).** Session dedup now keys on **subscriber identity (MAC + circuit-id + remote-id)**, independent of the DHCP `xid` , so a re-DORA reuses the in-flight session instead of spawning a duplicate. This eliminates the duplicate-session accumulation seen under storm and churn load; storm tests reconcile cleanly at 1000/1000.
- **Lease vs Session-Timeout conflation.** The DHCP lease had been derived from the RADIUS Session-Timeout, collapsing two independent clocks onto the same instant and (with the new lease timer) pre-empting the CONNECTED → CAPTIVE transition. The lease is now pool-driven and independent — see Lease-Expiry Timers above.

Breaking Changes

- **Management API now requires authentication.** v1.1.0 shipped with auth *disabled* by default; v1.2.0 ships it **enabled**. Clients must present a Bearer API key, and at least one key must be provisioned

(`create_api_key.sh`) before the API will accept requests. `AUTH_REQUIRED_CONFIGURABLE` remains true, so an operator can still opt out via runtime config for a trusted environment.

- **Test endpoints are off by default.** The RADIUS stub and synthetic-session routes are no longer available in the shipped build unless explicitly re-enabled (`/data/runtime_overrides.json` → `{"testEndpoints": true}`).
- `/bng/radius` is inert. Configuring RADIUS via the API now returns 503; local-AAA is the AAA backend for this release.

Resolved Issues

Closed under the v1.2.0 milestone; see the sections above for detail.

• 20 — onboarding performance under storm (identity/xid dedup)

• 21 — duplicate sessions on re-DORA

• 22 — DHCP reply-path counters / session-cap observability

Known Issues (controller-side)

- **Churn over-drain (#24), reported against v1.1.0** — believed **addressed** by the #20/#21 identity-dedup fix in this release (the v1.1.0 binary lacked it), but awaiting confirmation from a re-run of the churn-longevity harness against a v1.2.0 build. Note: session-dedup currently does offer-reuse; full tear-down replace-on-duplicate is a follow-on.
- **Production RADIUS client not yet shipped.** Standalone deployments use local-AAA and need no RADIUS server; production RADIUS integration (auth, accounting, CoA) is targeted for a follow-on release. The `/bng/radius` API is inert until then.
- `sessions.max` / `sessions.at_cap` require the plugin + vpp_bridge built with this release (the `BngCapStatus` RPC); against an older dataplane those two fields are omitted and `/bng/health` otherwise degrades gracefully.
- **Out-of-contract DSCP marking not yet implemented (issue #19).** Traffic is metered and counted into green / yellow / red bands, but out-of-contract (yellow) packets are not yet DSCP-remarked. Forwarding is unaffected.