DandyBNG User Guide

Introduction	2	CLI Users and Levels	34
Recommended Hardware and Settings	2	Log Server (NAT logging)	35
Useful Tips	5	Local Users (without RADIUS)	36
First-time setup checklist	6	SNMP (PRTG / Zabbix / LibreNMS)	37
What is new in V7.3.3	8	Web Panel	38
Configuring from the web (Config tab)	9	Traffic Split and Top organisations	42
Service-name rules	20	System health and auto-tune	44
Interface Configuration	12	Backup / Restore	46
PPPoE hardware spread (DDP)	15	DNS Cache Server	48
PPPoE and RADIUS	16	DNS redirect	49
IP Pools (IPv4 / IPv6)	17	Session license	51
IPv6 dual stack	18	Installing from the ISO	52
VLANs, VLAN ranges and Service Names	21	Ookla Speedtest Server	50
PPPoE packet sizes (MTU / MRU)	22	Time zone	54
FastPADO	24	Routing Mode (BGP, OSPF, IS-IS)	55
Pinging subscribers	26	IPoE Configuration	56
Why is a user not connected	27	High availability (HA)	57
NAT (SNAT / DNAT)	29	Displaying Configuration	61
NAT session limit per user	30	Maintenance commands	62
QoS: destination-group speeds	31	Command reference (all commands)	63
Static routes and the IX route policy	33	Troubleshooting	74

DandyBNG User Guide

Software V7.3.3 (November 2026).

This guide explains how to configure the BNG from its command line: interfaces, PPPoE and RADIUS, IP pools, VLANs, NAT, QoS, logging, SNMP, the web panel and day-to-day maintenance. Everything in the **bold** boxes is a command you will use on every installation.

Recommended Hardware and Settings

The BNG forwards in the Linux kernel with per-user policers. Throughput is decided by three things: how many CPU cores can work on packets, whether the PPPoE port can spread its traffic over several receive queues, and whether a packet stays on one CPU socket. The recommendations below follow from that.

Server

Part	Recommended	Why
CPU	Intel Xeon, 2.2 GHz or faster, 12 or more cores per socket; one or two sockets	Roughly one core per 3 to 4 Gbit/s of download including the ACK stream. Clock speed matters more than core count beyond 16 cores.
Memory	32 GB (16 GB minimum), all memory channels populated, same DIMMs on both sockets	Connection table (about 350 bytes per NAT flow; a busy box holds one to two million flows), PPPoE daemon, NetFlow export, web collector.
Disk	Enterprise SATA or NVMe SSD, 120 GB or more; hardware RAID 1 optional	Logs rotate; the system itself uses under 5 GB. Spinning disks slow down the boot and the log export.
WAN NIC	Intel X710 / XL710 (10G/40G) or Intel X520 (10G)	WAN traffic is plain IP and spreads over all queues on any Intel card.
PPPoE (LAN) NIC	**Intel X710 / XL710 with DDP enabled**	Only these cards can hash PPPoE frames across queues (DDP). On an X520 every PPPoE packet lands on one queue: plan about 1,500 to 2,500 users or 8 to 10 Gbit/s per X520 port.
Management NIC	Any onboard 1G port	Not in the forwarding path.
PCIe slots	PCIe 3.0 x8 or x16, electrically x8 at least	A 40G card is an x8 card; an x4 slot halves it. `show system health` reports the negotiated width.
Optics	Same type and speed at both ends (SR with SR, LR with LR); 10G optics in 10G ports	The PHYSICAL PORTS card shows light levels, module type and link drops.
Avoid	Realtek, onboard "LOM" cards for forwarding, Broadcom NetXtreme for PPPoE, USB NICs	Small rings, no PPPoE hashing, driver limits.

Card placement (CPU sockets)

On a two-socket server, put the **WAN card and the LAN (PPPoE) cards on different CPU sockets**: WAN on one socket, all LAN ports on the other. The health check shows which socket

("NUMA node") each card is in; its "NICs span NUMA nodes" line is expected with this layout. On a single-socket server there is nothing to choose.

Planning figures (per box)

Build	Download / Upload	Notes
1 socket, 2 x X520 (WAN + PPPoE)	8 to 10 Gbit/s / 1.5 Gbit/s	The single PPPoE receive queue is the limit; turn flow control off on the PPPoE port.
1 socket, dual-port XL710, DDP on	15 to 20 Gbit/s / 3 Gbit/s	Same-socket forwarding.
2 sockets, XL710 WAN on one socket, XL710 PPPoE port(s) on the other, DDP on	24 Gbit/s / 3.5 Gbit/s	With two PPPoE ports split subscribers by VLAN. Keep the public NAT pool under 128 users per address.

BIOS settings

Important for high performance: set the BIOS before the first boot of the installed box. Three settings decide how many packets each core can forward; a wrong one costs 20 to 40 percent of the box and is invisible in normal operation:

- **Virtualization: VT-x, VT-d (Intel) / SVM, IOMMU (AMD) and SR-IOV → Disabled.** With VT-d on, the chipset's IOMMU translates every network buffer on every packet: measured **about 19 percent of the forwarding CPU** on a production box (Xeon E5 v4, 40G ports). Since V7.3.2 the system boots with `iommu=pt`, which removes the translation even when VT-d is left on, so a box whose BIOS has no VT-d option or cannot be changed is still safe. `show system health` reports IOMMU: not active or passthrough when all is well, and a WARN with the fix when it is not.
- **Hyper-Threading / Logical Processor → Disabled.** Two threads share one core's packet-processing units; the box shows twice the CPUs and forwards less.
- **Power profile → Maximum Performance, C-states disabled, Turbo enabled.** A core that sleeps between packets wakes up too late for a full receive ring at peak.

A box that was installed with these settings wrong only needs the BIOS change and one reboot; nothing has to be reinstalled.

Setting	Value	Why
Hyper-Threading (Logical Processor)	**Disabled**	Two threads on one core fight for the same packet-processing units; real cores only.
Virtualization (VT-x, VT-d, SR-IOV; AMD: SVM, IOMMU)	**Disabled**	Not used by the BNG; VT-d puts the IOMMU in front of every NIC buffer (~19% of forwarding CPU measured). The kernel line `iommu=pt` (V7.3.2) neutralises it when it cannot be switched off.
Power profile / Energy policy	**Maximum Performance** (OS control off)	Cores must not sleep between packets; the OS also runs with idle polling and the performance governor.
C-states (C1E, C3, C6)	Disabled	Wake-up latency causes ring overruns at peak.
Turbo Boost	Enabled	Higher clock on the busy cores.
NUMA / Node interleaving	NUMA **enabled**, interleaving **disabled**	The OS must see the real socket layout to pin queues next to the cards.

Setting	Value	Why
PCIe link speed	Gen3 (not Auto)	Some boards negotiate Gen1/Gen2 on Auto after a power cycle.
Memory mode	Performance / Independent, not Mirroring or Sparing	Full memory bandwidth.
Boot	UEFI or legacy as the ISO installer offers; disable PXE and other boot devices after installation	Unattended reboots must land on the installed disk.
IPMI / iDRAC / iLO	Enabled, on the management network, with a changed password	Console access for upgrades and for any WAN change that goes wrong.
Watchdog	Disabled	A BIOS watchdog can reset a busy but healthy box.

After any hardware change (card moved, optic swapped, memory added) run `set system auto-tune apply` and then `show system health`: auto-tune re-pins the queues to the right socket and the health check confirms slot width, optics and losses.

Useful Tips

1. In the BIOS disable **Hyper-Threading** and **Virtualization**. On a two-socket server put the WAN card and the PPPoE cards on **different CPU sockets** (see Card placement); `show system health` shows the socket of each card.
2. Use Intel **X710 / XL710** cards for the PPPoE side when you plan more than a few Gbit/s. See PPPoE hardware spread. X520 cards work, but every PPPoE packet lands on one receive queue.
3. Default CLI login: username **admin**, password **Dandy@123**. Change it before the box goes live: `set system login user admin authentication password <new password>`, and delete users you do not need.
4. Factory web login, on the address and port set with `set system web address ...` (default port 8000): user **super**, password **Super@123** (level `superuser`, with the Config tab). Change this password before the box goes live: web **Config > Users**, or `set system web login user <name> authentication password <new password>` and `commit` in the CLI. See Web users and levels.
5. **commit** saves and applies; **commit discard** throws the pending changes away. Nothing takes effect before `commit`.
6. A commit applies **live**: PPPoE sessions stay up through interface, pool, RADIUS, NAT, QoS and SNMP changes. Only an IP-pool change needs re-login of the users to pick up new addresses, and a changed rate-limit reaches a user at the next login.
7. Press **TAB twice** to see the next available words. Press **q** to leave a long listing.
8. Keep a copy of `show configuration set` in a text file after every change. Pasting those lines back into a fresh box restores it.
9. Run **show system health** after any change. It checks about fifty things (forwarding, NAT rules, upload policer, SNMP agent answering, clock, disk, optics, ring losses) and tells you what to do about a WARN or FAIL. The same report is on the web panel's Stat tab and runs every few minutes by itself.
10. **Delete** is the mirror of **set**: repeat the same line with `delete` in front. For VLANs the `name` part offered by TAB can be included or left out.
11. Changes to the WAN address restart the web panel for about a minute; this is normal.
12. Keep the box's clock synchronised and set the time zone: the Traffic Split card derives your country from it.
13. Logs under /var/log are kept for 7 days: every log rotates daily at 05:40, or sooner when it passes 200 MB. Send what you must keep longer (NAT logs) to a log server.
14. Leave GRO, GSO and TSO off on the WAN and PPPoE ports; auto-tune sets them off at every boot and commit. Measured on V7.3.2: GRO on the WAN port costs 8 to 11 percent more CPU for the same download, because packets merged on the WAN side must be split again in software before they go out inside PPPoE.
15. Do not run `apt upgrade`, change the kernel, or edit files under `/etc` by hand; the CLI regenerates them. Updates come as a new ISO or from support.

First-time setup checklist

The addresses below are examples (203.0.113.x is a documentation range): use your own WAN port name, address, mask and gateway.

1. **Find the WAN port** (on the server console or KVM, log in as **admin**; the box has no address yet): list the ports and pick the one cabled to your upstream (state **up**, the expected speed):

   ```
   show interfaces
   ```

   Ports are named by position and speed, for example **eth0-10G**, **eth1-10G**, **eth0-40G**. Use that name in the next step.

2. **WAN address and gateway**, then commit:

   ```
   set interfaces ethernet eth0-40G address 203.0.113.10/29
   set interfaces ethernet eth0-40G gateway 203.0.113.9
   commit
   ```

   From now on the box is reachable on its WAN address. Connect to the CLI with SSH on the default port 22, user **admin**: `ssh admin@203.0.113.10`

3. **Web panel address**, then commit:

   ```
   set system web address 203.0.113.10 port 8000
   commit
   ```

4. **Your web login** with the superuser level (needed for the Config tab, where the router is configured from the web), then commit:

   ```
   set system web login user <username> authentication password <your password>
   set system web login user <username> level superuser
   commit
   ```

   Open `http://203.0.113.10:8000` in a browser and log in with that username and password. Staff who only watch get the viewer level: `set system web login user <name> level viewer`.

5. **Host name, CLI password, time zone and auto-tune** (auto-tune sets the network cards and CPU for forwarding):

   ```
   set system host-name <name>
   set system login user admin authentication password <new password>
   set system time-zone Asia/Kolkata
   set system auto-tune apply
   commit
   ```

6. **License**: the box runs as a demo with 100 sessions until a license key is set. Show the license state and the box's **Fingerprint**:

   ```
   show system license
   ```

   Send the **Fingerprint** line to Dandy Network (WhatsApp / call +91 7304391953 or info@dandynetwork.in) to get your license key, then install it: `set system license key <key-string>`. The new limit applies within a minute, without a restart. See Session license.

7. **PPPoE**: access-concentrator, RADIUS servers, NAS address, DAE server, DNS servers, IP pool(s), pool gateway-address (see PPPoE and RADIUS and IP Pools).
8. **PPPoE listening interface**:

   ```
   set service pppoe-server interface <port> vlan-range 2-4094
   ```

   Or list specific VLANs with `vlan-id`.

9. **NAT**:

   ```
   set system snat source-address-pool <private pool> translation pool <public pool> out-interface <wan>
   ```

   The out-interface must be the real WAN port name.

10. **Log server** (mandatory for CGNAT), **SNMP** community and allowed monitoring host.
11. **On X710 / XL710 PPPoE ports**:

    ```
    set interfaces ethernet <port> ddp enable
    ```

12. **Commit and check**:

    ```
    commit
    show system health
    ```

    Everything should be PASS or an explained WARN before the first subscriber.

What is new in V7.3.3 since V7.3.2

- **High availability (HA)**: two boxes work as one pair. One serves, the other holds a live copy of every subscriber and connection and takes over in about a second; **the subscribers do not re-dial**. PPPoE (IPv4 and IPv6) and IPoE, with FastPADO. Two lines per box turn it on. See High availability.
- **Configuration copy**: on a pair you configure only the box that serves; the standby takes every change at the commit. See Configuration copy.
- **Config tab > High Availability (HA)**: set up, live state, configuration copy state, switch over. See From the web.
- **Backup / Restore** new: the configuration as one file, from the CLI (`backup ...`) and the web (Config > Backup / Restore). A copy is saved by itself at every commit; a restore on another server asks for the ports it does not have. See Backup / Restore.
- **Config > Users** new: CLI users and web users in one web page (add, delete, password, level). `admin` and `super` cannot be deleted. The factory web login is `super`. See Web users and levels.
- **Config > PPPoE Interface** new (was PPPoE VLAN): PPPoE on a whole port can be added and removed from the web; IP pools can be deleted from the web; the restart of the PPPoE/IPoE server asks for a typed Yes; all questions are asked on the page itself.
- **Port names for added cards** new: a network card fitted after the installation gets its `eth<N>-<speed>` names from auto-tune. See System health and auto-tune.

Items marked new were added after V7.3.3 build 1, the ISO on the Downloads page. They come with the next build.

Changes in earlier versions: release notes.

Configuring from the web (Config tab) new

Since V7.3.2 the web panel has a **Config** tab, visible only to web users with the `superuser` level (see Web users and levels). Everything you can configure there you can also configure in the CLI: the web builds the same `set` / `delete` lines and hands them to a small service on the box (`dandy-confd`) that runs the CLI code itself. Validation, the saved configuration, the generated daemon settings and the apply steps are therefore identical, and `show configuration set` in the CLI always shows what the web did.

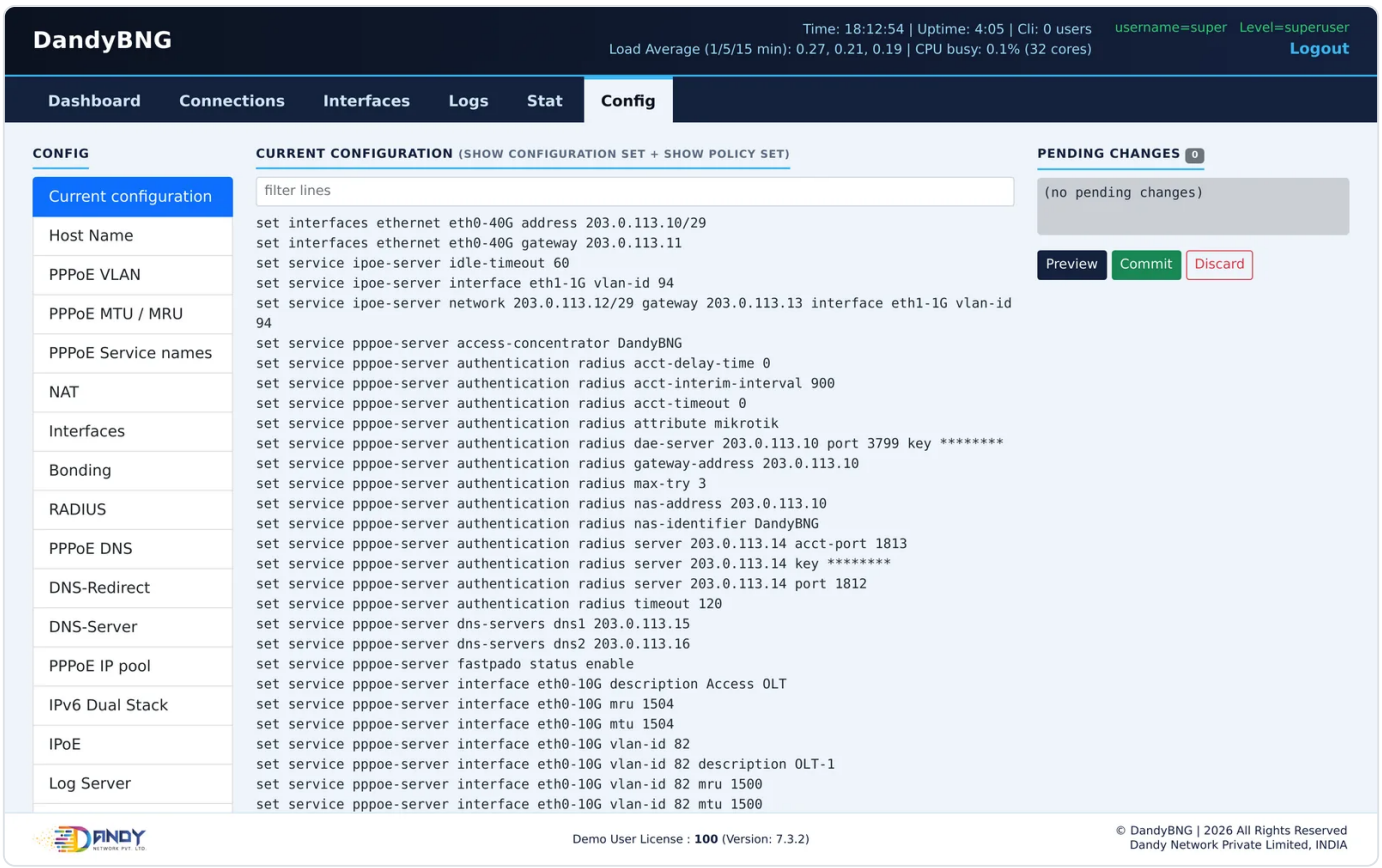
How a change is made: pick a page, fill the form, press its button. The CLI line appears under **Pending changes** on the right (a refused line shows the CLI's reason). **Preview** lists the configuration lines that will be added and removed. **Commit** saves and applies them; **Discard** drops them. Rows a pending change adds are green, rows it deletes are struck through. Only one superuser holds pending changes at a time; another one sees who it is and may discard them.

Pages:

- **Current configuration**: `show configuration set` and `show policy set` with a filter box.
- **Host Name**: shown in the web header and the CLI banner.
- **Users** (below Host Name) new: the CLI users and the web users in one page: add, delete, change password, change level. A change is applied at once (no Commit step). The CLI user `admin` and the web user `super` cannot be deleted and keep their level; their password can be changed. See Web users and levels.
- **PPPoE Interface** (PPPoE on a whole port, on a VLAN or on a VLAN range of a port; before this version the page was called PPPoE VLAN), **PPPoE Service names** (VLAN chosen from a list; the blank service name switch is global), **PPPoE DNS**, **PPPoE IP pool**, **PPPoE FastPADO** (enable, live status, whitelist, login-fail ban with explanations).
- **NAT**: part 1 SNAT (subscriber pool to public pool or IP on the WAN port, NAT sessions per subscriber), part 2 DNAT port forwarding (public address or `any`, tcp/udp, public port, inside host, inside port). Every rule is checked before it is accepted: network addresses, no private public pools, existing out interface, no overlap with another rule or with the box's own address, inside hosts inside a subscriber pool, public ports not used by the box itself (web, CLI, API, SNMP, DNS, speedtest, RADIUS CoA), no duplicate forwarding.
- **PPPoE MTU / MRU**: packet sizes in three steps. 1 the PPP size for new logins; 2 the size of a PPPoE port, a VLAN or a VLAN range; 3 the live state with an **Apply** button where a size is waiting.
- **Interfaces**: address, gateway, IPv6, MTU, description on a port or on a VLAN of it (`eth0-1G.97`).
- **Bonding**: bonds in three steps. 1 the bonds of this box with their settings; 2 create or change a bond: member ports, mode and LACP settings, address and gateway; 3 the live state (mode in use, member ports, link). The first port of a new bond also adds the default settings. See From the web.
- **IPv6 Dual Stack**: IPv6 for the PPPoE users in the order of the CLI steps: DNS servers, WAN pool, delegated (LAN) pool, gateway, pools in use. See From the web.

- **RADIUS**: the server (the box uses one RADIUS server: a new address replaces the current one), DAE (CoA/Disconnect) listener, general settings; the page tells you when a change needs a service restart.
- **DNS-Redirect**, **DNS-Server** (port 53 and threads = CPU cores by default), **IPoE** (access VLANs, subscriber networks on those VLANs, required idle timeout, default 60 s; IPoE and PPPoE never share an interface), **Log Server**, **SNMP**, **Ookla Speedtest Server** (with the DNS records, ports and submission steps).
- **QoS** (Policy Bypass Rate-limit): the QoS groups in five steps. 1 choose a group or create a new one; 2 add its networks, one at a time or from a .txt file (one network per line, # comments; an example file is shown on the page); 3 add ports (TCP and UDP); 4 an optional priority; 5 the speed (all / plus / multiply, always or a daily window), which makes the group active. Each network is checked: it must be a network address of /8 or longer, not loopback, multicast or reserved space, not a subscriber pool and not overlapping a network of any group. Live state shows the groups as the dataplane applies them and the subscribers using a group with passed and dropped counters. Deleting a group is refused while it has pending changes (commit or discard them first).
- **High Availability (HA)** V7.3.3: the pair of two boxes: settings of this box, set up or change, live state, configuration copy, switch over. See From the web.
- **Backup / Restore** (below System Health) new: save, download, upload and restore the configuration; a port of the backup that this server does not have is asked for, with the list of the ports here.
- **Restart**: PPPoE/IPoE server (shows the active session count first), web, DNS, SNMP (greyed out when the service is disabled), server reboot and shutdown. Every restart asks on the page itself; the PPPoE/IPoE server restart, the reboot and the shutdown go ahead only after you type **Yes**.
- **Audit Log**: every configuration line, commit, discard, restart and disconnect made from the web with user, address and time, newest first, 50 per page.

Changes to PPPoE DNS, IP pools or the DAE listener are saved at commit but the running service only picks them up after **Restart PPPoE/IPoE server**; the page shows a banner with the restart button after such a commit. Nothing restarts without your confirmation.



Config tab: the page list on the left, the current configuration in the middle, pending changes with Preview / Commit / Discard on the right.

Interface Configuration

Port names follow speed: `eth0-1G`, `eth0-10G`, `eth0-40G` ... Use `show interfaces` to list them and `show interfaces ethernet <port> blink count 10` to find a port by its LED. The ports of a card that is fitted later are named by auto-tune.

```
set interfaces ethernet <port> address <address/mask>
```

```
set interfaces ethernet <port> gateway <address>
```

```
set interfaces ethernet <port> ipv6-address <address/mask>
```

```
set interfaces ethernet <port> ipv6-gateway <address>
```

```
set interfaces ethernet <port> dns1 <address>   |   dns2 <address>
```

DNS servers the BNG itself uses (RADIUS by name, updates).

```
set interfaces ethernet <port> mtu <number>
```

```
set interfaces ethernet <port> description <text>
```

```
set interfaces ethernet <port.vlan> address <address/mask>
```

An address on a VLAN sub-interface (for example `eth0-10G.100`).

```
delete interfaces ethernet <port> address <address/mask>
```

Bonding (LACP) reworked

A bond joins two or more ports into one link to the same switch or router: more capacity, and the link stays up when one port or cable fails. **Bonding is recommended for the WAN link only**, not for subscriber (PPPoE / IPoE) ports. The address and the gateway belong on the bond, not on its member ports.

```
set interfaces bonding bond8 member interface eth0-10G
set interfaces bonding bond8 member interface eth1-10G
set interfaces bonding bond8 mode 4
set interfaces bonding bond8 lacp-rate 1
set interfaces bonding bond8 hash-policy layer2+3
set interfaces bonding bond8 miimon 100
set interfaces bonding bond8 updelay 800
set interfaces bonding bond8 downdelay 200
set interfaces bonding bond8 address 103.100.100.2/30
set interfaces bonding bond8 gateway 103.100.100.1
commit
```

The address and gateway above are an example: use your own. The switch or router on the other side must have the same ports in one LACP group before you commit.

Setting	Values	Meaning
`member interface`	a port name	A port of the bond; repeat the line for every port. The bond name is free (a letter first, at most 15 characters, for example `bond0`).
`mode`	`4` or `802.3ad` (LACP); also 0-6 or balance-rr, active-backup, balance-xor, broadcast, balance-tlb, balance-alb	How the ports work together. Use mode 4 with a switch that runs LACP. **Always set the mode**: a bond without a mode line runs as round-robin (mode 0), which an LACP switch does not accept.
`lacp-rate`	`1` / `fast` or `0` / `slow`	LACP messages every second (fast) or every 30 seconds (slow); mode 4 only. Fast finds a dead port sooner.
`hash-policy`	layer2, layer2+3, layer3+4, encap2+3, encap3+4, vlan+srcmac	How traffic is spread over the member ports. `layer2+3` (MAC and IP addresses) is the usual choice.
`miimon`	1-10000 (milliseconds)	How often the link of every port is checked. Usual value 100.
`updelay` / `downdelay`	0-60000 (milliseconds)	How long a port must be up before it is used again, and how long it must be down before it is taken out. Multiples of `miimon`.
`address` / `gateway` `ipv6-address` / `ipv6-gateway`	address/mask, address	The address and gateway of the bond. A bond may also have no address.

Rules for a member port (a refused line says which one):

- the port exists and is not a member of another bond;
- it has no address or gateway of its own: delete them first, they go on the bond;
- it is not a PPPoE or IPoE port;
- once it is a member it takes no address, gateway, PPPoE or IPoE of its own.

Moving the present WAN port into a bond: delete the address and gateway of the port, add the port (and the second one) to the bond, put the same address and gateway on the bond, then `commit` once. The WAN link drops while it changes, and it only comes back when the LACP group on the other side is ready: do it in a maintenance window with console access at hand.

```
delete interfaces bonding bond8 member interface eth1-10G
```

Takes one port out of the bond.

```
delete interfaces bonding bond8
```

Removes the whole bond with all its settings.

```
show interfaces bonding   |   show interfaces bonding <name>
```

Shows the committed bonds. With none it answers `no bond configured`, and for a name that does not exist `bonding <name> is not configured`. A bond that was typed but not committed yet is not listed.

`commit` rebuilds a bond whose members or settings changed: its link drops for a moment. A commit that changes something else leaves the bond alone. If the bond could not be built as configured (a port did not join, another mode is running), commit says so.

Bonding from the web

Config tab > **Bonding** (superuser level), directly below Interfaces, has the same settings and rules as the CLI:

- **1. Bonds**: every bond of the configuration with its member ports and settings; a line or the whole bond can be deleted here.
- **2. Create or change a bond**: pick the bond (or new bond and a name) and add its ports. Only free ports are offered. **The first port of a new bond also adds the default settings**: mode 4 (LACP), LACP rate 1 (fast), hash policy layer2+3, link check 100, up delay 800, down delay 200. Change a value when the switch is set differently; **Fill the default values** puts the defaults into the form. Then add the address and the gateway.
- **3. Live state**: the committed bonds against what runs now: the mode in use, each member port with its link, and for LACP whether the switch answers.

Nothing changes on the box until **Commit**. A new bond that is not committed yet is taken back line by line; a committed bond is deleted as a whole.

PPPoE hardware spread (DDP) new

Network cards cannot normally look inside PPPoE frames, so every PPPoE packet from subscribers lands on **one** receive queue and one CPU drains the whole access port (the box spreads the follow-up work with RPS, but the queue itself stays the limit). Intel X710/XL710 cards accept a "Dynamic Device Personalization" profile that teaches the card to hash PPPoE on the inner addresses, spreading subscribers over all queues.

```
set interfaces ethernet <pppoe-port> ddp enable
```

```
set interfaces ethernet <pppoe-port> ddp disable
```

```
show interfaces ethernet <port> ddp
```

- Only for the **PPPoE** port(s). The WAN carries plain IP and gains nothing.
- Other cards (X520, Realtek, Broadcom) are refused at `set` time with a clear message.
- The profile is re-loaded automatically at every boot while `enable` is configured. Disabling takes effect at the next reboot (the running card keeps the profile until then; `show` says so).
- Needs NIC firmware 6.01 or newer; `show ... ddp` prints the firmware version.

PPPoE and RADIUS

```
set service pppoe-server access-concentrator <name>
```

```
set service pppoe-server authentication radius nas-address <wan ip>
```

```
set service pppoe-server authentication radius nas-identifier <name>
```

```
set service pppoe-server authentication radius gateway-address <wan ip>
```

```
set service pppoe-server authentication radius attribute mikrotik
```

Rate-limit attribute format: `mikrotik` (Mikrotik-Rate-Limit "50M/50M") or `dandy`.

```
set service pppoe-server authentication radius server <address> key <secret>
```

```
set service pppoe-server authentication radius server <address> port 1812
```

```
set service pppoe-server authentication radius server <address> acct-port 1813
```

```
set service pppoe-server authentication radius server <address> fail-timeout | max-fail | req-
limit | weight <value>
```

Add a second `server` line for a backup RADIUS; `weight` balances between them.

```
set service pppoe-server authentication radius dae-server <wan ip> port 3799 key <secret>
```

Incoming CoA / Disconnect from the billing system.

```
set service pppoe-server authentication radius timeout | max-try | acct-timeout | acct-delay-time
<value>
```

```
set service pppoe-server authentication radius acct-interim-interval <seconds>   (default 900)
```

```
set service pppoe-server authentication radius acct-interim-jitter <seconds>
```

```
set service pppoe-server dns-servers dns1 <address>   |   dns2 <address>
```

```
set service pppoe-server ipv6-dns-servers dns1 <address>   |   dns2 <address>
```

```
set service pppoe-server pado-delay <ms> sessions <count>
```

Slow the PADO answer once this many sessions exist (load sharing between BNGs on the same VLAN).

IP Pools (IPv4 / IPv6)

```
set service pppoe-server ip-pool address <network/mask> name <name>
```

```
set service pppoe-server ip-pool address <network/mask> name <name> next <other pool>
```

Addresses are given lowest first: the first user gets the first address of the pool, and an address that becomes free is used again before a higher one. A `network/mask` pool starts at its first host address and never gives out the network or broadcast address. A pool change is read when the PPPoE service starts.

Chains pools: when one is full the next one is used, and a free address in the first pool is used again before the next pool. A pool named `expired-pool` (or any name without a NAT rule) gives expired accounts an address but no internet.

```
set service pppoe-server ip-pool gateway-address <wan ip>
```

```
set service pppoe-server ip-pool-name <name>
```

The default pool used when RADIUS sends no Framed-Pool.

```
delete service pppoe-server ip-pool address <network/mask>
```

IPv6

```
set service pppoe-server ipv6-pool address <prefix/mask,length> name <name>
```

```
set service pppoe-server ipv6-pool delegated-address <prefix/mask,length> name <name>
```

```
set service pppoe-server ipv6-pool gateway-address <wan ipv6>
```

```
set service pppoe-server ipv6-pool-name <name>   |   ipv6-pool-delegate-name <name>
```

A complete example is in IPv6 dual stack.

Size the pool for growth: a /19 holds 8,190 addresses, a /18 holds 16,382.

IPv6 dual stack

PPPoE users get IPv6 next to their IPv4 address: one prefix for the WAN side of the user's router and one delegated prefix for the LAN behind it. The IPv4 setup (IP pools, NAT) stays as it is.

The addresses below are examples only (`2001:db8::/32` is the range reserved for documentation). Replace every one of them with your own IPv6 DNS servers, your own prefixes and the IPv6 address of your WAN port.

Step	Command	What it sets
1	`set service pppoe-server ipv6-dns-servers dns1 2001:db8:53::1` `set service pppoe-server ipv6-dns-servers dns2 2001:db8:53::2`	IPv6 DNS servers given to the PPPoE users.
2	`set service pppoe-server ipv6-pool address 2001:db8:3::/48,64 name ipv6pool1`	IPv6 WAN for the PPPoE users: each user's router gets a /64 out of this /48 on its WAN side.
3	`set service pppoe-server ipv6-pool delegated-address 2001:db8:2::/48,64 name ipv6delegate1`	IPv6 LAN for the PPPoE users: each user's router is delegated a /64 out of this /48 for the devices behind it.
4	`set service pppoe-server ipv6-pool gateway-address 2001:db8::12`	The gateway of the PPPoE users: the IPv6 address of your WAN port.
5	`set service pppoe-server ipv6-pool-name ipv6pool1` `set service pppoe-server ipv6-pool-delegate-name ipv6delegate1`	Which pools are used: the names given in steps 2 and 3.
6	`commit`	Applies the change.

In `<prefix>/48,64` the first number is the size of the whole pool and the number after the comma is the size each user gets: a /48 cut into /64 serves 65,536 users. The WAN pool and the LAN pool must be two different prefixes, and both must be routed to this box by your upstream router.

```
show service pppoe-server ipv6-pool   |   ipv6-pool-name   |   ipv6-pool-delegate-name   |
ipv6-dns-servers
```

```
delete service pppoe-server ipv6-pool address 2001:db8:3::/48,64 name ipv6pool1
```

A pool is deleted with its whole line, the name included (and `next <pool>` if it has one). First take it out of use: `delete service pppoe-server ipv6-pool-name ipv6pool1`.

The WAN port itself gets its IPv6 address and gateway in Interface Configuration (`ipv6-address`, `ipv6-gateway`). The IPv6 address of an online user is shown on the web **Connections** page.

IPv6 dual stack from the web

Config tab > **IPv6 Dual Stack** (superuser level) has the same settings in the same order: IPv6 DNS servers, the WAN pool, the delegated (LAN) pool, the gateway and the pools in use. Every value is checked before it is accepted: a real IPv6 address or prefix, a user size that fits inside the pool, no overlap between two pools, no pool name used twice, a WAN pool not chosen as the delegated pool or

the other way round. A pool that is in use, or that is the next pool of another one, cannot be deleted until that is changed.

The values are saved at **Commit**; the PPPoE server uses them after **Restart PPPoE/IPoE server** (the page shows the banner with the button). Nothing restarts without your confirmation.

Service-name rules V7.3.2

A device announces a **service name** in its PPPoE discovery (PADI); many devices send none. The box answers only what the rules allow. Verified on the daemon:

Rule on the port/VLAN	Device sends no name	Device sends `test`	Device sends another name
`service-name test`	refused	**accepted**	refused
`service-name blank` (per VLAN)	**accepted**	refused	refused
`service-name anyname`	**accepted**	**accepted**	**accepted**
`test` + `blank`	**accepted**	**accepted**	refused
no rule on this VLAN, rules elsewhere	refused	refused	refused
no rule here, global `blank enabled`	**accepted**	refused	refused
no service name anywhere on the box	**accepted**	**accepted**	**accepted**

```
set service pppoe-server service-name blank enabled | disabled        (global: devices without a
service name, every port and VLAN)
set service pppoe-server service-name <name> | blank | anyname interface <port> vlan-id <number>
show service pppoe-server service-name                                (what each port/VLAN
accepts)
show subscriber <user|mac>                                            (why a device is refused
and the line that fixes it)
```

Rules apply to new logins only: changing them never disconnects a connected session. Deleting the last service name of a VLAN prints a warning with the effect after commit; the web page asks for confirmation. Refused discovery requests appear in the Logs tab under "Service-name mismatch" with the device's last known username.

VLANs, VLAN ranges and Service Names

```
set service pppoe-server interface <port>
```

PPPoE on the whole port, without a VLAN (untagged users). One port can carry both: its own untagged users and VLANs. `delete service pppoe-server interface <port>` ends PPPoE on the whole port only; the VLANs of that port stay. In the web: **Config > PPPoE Interface**.

```
set service pppoe-server interface <port> vlan-range <low-high>
```

Listen for PPPoE on every VLAN in the range (1 to 4094). Example: `vlan-range 2-4094`.

```
set service pppoe-server interface <port> vlan-id <number>
```

```
set service pppoe-server interface <port> vlan-id <number> description <text>
```

```
delete service pppoe-server interface <port> vlan-range <low-high>
```

```
delete service pppoe-server interface <port> vlan-id <number>
```

Both delete forms work with or without the trailing `name ...` that TAB completion offers.

Service names

```
set service pppoe-server service-name blank enabled   |   disabled
```

Accept PPPoE requests with an empty service name on all interfaces (enabled is the usual setting).

```
set service pppoe-server service-name <name> interface <port> vlan-id <number>
```

```
set service pppoe-server service-name anyname interface <port> vlan-id <number>
```

```
set service pppoe-server service-name blank interface <port> vlan-id <number>
```

```
delete service pppoe-server service-name <name> interface <port> vlan-id <number>
```

A VLAN keeps working when its description, mtu or mru is deleted: `delete service pppoe-server interface <port> vlan-id <number> description | mtu | mru` removes only that one entry.

Service-name changes apply at `commit`, without a restart. When no service name is configured at all, every service name is accepted. A refused request is logged as "Service-Name mismatch" with the service name the device sent, its MAC, the port and the VLAN: `show connections log`.

PPPoE packet sizes (MTU / MRU) new

There are two sizes. The **PPP size** is the largest packet of a subscriber session (the usual PPPoE value is 1492). The **link size** is what the access port or VLAN under it carries (a plain Ethernet link is 1500). PPPoE needs 8 bytes of the link, so a login never gets more than **the link size minus 8**. On a standard network nothing has to be set: without these commands every login gets 1492.

PPP size for all users

```
set service pppoe-server mtu <1280-1500>
```

```
set service pppoe-server mru <1280-1500>
```

```
delete service pppoe-server mtu   |   delete service pppoe-server mru
```

`mtu` is the largest packet the BNG sends to a user, `mru` the largest packet it asks the user's router to send. Default 1492 each. The `mtu` is a strict maximum: a router that can receive more is accepted, but is never sent more than this; a router that asks for less gets the smaller size it asked for. The value is used for new logins after `commit`; users who are online keep their size until they log in again, nobody is disconnected.

Link size of a port, a VLAN or a VLAN range

```
set service pppoe-server interface <port> mtu <1288-9216>   |   mru <1288-9216>
```

```
set service pppoe-server interface <port> vlan-id <number> mtu <number>   |   mru <number>
```

```
set service pppoe-server interface <port> vlan-range <low-high> mtu <number>   |   mru <number>
```

```
delete service pppoe-server interface <port> [vlan-id <number> | vlan-range <low-high>] mtu   |
mru
```

Only needed when the access network is not a plain 1500 link: a VLAN that carries less (a radio link, a tunnel) gets its real size, and the users on it get that size minus 8; a port that needs room for extra tags gets a size above 1500. A VLAN without its own size follows the port, but never above 1500. The CLI refuses a VLAN larger than its port and an `mru` larger than the `mtu`, and names what to change first. When the PPP size does not fit in a link, `set` prints a note with the size the users of that link really get.

What is in effect

```
show service pppoe-server interface
```

```
interface     mtu    mru    ppp-size   live-mtu users  state
eth0-10G      1504   1504   1480/1480  1504     1      applied
eth0-10G.82   1500   1500   1480/1480  1500     1      applied
PPP size for new logins: mtu 1480, mru 1480. ppp-size = mtu/mru a login on that line really gets
(never above the link size minus 8)
```

`mtu` / `mru`: the configured link size. `ppp-size`: what a new login on that line gets. `live-mtu`: the size the interface has now. `state`: applied, or PENDING with the command that applies it. The same table is printed by `commit` when a size was changed.

Commit never disconnects users for a size change. A link size is changed only while nobody is logged in on that link (changing it under running sessions would drop all of them). A VLAN with users online stays on its old size and shows PENDING: it takes the new size when nobody is on it, or at once with the restart command below. A new size of the port itself is applied only by restarting the port.

Applying a size at once

```
restart service pppoe-server interface <port> vlan-id <number>
```

```
restart service pppoe-server interface <port>
```

Disconnects the users of that one VLAN (or of every VLAN on the port), applies the waiting sizes and lets the users log in again within seconds. The command first prints how many users will be disconnected and continues only after you type `yes`. When the size of the port itself changes, the link of the whole port resets for a few seconds; the warning says so. Do it in a maintenance window on a busy port.

The MTU column of the web Connections page is the size each session really sends with.

From the web

Config tab > **PPPoE MTU / MRU** (superuser level) has the same settings and rules as the CLI:

- **1. PPP size for new logins**: MTU and MRU, 1280 to 1500 (empty = default 1492). **Back to the default** removes both.
- **2. Port and VLAN sizes**: pick the PPPoE port and either the whole port or one of its VLANs or VLAN ranges, enter MTU and/or MRU (1288 to 9216). A refused value shows the reason, for example a VLAN larger than its port or an MRU larger than its MTU.
- **3. Live state**: the committed sizes against the interfaces: configured MTU / MRU, the size running now, the PPP size a new login gets, users and state. A size that waits shows an **Apply** button; it says how many users will be disconnected and asks for confirmation (twice for a whole port with users). Apply on a VLAN disconnects only that VLAN; Apply on a port disconnects every VLAN of it and, if the port size changes, resets its link for a few seconds. Users log in again within seconds.

As in the CLI, Commit never disconnects anyone: a VLAN size goes in by itself as soon as nobody is on that VLAN, a port size waits for Apply or the next reboot. Every Apply is in the Audit Log.

FastPADO

FastPADO answers PPPoE discovery in the network card driver so this BNG wins the PADI/PADO race against other servers on the same VLAN. It is **off by default**; enable it per box once the PPPoE ports are configured. It uses exactly the ports given by `set service pppoe-server interface ...` (plain port or VLAN form) and, on the VLAN form, answers only on the VLANs you configured.

```
set service pppoe-server fastpado enable   |   disable
```

```
show service pppoe-server fastpado
```

Shows each port (armed, VLANs or plain port) and how many known subscribers, strangers and banned devices there are. Known subscribers (anyone who authenticated here) get the instant answer; an unknown device is answered by the daemon after 1.5 s; a device whose RADIUS reject says it is not your subscriber is ignored for 10 minutes (a "stranger"). Flooding devices are slowed down, never cut off.

```
set service pppoe-server fastpado whitelist <mac>   |   delete service pppoe-server fastpado
whitelist <mac>
```

```
show service pppoe-server fastpado whitelist   |   banned
```

A whitelisted MAC is never slowed down or banned. Enabling or disabling FastPADO re-initialises Intel X520 ports (about 2 s of link loss); sessions normally survive it.

Login-fail ban

Optional and **off by default**; it works only while FastPADO is enabled. A device whose logins RADIUS rejects again and again (wrong password, unknown account) gets no PADO for a while, so it stops loading RADIUS. Rejects such as "already logged in" or "session limit" never count, a whitelisted MAC is never banned, and a subscriber who has authenticated on this box is banned for at most 2 minutes.

```
set service pppoe-server fastpado login-fail-ban enable   |   disable
```

```
set service pppoe-server fastpado login-fail-ban attempts <3-20>           (default 5
rejects ...)
set service pppoe-server fastpado login-fail-ban window <60-1800>           (... within 300 s)
set service pppoe-server fastpado login-fail-ban ban-time <60-3600>         (first ban 120 s)
set service pppoe-server fastpado login-fail-ban max-ban-time <60-7200>     (repeat bans grow up
to 1800 s)
set service pppoe-server fastpado login-fail-ban known-max-ban-time <60-600> (known subscribers,
max 120 s)
```

```
show service pppoe-server fastpado login-fail-ban
```

```
delete service pppoe-server fastpado login-fail-ban [attempts|window|ban-time|max-ban-time|known-
max-ban-time]
```

`show` prints the settings, the ban ladder (first, second, third ban) and the MACs banned right now. Deleting one setting returns it to its default; deleting `login-fail-ban` turns the ban off. Apply with `commit`, no restart.

Pinging subscribers new

A ping from the BNG to a user goes out with the PPPoE gateway address (`gw-ip-address`) as its source. When that address is private, or is also the WAN address, many routers never answer it although the session works. With subscriber-ping, pings to users leave with the WAN public address instead, which every router answers. It covers pings from the BNG itself and pings from your NOC or monitoring that pass through the BNG. It is **on by default**.

```
set service pppoe-server subscriber-ping enable   |   disable
```

```
set service pppoe-server subscriber-ping source auto   |   <x.x.x.x>
```

```
show service pppoe-server subscriber-ping
```

```
delete service pppoe-server subscriber-ping   (back to the default: enabled, source auto)
```

`auto` uses the first public IPv4 address on the WAN port. A box with only a private WAN address does not apply it (commit says so): set `source` to an address the routers answer. Only ICMP echo requests towards users are rewritten: user traffic, NAT and FastPADO are not affected, and a ping from one user to another keeps its own source. When the gateway address already is the WAN public address nothing changes. `show` prints the source in use and how many pings were rewritten (BNG and NOC). Applies at `commit` and at boot, no restart.

Why is a user not connected new

When a subscriber calls, one command tells you whether the problem is the account, the device or the BNG. It looks the user up by username, device MAC or IP address and combines the connection log, the PPPoE discovery watcher, FastPADO, the live sessions and the configuration into a verdict, what to do about it, and the timeline of the last attempt.

```
show subscriber <username>   |   show subscriber <mac>   |   show subscriber <ip>
```

```
show subscriber failures   |   show subscriber failures last <minutes>   (default 60)
```

```
mahendra1  mac B4:B0:24:EC:9E:9E  eth0-10G vlan 82  last login attempt 4 min ago
status : NOT CONNECTED
reason : RADIUS rejected the login: "Account expired mahendra1" (4 min ago)
what to do: activate or renew the account in the billing system; device and link are fine
discovery: PADI 3 (last 4 min ago, service-name any)  PADO 3 (4 min ago)  PADR 3 (4 min ago)
PADS 3 (4 min ago)
timeline of the last attempt:
  10:21:05  login request sent to RADIUS (user mahendra1, device b4:b0:24:ec:9e:9e, port
eth0-10G.82)
  10:21:05  RADIUS Access-Reject "Account expired mahendra1"
  10:21:05  authentication failed
```

The verdicts you can see, and what they mean:

Reason	Where the problem is
CONNECTED (with IP, session time, port and VLAN)	Nowhere on the BNG: the user is online. If the customer still has no service, see Troubleshooting ("Users log in but have no internet").
wrong password / account not active or expired / unknown account / already logged in	Billing or RADIUS: the reply message is quoted. "Already logged in" means another session of the account is still up, here or on another BNG.
RADIUS not answering	The RADIUS server or the route to it (`show system health`).
IP pool full	The pool: add addresses or free stale sessions.
service name refused / VLAN not served	The configuration: the device asks for a service name that is not configured on its port, or sends PADI on a VLAN that is not in the PPPoE configuration (seen by the watcher, never answered).
no PPPoE discovery from the device since ...	The device or the last mile: nothing arrives from that MAC. The BNG would answer; check the device's PPPoE client, its WAN cable and the switch port.
PADI seen, no PADR	The device got the PADO and stopped: usually a device restarting or a second PPPoE server on the segment.
held back by FastPADO	The device is banned for flooding or repeated rejects (`show service pppoe-server fastpado banned`); it is released automatically.
the last session ended ... with <cause>	The cause names who ended it: Lost-Carrier (link or device dropped), User-Request (device logged out), Admin-Reset/NAS-Request (disconnected from the BNG or RADIUS, or the device re-dialled and replaced its own session), NAS-Reboot (the BNG restarted), Session-Timeout / Idle-Timeout (RADIUS limits).

Discovery is watched by a small service on the PPPoE access ports; it keeps the last 24 hours per device and costs no measurable CPU. The box keeps one day of the full connection log (it is emptied every morning at 05:18, and earlier if it passes 1 GB); before that, the last attempt of every user and device is kept for two days in a small digest, so the reason for last night's failure is still there in the morning. `show subscriber` says how far back its records go; `failures` counts only the current day's log in full. The lookup takes about a second on a busy box. The web panel has the same text under Logs, Subscriber Logs.

NAT (SNAT / DNAT)

```
set system snat source-address-pool <private pool> translation pool <public pool> out-interface
<wan-port>
```

This is the carrier-grade NAT (CGNAT) of the box: subscribers in the private pool leave the WAN with addresses from the public pool (persistent: a subscriber keeps the same public address). The out-interface **must be the WAN port name of this box** (for example `eth0-40G`); a rule on a port that does not exist never matches and the users get no internet.

```
set system snat source-address-pool <private pool> translation ip <public ip> out-interface <wan-
port>
```

```
set system snat source-address-pool <private pool> translation pool <pool>,<pool> out-interface
<wan-port>
```

```
delete system snat source-address-pool <private pool> translation pool <public pool> out-
interface <wan-port>
```

```
set system snat destination-address-pool <pool> translation ip|pool <...> out-interface <port>
```

```
set system dnat destination <ip|any> source ip|pool <...> protocol <tcp|udp> port <n> to-port <n>
```

```
show system snat   |   show system dnat
```

Ratio guidance: up to about 128 subscribers per public address is comfortable; `show system health` warns above that. The NAT guard drops subscriber packets that would leave un-NATed and reports them in the health check ("un-NATed private source"), which is how a missing rule shows up.

NAT session limit per user new

```
set system snat nat-session-per-user limit 5000
```

```
delete system snat nat-session-per-user
```

Caps the number of simultaneous NAT connections one subscriber address may hold (default 5000; `0` turns it off). New connections above the cap are refused until old ones expire; existing flows are untouched. A normal household uses a few hundred; an infected device or a scanner opens tens of thousands and would otherwise fill the table for everyone. Subscribers at the limit appear in the web panel's **DDoS WATCH** card and in `show system health`.

QoS: destination-group speeds

Give chosen destination networks (an IX, a cache, a local server farm) a different speed than the subscriber's plan, without touching the plan itself.

```
set policy qos group <name> network <network/mask>
```

```
set policy qos group <name> speed all 200M
```

Every subscriber gets 200 Mbit/s towards these networks (one value = same download and upload).

```
set policy qos group <name> speed plus 50M
```

Plan + 50 Mbit/s towards these networks.

```
set policy qos group <name> speed multiply 2
```

Twice the plan towards these networks.

```
set policy qos group <name> speed all download 200M upload 100M
set policy qos group <name> speed plus download 100M upload 50M time 00:00 to 06:00
```

Separate download and upload values, and an optional time window (24 h clock; default `time any`). Rates are written like 200M, 51200k or 1G; `multiply` takes a factor such as 2 or 1.5.

```
delete policy qos group <name> network <network/mask>   |   delete policy qos group <name>   |
delete policy qos
```

```
show policy qos   |   show policy qos group <name>   |   show policy set
```

```
set policy qos group <name> port <1-65535>
```

```
set policy qos group web port 443
set policy qos group web port 8080
set policy qos group web speed all download 20M upload 20M
```

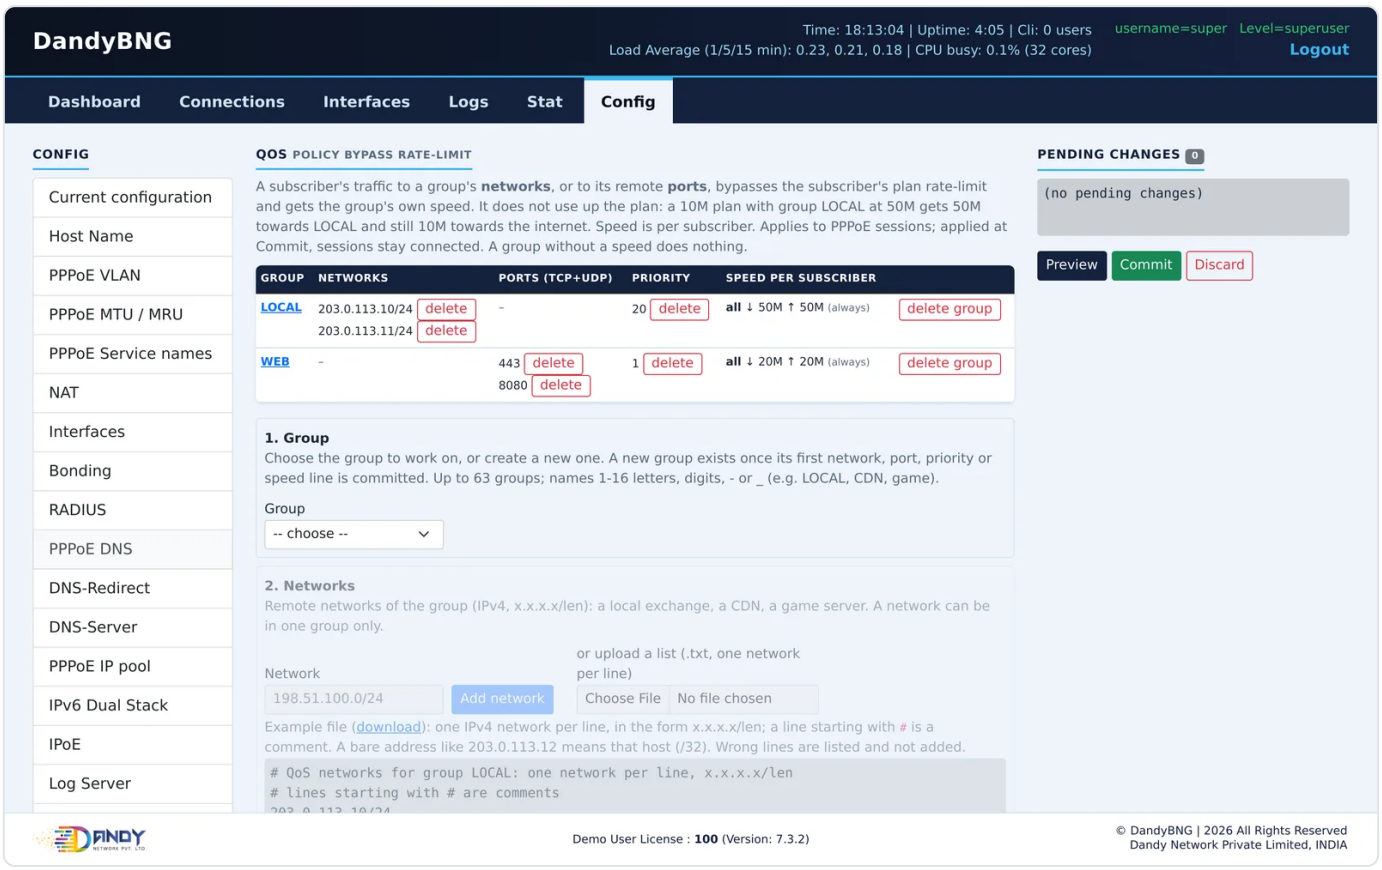
Port groups: traffic to and from these TCP or UDP ports (UDP 443 is QUIC) gets the group's speed per user, on its own lane next to the plan. The port is the remote side's port. A network group wins over a port group: traffic to a network in another group keeps that group's speed. A port can be in one group only; up to 256 ports. IPv4 only, like network groups.

```
delete policy qos group <name> port <1-65535>
```

```
set policy qos group <name> priority <1-63>   |   delete policy qos group <name> priority
```

Optional. When a destination is in a network group and its port is in a port group, the group with the **lower** priority number wins. Without priorities (or on a tie) the network group wins, so a port group never changes what a network group already allows. Example: `web` (ports 443/8080, 20M) priority 1 and `LOCAL` priority 10 gives HTTPS to local networks the 20M web lane instead of LOCAL's speed; `commit` prints a notice when the winning group is the slower one.

A group is active as soon as it has a speed; delete the group to switch it off. `commit` applies changes to every connected user at once, without a re-login or restart. Up to 63 groups.



Config > QoS: groups with networks, ports, priority and speed, built in numbered steps.

Static routes and the IX route policy

```
set ip route <network/mask> gateway <address> out-interface <port>
```

```
show ip route
```

```
set policy route ix network <network/mask>   |   network6 <prefix>   |   address <ip>
```

Destinations reached through a local exchange; used by the legacy per-connection `ix qos cbq` settings below. Prefer the QoS groups above for new deployments.

```
set policy route ix qos cbq all <rate>   |   plus <rate>   |   multiply <factor>
```

```
show policy route ix
```

CLI Users and Levels

```
set system login user admin authentication password <new password>
```

```
set system login user <name> authentication password <password>
```

```
set system login user <name> level user   |   admin   |   superadmin
```

`user` can only show; `admin` configures; `superadmin` can also enter routing mode.

```
delete system login user <name>
```

```
show system login
```

The user `admin` always has the level superadmin and cannot be deleted; its password can be changed. A name that the operating system already uses (root, daemon ...) is refused. The same work can be done in the web: **Config > Users**.

Log Server (NAT logging)

Every NAT translation (subscriber, private address, public address and port, time, user name) is exported in real time. This is a legal requirement for CGNAT; configure it before the first subscriber.

```
set system log-server address <collector> port <port> netflowv9
```

```
set system log-server address <collector> port <port> ipfix
```

```
set system log-server address <collector> port <port> ipfix all   |   netflowv9 all
```

`all`: when a NAT entry closes, the record also carries the connection's total bytes and packets (useful for usage questions); without it only the open/close events are sent. Re-enter the command without `all` to go back.

```
set system log-server address <collector> port <port> syslog
```

```
set system log-server address <a>,<b> port <p1>,<p2> netflowv9
```

```
set system log-server source-ip-pool <network/mask>   |   source-ipv6-pool <prefix>
```

Pools whose traffic must be logged (normally your private pools). The last records sent are visible on the web panel's Stat tab ("Log Server") to confirm the export is flowing. NetFlow v9 or IPFIX are recommended; each record carries the user name.

Local Users (without RADIUS)

```
set service pppoe-server local-user username <user> password <pass> speed download 100M upload
100M
```

```
set service pppoe-server local-user username <user> password <pass> speed download 100M upload
100M address <ip>
```

```
delete service pppoe-server local-user username <user>
```

```
show service pppoe-server local-user
```

A local user is answered locally; no RADIUS request is sent for that name.

SNMP (PRTG / Zabbix / LibreNMS) reworked

The agent exposes exactly what a monitoring system graphs on a BNG: system (uptime, name), **physical ports, VLANs and bonds** with 32/64-bit counters, CPU, memory, disks and temperatures. PPPoE sessions are deliberately not listed (an agent walking thousands of ppp rows stops answering).

```
set system snmp community <string>
```

```
set system snmp allow <monitoring host or network>
```

```
set system snmp location <text>   |   contact <text>
```

```
set system snmp enable   |   disable
```

```
set system snmp port <number>              (default 161)
```

```
set system snmp version 1 | 2c | 3          (default 2c)
```

```
set system snmp listen-address <ip>         (default: all addresses)
```

```
set system snmp v3 user <name>
set system snmp v3 auth sha <password>      (or md5)
set system snmp v3 priv aes <password>      (or des)
```

```
set system snmp trap-target <ip> community <string>
```

Traps: cold start and authentication failures. Link up/down traps are deliberately off (thousands of ppp interfaces would flood the NMS).

```
show system snmp
```

Prints the configuration, the sockets the agent listens on, and "requests received / replies sent" since start. If requests stay at zero while the NMS reports "no response", the NMS traffic is not reaching the box (allow list, firewall, wrong source address).

PRTG: add the device with SNMP v2c, the community above, port 161; use SNMP Traffic on the WAN, PPPoE ports and VLANs, SNMP CPU Load, SNMP Memory, SNMP Disk Free.

Web Panel

```
set system web address <wan ip> port <port>
```

```
set system web login user <name> authentication password <password>
```

```
set system web login user <name> interface <port> vlan-id <number>
```

A user restricted to one interface/VLAN sees only those subscribers.

```
delete system web login user <name>
```

```
restart system web
```

Web users and levels V7.3.2

```
set system web login user <name> level viewer | operator | superuser
```

```
delete system web login user <name> level      (back to viewer)
```

```
show system web login
```

`viewer` sees Dashboard, Connections, Interfaces, Logs and Stat. `operator` may also disconnect a user or an interface and lift a FastPADO ban. `superuser` also gets the **Config** tab (Configuring from the web). A user created with a password only is a viewer. Web users are added, changed and deleted by a CLI superadmin or, in the web, by a superuser in **Config > Users**. Passwords must be 6 characters or more and are stored as salted hashes; accounts from earlier versions keep working with their old hash until the password is set again (`show system web login` marks them "old-hash").

Config > Users new (superuser only): two lists, CLI users and Web users, each with Add, Change password, Change level and Delete. A change is applied at once and written to the audit log (never the password). The web user `super` always has the level superuser and cannot be deleted, in the web or in the CLI; the same holds for the CLI user `admin` (always superadmin). You cannot delete the user you are logged in as, or change your own level. A deleted web user is logged out at the next click; a new level counts at once, without a new login. Passwords in this page: 6 to 64 characters, letters, digits and `! # % ( ) * + , - . / : = ? @ [ ] ^ _ { } ~`.

All web passwords lost? Log in to the CLI as `admin` and set one again: `set system web login user super authentication password <new password>`, then `commit`. Factory web user: `super` (superuser); the old factory user `test1` is gone.

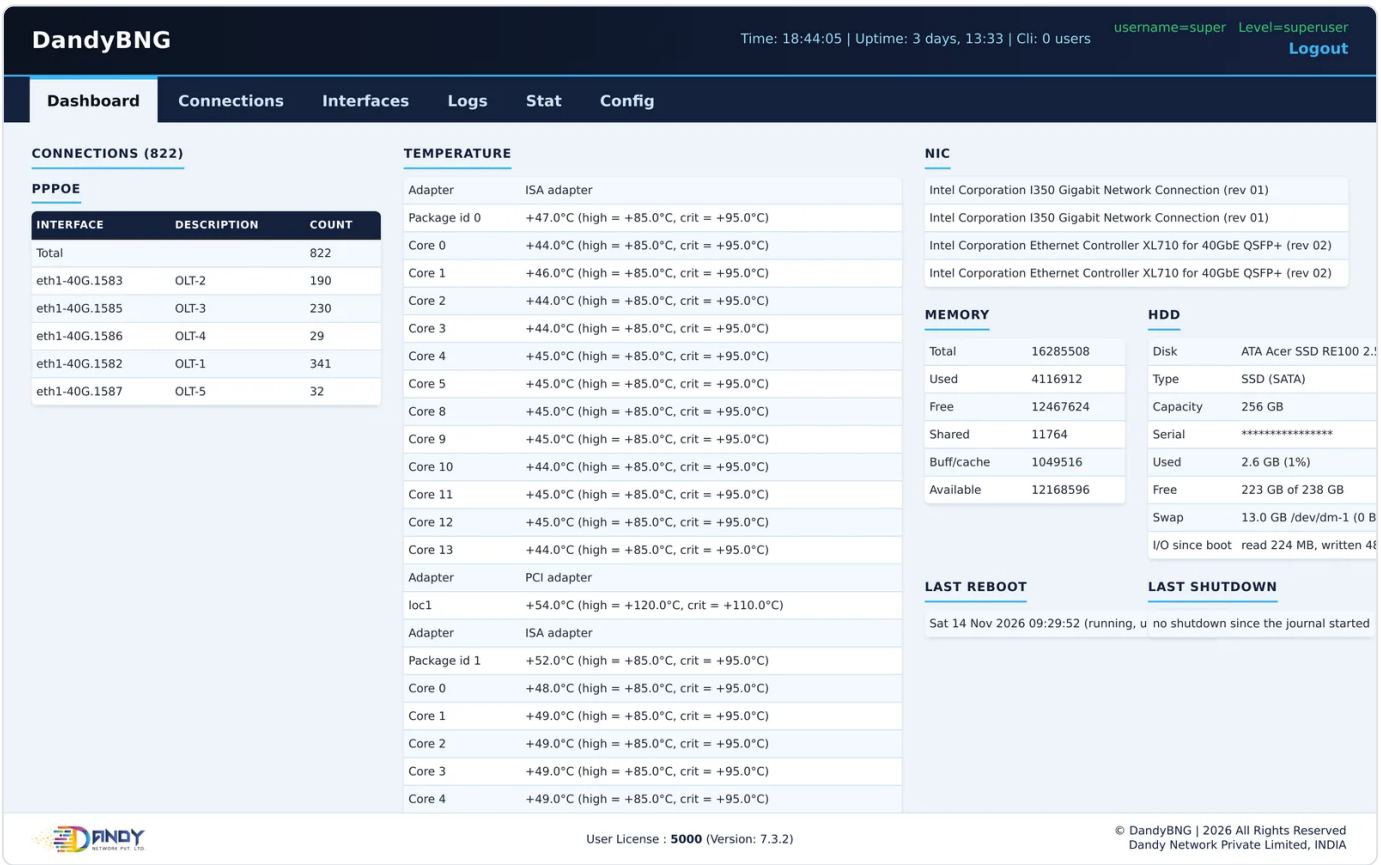
Five failed logins from one address lock that address for five minutes. The header shows **username=... Level=...** above Logout. Logins, failures and every web action are logged (web actions are shown in Config > Audit Log) and follow the 7-day log retention.

Pages: **Dashboard** (CPU, memory, NIC, **HDD** make/model/capacity/used/free, last reboot and last shutdown), **Connections** (search by Username, MAC, IPv4, IPv6, Interface, Service-name or IFNAME (ppp0, ppp1 ...); disconnect; live per-user traffic window), **Interfaces**, **Stat**: PHYSICAL PORTS (optic type, light levels, errors, link drops, PCIe width, slot/NUMA advice), System Health, NAT SESSIONS and Top 20 NAT users, DDoS WATCH (suspect flooders, inbound floods, users at the connection limit), Top 20 download / upload / international / data usage, TRAFFIC SPLIT, Log Server (last records exported), **Tools**, **Logs**.

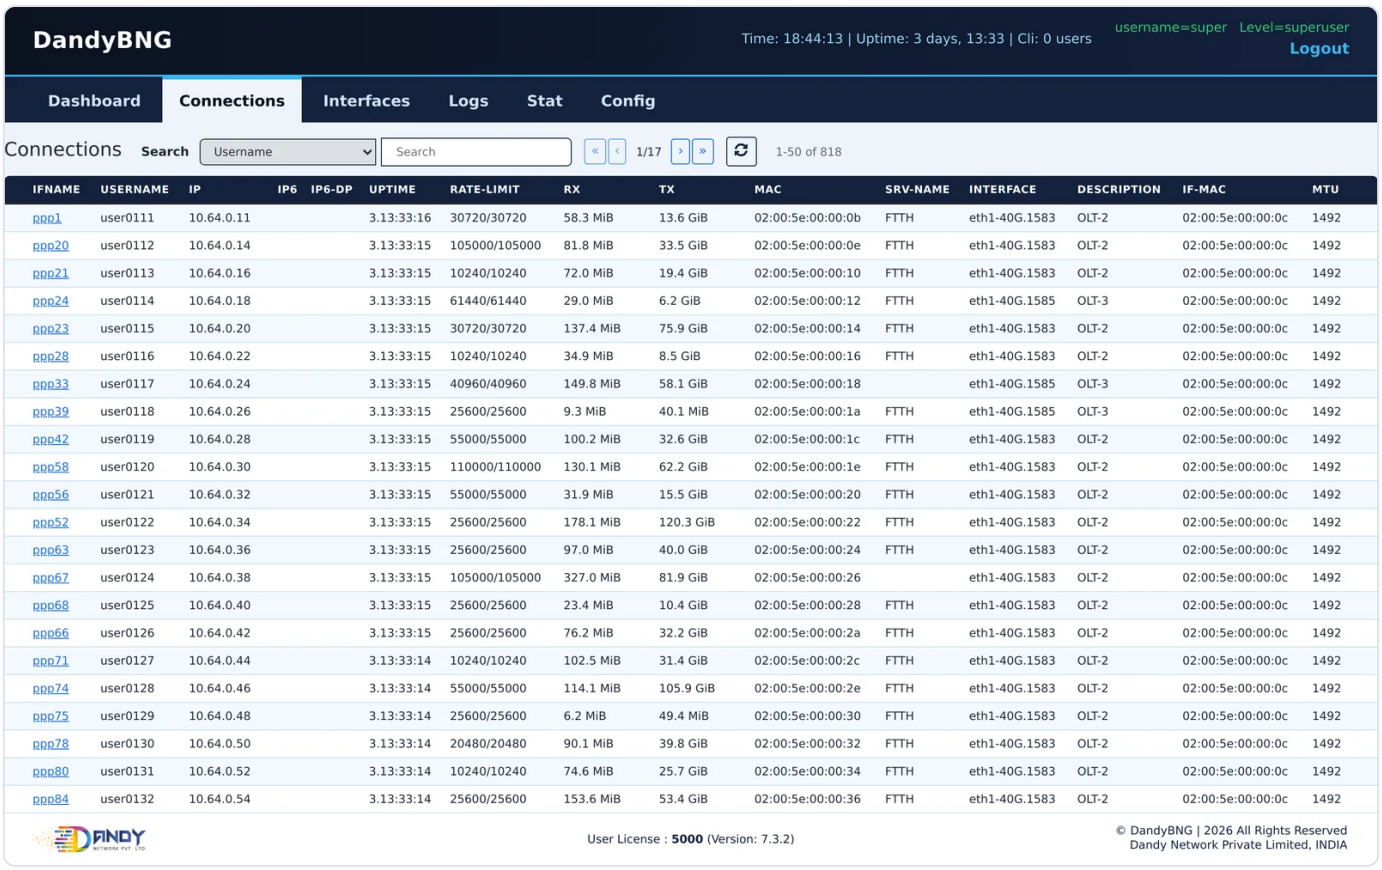
Logs tab

One line of controls: **Search logs**, **Pause Logs**, then the filters **ALL** (default), **Failed**, **Account expired**, **Success**, **LCP time-out**, **Service-name mismatch**, then **Subscriber Logs** with its username / MAC / IP box. A filter shows only its lines, fetched from the server (the last 1,000 matching lines of the last 10 MB of the log, not only what is on screen), with the count in the button; new lines keep streaming under the same filter. **Subscriber Logs** shows the same text as `show subscriber` for the name, MAC or IP typed in the box (Enter works too), with a "Back to logs" link.

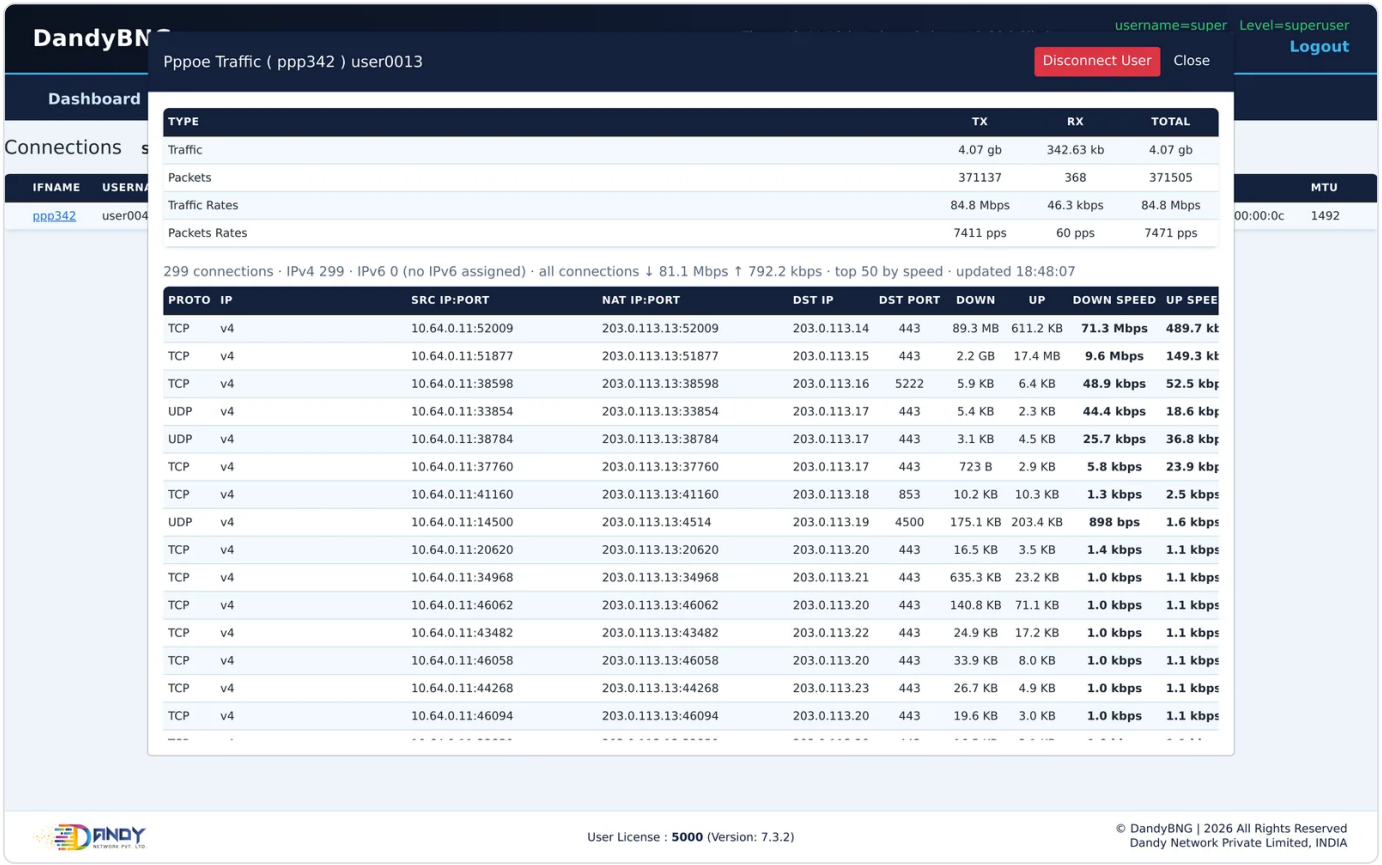
Every line that can be attributed ends with the username: RADIUS requests from their User-Name, RADIUS replies paired to their request by channel and id, lines that only name a MAC (Service-Name mismatch, FastPADO) with the MAC's last login name as "(last login as)" or "no login seen from this device". "Mikrotik-Rate-Limit" is shown as "Rate-Limit"; the "connect: pppN <--> pppoe(MAC)" lines are not shown. The Search box and the filters see the suffix, so typing a username shows that user's whole story, rejects included.



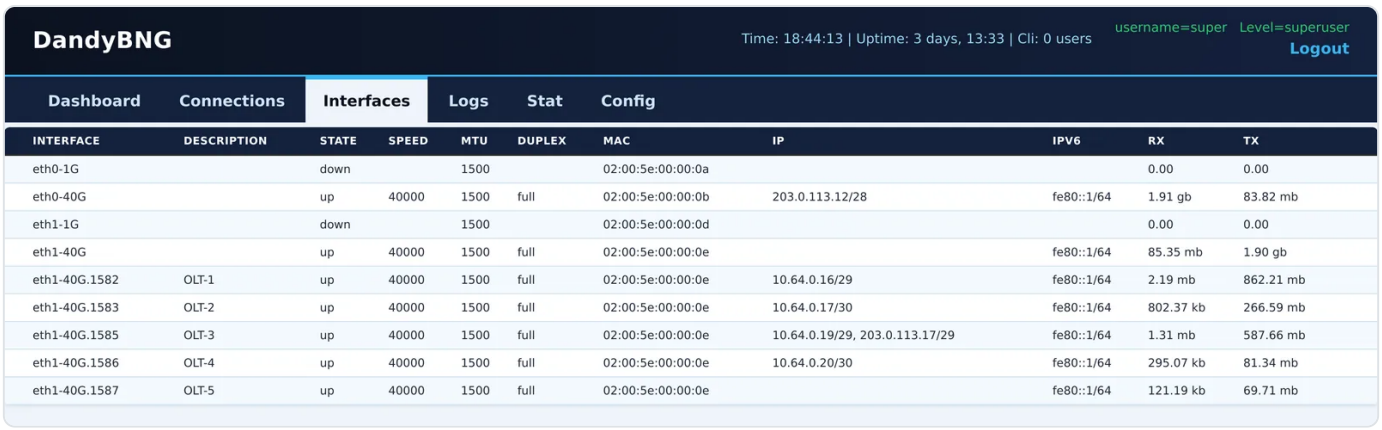
Dashboard: sessions per port and VLAN, temperatures, network cards, memory and disk.



Connections: every online subscriber; click an IFNAME (ppp...) to see that subscriber's live traffic.

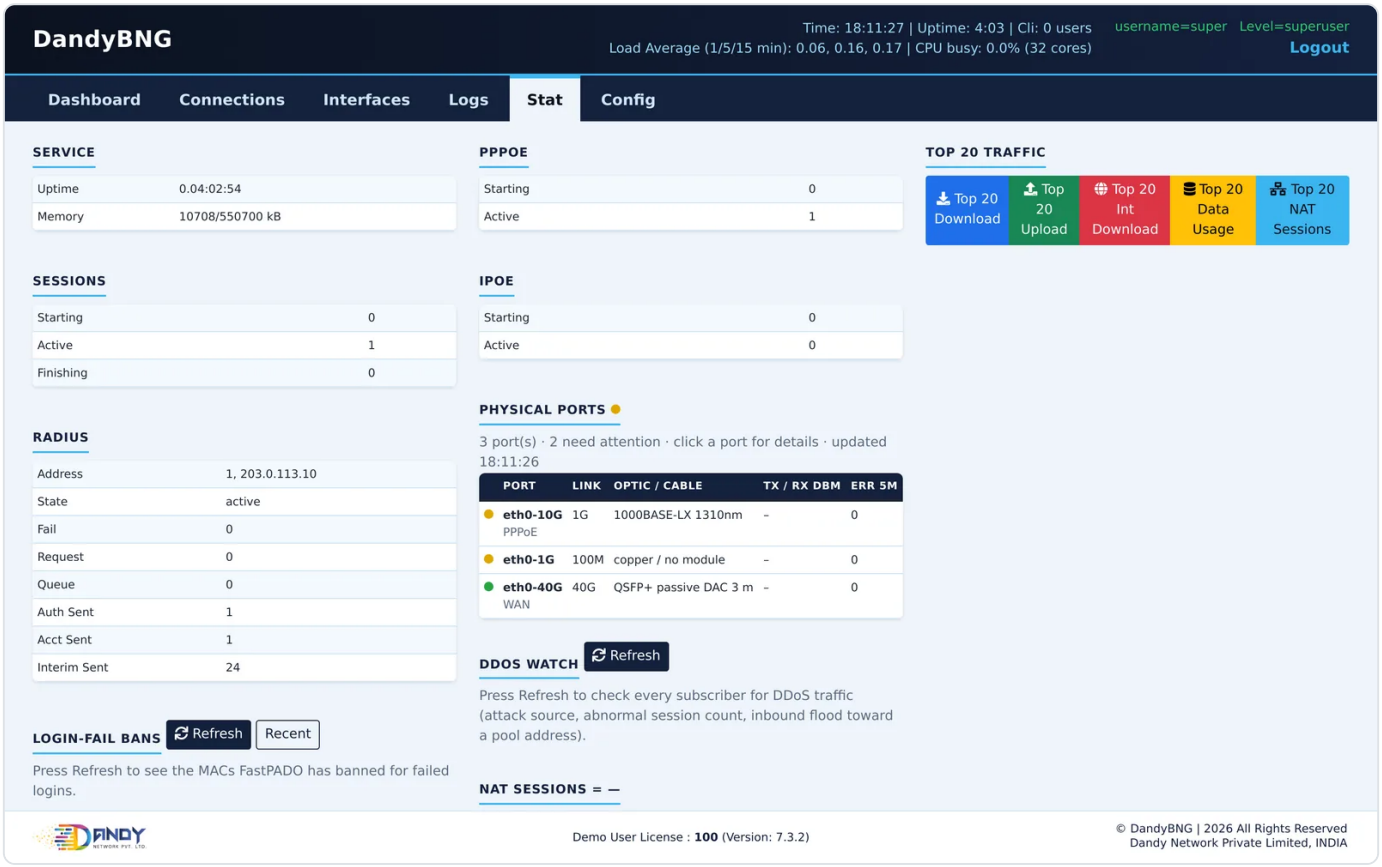


Live traffic of one subscriber: rates every 2 seconds and every open connection with NAT address, destination and speed.



INTERFACE	DESCRIPTION	STATE	SPEED	MTU	DUPLEX	MAC	IP	IPV6	RX	TX
eth0-1G		down		1500		02:00:5e:00:00:0a			0.00	0.00
eth0-40G		up	40000	1500	full	02:00:5e:00:00:0b	203.0.113.12/28	fe80::1/64	1.91 gb	83.82 mb
eth1-1G		down		1500		02:00:5e:00:00:0d			0.00	0.00
eth1-40G		up	40000	1500	full	02:00:5e:00:00:0e		fe80::1/64	85.35 mb	1.90 gb
eth1-40G.1582	OLT-1	up	40000	1500	full	02:00:5e:00:00:0e	10.64.0.16/29	fe80::1/64	2.19 mb	862.21 mb
eth1-40G.1583	OLT-2	up	40000	1500	full	02:00:5e:00:00:0e	10.64.0.17/30	fe80::1/64	802.37 kb	266.59 mb
eth1-40G.1585	OLT-3	up	40000	1500	full	02:00:5e:00:00:0e	10.64.0.19/29, 203.0.113.17/29	fe80::1/64	1.31 mb	587.66 mb
eth1-40G.1586	OLT-4	up	40000	1500	full	02:00:5e:00:00:0e	10.64.0.20/30	fe80::1/64	295.07 kb	81.34 mb
eth1-40G.1587	OLT-5	up	40000	1500	full	02:00:5e:00:00:0e		fe80::1/64	121.19 kb	69.71 mb

Interfaces: ports and VLANs with state, speed, MTU, addresses and traffic.



Stat: sessions, RADIUS counters, physical ports, DDoS watch, NAT sessions and Top 20 buttons.

Traffic Split and Top organisations redesigned

Stat tab, button **Traffic Split**. Nothing is computed in the background; each click measures and shows:

- **Local (Local AS<number>)**: traffic to prefixes announced by your own AS (detected from the WAN address, refreshed daily). If the detection is wrong, Dandy support can set your AS number or a prefix list by hand.
- **<COUNTRY> (within <country>)**: destinations that answer from inside the country, judged by round-trip time from the BNG (below 60 ms by default, adjustable in the same file). Content caches of Meta, Akamai, Google and CloudFront hosted in-country land here even though their addresses are registered abroad.
- **International (transit outside <country>)**: everything farther away.
- **Top 10 within <country>** and **Top 10 International**: organisations (Google, Meta, Akamai via your upstream, Netflix ...) with share and bandwidth.

The country name comes from the system time zone. Shares are byte-weighted over all live NAT flows; bandwidth is the WAN rate since the previous click. The first click after a restart takes 10 to 20 seconds while destinations are measured.

Traffic Split

REGION	SHARE %	BANDWIDTH
Local (Local AS64500)	35.35%	532.47 Mbps
INDIA (within India)	63.05%	949.70 Mbps
International (transit outside India)	1.6%	24.05 Mbps
Total	**100%**	**1.51 Gbps**

TOP 10 WITHIN INDIA

ORGANISATION	AS	SHARE	BANDWIDTH
Datacamp Limited	AS212238	84.87%	805.98 Mbps
Airnet Cable And Datacom Pvt Ltd	AS133001	5.4%	51.32 Mbps
Amazon (CloudFront/AWS)	AS16509	2.32%	22.07 Mbps
Google	AS15169	1.89%	17.98 Mbps
Meta (Facebook)	AS32934	1.25%	11.89 Mbps
Reliance Jio	AS55836	0.72%	6.86 Mbps
Threesa Infoway Pvt.Ltd.	AS132996	0.61%	5.82 Mbps
Cloudflare	AS13335	0.51%	4.87 Mbps
Google (via Google Cloud)	AS396982	0.48%	4.52 Mbps
Fastly	AS54113	0.37%	3.54 Mbps

Stat > Traffic Split on a production server: own AS, within India, international, and the top 10 organisations.

System health and auto-tune

```
show system health
```

PASS / WARN / FAIL lines with the fix for each. Also on the web panel.

```
set system auto-tune apply   |   status   |   off
```

Auto-tune sets NIC rings and offloads, pins interrupts to the card's CPU socket, sizes the connection table and re-applies at every boot and commit. `apply` is only needed after changing hardware. A port that becomes a PPPoE or IPoE interface is tuned at the commit that adds it.

Card fitted after the installation new: its ports first show plain names (`eth0`, `eth1` ...). Auto-tune gives them the product names (`eth2-10G`, `eth3-10G` ..., the next free number of that speed) at the next boot, at a `commit` or with `set system auto-tune apply`; no reboot is needed. Ports that already have a name never change, and `set system auto-tune status` shows a "port names" line.

```
set system auto-tune irq-spread auto   |   on   |   off
```

```
set system auto-tune napi-defer on   |   off
```

Interrupt suppression on the dataplane ports (default `on`): after an empty poll a port is re-polled by a 20 µs timer up to 20 times before its interrupt is re-armed. Measured in the lab at the same throughput: 90 percent fewer network interrupts, 42 percent less packet-processing CPU. Live, no restart.

```
set system auto-tune idle c1   |   poll
```

What idle cores do. `c1` (default) lets them halt: about 25 W less per CPU socket and a much cooler CPU, same forwarding cost, wake-up under 2 µs. `poll` keeps them spinning: zero wake-up delay, every core at the all-core turbo clock, but the CPUs run hot even with no traffic. Deeper sleep states are never used on a BNG. The change is written to the boot line and takes effect at the next reboot; `auto-tune status` and `show system health` show "reboot pending" until then.

Where the network interrupts run. `auto` (default) keeps them on each card's own CPU socket, and spreads them over all cores only when every network card sits on one socket of a two-socket server. `on` always uses all cores, `off` always stays socket-local. Applies at `commit`, sessions stay up.

The health check also reports a first-boot setup error (code F21/F22 - contact support) and network ports that have no driver or firmware; both come from the installer and are shown at every login until fixed.

```
set system upload-mode police   |   htb
```

```
show system upload-mode   |   show system cpu   |   show system status
```

Upload shaping: `police` (default, per-user token bucket, scales to thousands of users at 10G and above) or `htb` (per-user classes, legacy).

Backup / Restore of the configuration new

A backup is one text file with the whole configuration and the list of the server's ports. Backups are kept on the server, can be downloaded and can be restored on the same server or on another one. A copy is also saved **by itself at every commit** (the last 20 are kept) and **before every restore** (the last 5 are kept); backups you make or bring in yourself stay until you delete them.

```
backup create [note]   |   backup list   |   backup show <name|number>
```

```
backup export <name|number>   |   backup import [name]   |   backup delete <name|number>
```

`export` prints the file to copy it to another place; `import` takes a pasted file (finish with a line that holds only `end`). The number is the one `backup list` shows. The file holds RADIUS secrets and the SNMP community: keep it in a safe place.

```
backup restore <name|number> [logins]
```

The restore works in steps and changes nothing before the last `yes`:

1. **Ports.** The ports the backup names are compared with the ports of this server. For a port that is not here the CLI shows the ports of this server (speed, link, what each is used for now, hardware address) and asks which one to use; `Enter` takes the suggestion, `s` leaves the lines of that port out, `q` stops. VLANs, bond members and service names follow the chosen port.
2. **Check.** Every line is checked by the CLI itself. Lines this server does not accept are listed with the reason; you can go on without them or stop.
3. **What changes.** The lines that will be added and removed are shown. Only lines that differ are touched; sessions on parts that do not change stay up.
4. **Restart question.** When PPPoE DNS, an IP pool or the DAE listener changes, the CLI tells how many subscribers are online and asks whether to restart the PPPoE / IPoE server as part of the restore. Without `yes` nothing is restarted.
5. **yes.** The running configuration is saved first as a before-restore copy, then the backup is applied.

```
backup keep   |   backup return   |   backup status
```

Safety timer. When a restore changes interface addresses, routing or the web address, it is on trial for **5 minutes**: `backup keep` (or the **Keep** button) makes it final; without it the configuration from before returns by itself, so a wrong address cannot lock you out. A reboot inside the 5 minutes does not lose the return. `backup return` goes back at once.

Not changed by a restore: the high-availability settings and the CLI users of the server, and the web / API logins unless you add `logins` (web: the tick box). The session license is bound to the hardware and is not part of a backup. A restore is refused while changes are waiting for commit.

From the web: Config > Backup / Restore (below System Health, superuser only) shows the list of backups with **Make a backup now**, **Bring in this backup file** (a file from your computer), and for each backup **View**, **Download**, **Restore** and **Delete**. Restore walks through the same steps: a table of the missing ports with a list of this server's ports for each, the lines that change, the restart tick

box, and the word `Yes` typed before anything is done. While a restore is on trial the page shows the time left with **Keep the restored configuration** and **Go back now**.

DNS Cache Server

```
set system dns enable   |   disable
```

```
set system dns address <wan ip>
set system dns address 127.0.0.1
set system dns port 53
set system dns allow-pool <private pool>
set system dns allow-pool <wan ip>/32
set system dns threads <cores>
set system dns forwarder <upstream resolver>
```

```
set service pppoe-server dns-servers dns1 <wan ip>
```

```
restart system dns
```

DNS redirect

Sends every subscriber DNS query (port 53) to the resolver you name, whatever the subscriber configured. Off until you enable it explicitly.

```
set system dns-redirect target <resolver ip>
set system dns-redirect state enabled
```

```
set system dns-redirect state disabled   |   delete system dns-redirect
```

The target is usually this box's own DNS cache (the WAN address with `set system dns enable`) or your central resolver.

Ookla Speedtest Server

```
set system speedtest server enable   |   disable
```

```
set system speedtest domain speedtest.yourdomain.in
```

Register the box's WAN address with Ookla as a host (Ookla host guide). The server starts at every boot. With a `domain` set, `commit` checks for an SSL certificate for that name and obtains one from Let's Encrypt if missing: the name must already point to this box (A and, if present, AAAA record), port 80 must be free, and the box needs internet access. The certificate renews itself; nothing to repeat. Without a domain the server keeps Ookla's own certificate, which is enough for the speedtest.net listing.

```
delete system speedtest domain
```

Session license

Each box is limited to the number of simultaneous subscriber sessions its license allows; without a key it runs as a demo with 100 sessions. Keys are bound to the hardware and signed by the vendor.

```
show system license
```

Prints the fingerprint to send to the vendor, the allowed session count and the temporary key window if any.

```
set system license key <key-string>
```

```
delete system license key   |   temp-key
```

A new limit applies within a minute, without any restart; a Telegram message is sent at every change when the WAN monitor is configured. Sessions above the limit are refused, existing ones are never cut.

Installing from the ISO

1. Write the ISO to a USB stick in **DD mode** (Rufus: choose "DD Image" when asked; or Etcher, or `dd`), or use the server's virtual media. A stick written in Rufus "ISO mode" boots but the installer cannot read it.
2. Boot it in **Legacy/BIOS (CSM) mode**, the proven mode. UEFI works as well when **Secure Boot** is disabled. No network is needed during the installation.
3. The **System check** screen appears after about a minute (see below). When everything is fine it continues by itself after 30 seconds, or at any key.
4. The installer picks the **largest fixed disk** (at least 40 GB) and erases it. It shows the step it is in, a timer and "Do not power off"; aprox 10-15 minutes.
5. After "Installation complete" remove the media. The system restarts aprox 3 times (aprox 5 minutes) showing "Finishing setup - step 1 of 2 / 2 of 2" under the logo, then the login prompt.
6. Log in as **admin**, change the password, follow the first-time setup checklist.

System check screen

Before the disk is touched, the installer checks the machine and shows the result. Warnings (orange "!") are advice and never stop the installation. A blocker (red "X") stops it with a code and the exact change to make; fix it and press **R** to check again (a cable, another stick, a plugged-in disk), or restart after a BIOS change.

Item	Warning (continues)	Blocker (stops)
Boot mode	UEFI: Legacy/CSM is the proven mode, enable CSM in the BIOS boot options if available	
Secure Boot		E12 enabled: BIOS/UEFI setup, Security, Secure Boot: Disabled
Install media		E13 the installer cannot read its own files: write the stick again in DD mode, try another USB port, or use virtual media
Target disk		E11 no fixed disk of 40 GB or more; the advice names what was found: SATA mode RAID/RST (set AHCI), a RAID/SAS card without a virtual disk (create one, or HBA/IT mode), IDE mode (set AHCI), a disk that is too small. USB disks are never used.
Processor	fewer than 4 threads	
Memory	below 8 GB: 16 GB or more recommended	E14 below 2 GB
Network ports	only one port (a BNG needs WAN + subscribers); a port whose firmware is not on the disc, named	E15 no port with a driver
Clock	wrong year: set the BIOS clock	
Machine	vendor, model, BIOS version and date, for support photos	

"Installation needs your attention"

The installer runs fully unattended and never asks a question on its own screen. When its log shows no progress for 5 minutes (code **E30**), or a step fails (**E20**), this screen shows the step, the last lines of the installer log, what to do, and how long the log has been silent. Keys: **Alt+F4** shows the live installer log, **Alt+F5** returns, **H** shows a help page with the keys and all codes, **Enter** restarts the machine. The screen clears by itself when the installation continues. The common cases are explained on the screen itself: old data on the disk stopping the partitioning (restart; the installer erases the first 2 GB and the last 1 MB of the target disk before partitioning, so a reinstall over a previous BIOS-mode or UEFI install is clean), a stick written in ISO mode (rewrite it in DD mode), a disk not seen (SATA mode AHCI or RAID volume), read errors on the stick or disk.

Codes: E11 no suitable disk, E12 Secure Boot, E13 media unreadable, E14 memory below 2 GB, E15 no network port with a driver, E20 an installer step failed, E22 to E25 a setup step failed, E30 no progress for 5 minutes. Photograph the screen and contact support with the code. A network port without driver or firmware is only reported ("Needs attention"), the installation continues. The check result is saved on the installed box.

Time zone

```
set system time-zone Asia/Kolkata
```

```
show date
```

Used for logs, the health check's clock test and the Traffic Split country name.

Routing Mode (BGP, OSPF, IS-IS)

```
set system routing
```

Superadmin only. Opens the router shell with Cisco-style commands. Do not change interface addresses from routing mode; use the CLI commands above.

IPoE Configuration

PPPoE and IPoE cannot share a physical port; use different ports or VLANs.

```
set service ipoe-server idle-timeout 60
```

```
set service ipoe-server interface <port> vlan-id <number> description <text>
```

```
set service ipoe-server interface <port> vlan-range <low-high>
```

```
set service ipoe-server network <network/mask> gateway <ip> interface <port> vlan-id <number>
```

```
show service ipoe-server interface
```

A host on an IPoE port gets nothing until RADIUS accepts it: traffic that goes straight through an IPoE port or VLAN without a login is blocked, and logged-in subscribers are not affected. Do not put an IPoE network on the IPoE port itself (an interface address or a static route there): logins on that network cannot start. A commit that would do this is refused, and `show system health` reports it in the lines **IPoE guard** and **IPoE network routing**.

High availability (HA) V7.3.3

Two boxes work as one pair. One box is **active** and serves all subscribers. The other is the **standby**: it holds a live copy of every subscriber session and of the connection (NAT) table, and keeps its WAN and subscriber ports closed. When the active box fails, is rebooted, is restarted or is switched over, the standby opens its ports and serves in about a second. **The subscribers do not re-dial**, RADIUS accounting goes on with the same session, and open connections continue.

What it needs

- two boxes with the same ports (same port names) and the same software version;
- one free port on each box, joined by a direct cable: the **HA link**. It must be a port of its own: not the WAN port, not a subscriber port, without an address and not in a bond;
- the WAN port and the subscriber ports of both boxes on the same switches and VLANs;
- a session license on each box for the number of subscribers of the pair (the copies on the standby count like sessions).

The pair uses one WAN address and one set of port addresses: the upstream router, the RADIUS server and the subscribers see one BNG. The standby does not answer on the WAN address; the active box answers there, also for the CLI and the web panel.

Turning it on

```
set system high-availability link <port>
set system high-availability role primary
commit
```

On the second box the role is `secondary`. High availability is on when both lines are committed; there is no separate enable. The role only decides which box serves when both start at the same moment: after a takeover the box that serves keeps serving, it does not move back by itself.

```
set system high-availability key <text>
```

Optional pair key (8 to 64 characters: letters, digits and `_ . @ -`). The two boxes accept each other only with the same key; without the line a built-in key is used. Set the same text on both boxes.

Order for a new pair

1. **The box that serves now** (primary): set the link port and the role, `commit`. It keeps serving as before. High availability starts with one `restart service pppoe-server`: its subscribers re-dial this one time, so do it in a quiet hour. A box without subscribers starts by itself at the commit.
2. **The second box**: connect the HA link cable. Set its WAN address and gateway, the link port, the role `secondary` (and the same pair key), `commit`. It takes the rest of the configuration from the serving box by itself and becomes the standby.
3. Connect the WAN and subscriber cables of the second box and wait until `show high-availability` says **READY** on both boxes.

State

```
show high-availability
```

```
High availability : READY
This box          : primary, ACTIVE since 16 Nov 01:48:36, 37 subscribers online
Other box         : secondary, STANDBY, holds 37 subscriber copies
Connections       : 46, copied to the standby
Configuration     : the same on the standby (last copy 16 Nov 01:45:49)
HA link eth2-10G  : up
Path through WAN  : seen
```

The first line is the answer: **READY** means the standby can take over without a re-dial right now. **NOT READY** comes with one reason, for example the other box is not reachable, the standby is still copying, or a port has no link on one box only. The lines below show both boxes, the HA link, the second path through the WAN switch, ports without link and the last change.

What makes the standby take over

Event on the active box	Result
Power loss, crash	Takeover in 1 to 2 seconds.
`reboot system`, `restart service pppoe-server`, PPPoE/IPoE service fails	The subscribers are handed over first; under half a second.
WAN port or a subscriber port loses its link	Handed over in about a second, when the other box has link on that port.
HA link cable alone is cut	No switch. The boxes still see each other through the WAN switch; the state shows NOT READY with the reason.

Two boxes never serve at the same time: a standby that cannot be sure the other box is down waits and says so in the state. A port that has no cable on both boxes is not counted.

Switch over by hand

```
high-availability switchover
```

Typed on the active box: moves the subscribers to the other box without a re-dial, for planned work on this box. Refused with the reason when the pair is not READY; `high-availability switchover force` moves anyway (subscribers the standby does not hold re-dial).

```
high-availability takeover force
```

Typed on a standby whose other box is switched off or unplugged and that did not take over by itself. Use it only when you are sure the other box does not serve.

Configuration copy

The box that serves owns the configuration. Make every change there; it does not matter whether that box is the primary or the secondary. At the `commit` the standby takes the change, usually within a few seconds, and commit says so:

```
high-availability: the standby box takes this change now (show high-availability)
```

- **A box that was switched off** takes the whole current configuration when it comes back as the standby: first the configuration, then the subscribers, then it shows READY.
- **If somebody changed a box while it was off the pair**, the configuration of the serving box wins. The replaced configuration is kept as a backup on that box (the last 20) and the copy is written to its HA log.
- **Changes on the standby are refused** with a message to make them on the box that serves, in the CLI and in the web panel, because the next copy would replace them.
- **Copied**: everything that is not in the next line: ports and VLANs, PPPoE and IPoE servers, IP pools, RADIUS, NAT, QoS, DNS, SNMP, routes, API and all other settings.
- **Not copied, set on each box (also on the standby)**: the high-availability lines (link port, role, key), the host name, the CLI and web logins, the license.
- **Settings that need a service restart** (IP pools, DNS for subscribers, RADIUS DAE): the standby restarts its own service by itself; no subscriber is touched. On the serving box they take effect at its next `restart service pppoe-server`, as always; on a READY pair that restart hands the subscribers to the other box first.

The `Configuration` line of `show high-availability` shows the copy on both boxes: the same (with the time of the last copy), copying, waiting for the HA link (the copy runs on the HA link only), or **NOT THE SAME** with the reason. While the copy runs or has failed the pair shows NOT READY and a planned switch over is refused; an automatic takeover after a failure still works. A failed copy is tried again at the next commit and every ten minutes; until it succeeds the standby accepts changes, so the reason can be repaired there.

IPoE on a pair

IPoE sessions on shared VLANs that start on the first packet (IPv4) are copied like PPPoE sessions. IPoE sessions started by DHCP and IPoE with IPv6 are not copied: after a takeover they log in again on their next packet.

Turning it off

```
delete system high-availability
commit
```

The box keeps serving on its own. Remove it on the standby first, with its WAN and subscriber cables unplugged: two boxes without high availability on the same switches would both answer.

From the web

Config tab > **High Availability (HA)** (superuser level), below QoS:

- **1. Settings of this box**: the committed settings.
- **2. Set up or change**: HA link port (only ports that can be used are offered; for the others the page says why not), role, pair key.
- **3. Live state**: the same lines as `show high-availability`, refreshed every few seconds, with a **Restart PPPoE/IPoE server** button while the first start is pending.

- **4. Configuration copy**: the state of the copy in one coloured line, and what is and is not copied.
- **5. Switch over / take over**: the same operations as in the CLI, each with a confirmation and an Audit Log entry.

Displaying Configuration

```
show configuration set
```

```
show interfaces   |   show interfaces ethernet <port>   |   show interfaces bonding <name>   |
show interfaces pppoe
```

```
show interfaces connections
```

Number of PPPoE sessions per interface and VLAN.

```
show connections   |   show connections <username>   |   show connections log
```

```
show subscriber <username|mac|ip>   |   show subscriber failures [last <minutes>]
```

Why is this user not connected: verdict, what to do and the timeline of the last attempt. See Why is a user not connected.

```
show traffic interface <port> bandwidth
```

```
show system status   |   show system health   |   show system cpu   |   show system host-name
```

```
show service pppoe-server   |   interface   |   ip-pool   |   ipv6-pool   |   authentication
radius server   |   dns-servers   |   service-name interface   |   local-user
```

```
show system snat   |   show system snmp   |   show policy set   |   show arp ip   |   show ip
route
```

Maintenance commands

```
disconnect user <username>   |   disconnect ifname <pppN>
```

```
restart service pppoe-server interface <port> vlan-id <number>   |   restart service pppoe-server interface <port>
```

Disconnects the users of one access VLAN (or of the whole port) and lets them log in again; the VLAN itself stays in service. With users online the command shows their number and asks for `yes` first. Waiting packet sizes are applied on the way.

```
restart system snat   |   restart system dnat
```

```
restart system web   |   restart system dns   |   restart system snmp
```

`restart system web` restarts the panel in about a second. A port name that does not exist on the box (for example a typo like `etth1-10G`) is refused by `set interfaces ethernet` and `set service pppoe-server interface` at once, so it can never reach `commit`.

```
reboot system   |   shutdown system
```

```
commit confirm   |   commit discard
```

Command reference V7.3.2

Every command of the CLI, grouped by area. Words in `<angle brackets>` are values you type; `|` separates alternatives. Press Tab in the CLI to complete a word or list what may follow.

- `set` and `delete` only stage a change. Nothing happens until `commit`; `commit discard` throws staged changes away.
- A command that is incomplete or wrong (for example `set system log-server address 103.5.5.10 port 8000` without `netflowv9`, `ipfix` or `syslog`, or a port name that does not exist) is refused with "Something went wrong" and is never committed.
- `show`, `restart`, `disconnect` and the tools work at once, without commit.
- Some commands need the `admin` or `superadmin` level; see CLI Users and Levels.

Interfaces (details)

```
set interfaces bonding <name> address <ip/mask>
set interfaces bonding <name> downdelay <value>
set interfaces bonding <name> gateway <ip>
set interfaces bonding <name> hash-policy <value>
set interfaces bonding <name> ipv6-address <ipv6/len>
set interfaces bonding <name> ipv6-gateway <ipv6>
set interfaces bonding <name> lacp-rate <value>
set interfaces bonding <name> member interface <port>
set interfaces bonding <name> miimon <value>
set interfaces bonding <name> mode <value>
set interfaces bonding <name> updelay <value>
set interfaces ethernet <port> address <ip/mask>
set interfaces ethernet <port> ddp disable
set interfaces ethernet <port> ddp enable
set interfaces ethernet <port> description <text>
set interfaces ethernet <port> dns1 <ip>
set interfaces ethernet <port> dns2 <ip>
set interfaces ethernet <port> gateway <ip>
set interfaces ethernet <port> ipv6-address <ipv6/len>
set interfaces ethernet <port> ipv6-dns1 <ipv6>
set interfaces ethernet <port> ipv6-dns2 <ipv6>
set interfaces ethernet <port> ipv6-gateway <ipv6>
set interfaces ethernet <port> mtu <number>
set interfaces ethernet <port> name <name>
delete interfaces bonding <name>
delete interfaces bonding <name> address <ip/mask>
delete interfaces bonding <name> downdelay <value>
delete interfaces bonding <name> gateway <ip>
delete interfaces bonding <name> hash-policy <value>
delete interfaces bonding <name> ipv6-address <ipv6/len>
delete interfaces bonding <name> ipv6-gateway <ipv6>
delete interfaces bonding <name> lacp-rate <value>
delete interfaces bonding <name> member interface <port>
delete interfaces bonding <name> miimon <value>
delete interfaces bonding <name> mode <value>
delete interfaces bonding <name> updelay <value>
delete interfaces ethernet <port> address <ip/mask>
delete interfaces ethernet <port> ddp
delete interfaces ethernet <port> description <text>
delete interfaces ethernet <port> dns1 <ip>
delete interfaces ethernet <port> dns2 <ip>
delete interfaces ethernet <port> gateway <ip>
delete interfaces ethernet <port> ipv6-address <ipv6/len>
delete interfaces ethernet <port> ipv6-dns1 <ipv6>
delete interfaces ethernet <port> ipv6-dns2 <ipv6>
delete interfaces ethernet <port> ipv6-gateway <ipv6>
delete interfaces ethernet <port> name <name>
show interfaces bonding <name>
show interfaces connections
show interfaces ethernet <port> blink count
show interfaces ethernet <port> ddp
show interfaces pppoe
```

PPPoE server, RADIUS, pools, VLANs, FastPADO (details)

```
set service pppoe-server access-concentrator <name>
set service pppoe-server authentication radius acct-delay-time <seconds>
set service pppoe-server authentication radius acct-interim-interval <seconds>
set service pppoe-server authentication radius acct-interim-jitter <seconds>
set service pppoe-server authentication radius acct-timeout <seconds>
set service pppoe-server authentication radius attribute dandy
set service pppoe-server authentication radius attribute mikrotik
set service pppoe-server authentication radius dae-server <ip> port <number> key <secret>
set service pppoe-server authentication radius gateway-address <ip>
set service pppoe-server authentication radius max-try <number>
set service pppoe-server authentication radius nas-address <ip>
set service pppoe-server authentication radius nas-identifier <text>
set service pppoe-server authentication radius server <ip> acct-port <number>
set service pppoe-server authentication radius server <ip> fail-timeout <number>
set service pppoe-server authentication radius server <ip> key <value>
set service pppoe-server authentication radius server <ip> max-fail <number>
set service pppoe-server authentication radius server <ip> port <number>
set service pppoe-server authentication radius server <ip> req-limit <number>
set service pppoe-server authentication radius server <ip> weight <number>
set service pppoe-server authentication radius timeout <seconds>
set service pppoe-server dns-servers dns1 <ip>
set service pppoe-server dns-servers dns2 <ip>
set service pppoe-server fastpado disable
set service pppoe-server fastpado enable
set service pppoe-server fastpado login-fail-ban <enable|disable>
set service pppoe-server fastpado login-fail-ban attempts <3-20>
set service pppoe-server fastpado login-fail-ban ban-time <60-3600>
set service pppoe-server fastpado login-fail-ban known-max-ban-time <60-600>
set service pppoe-server fastpado login-fail-ban max-ban-time <60-7200>
set service pppoe-server fastpado login-fail-ban window <60-1800>
set service pppoe-server fastpado whitelist <mac>
set service pppoe-server subscriber-ping <enable|disable>
set service pppoe-server subscriber-ping source <auto|x.x.x.x>
set service pppoe-server interface <port> description <text>
set service pppoe-server interface <port> mru <1288-9216>
set service pppoe-server interface <port> mtu <1288-9216>
set service pppoe-server interface <port> vlan-id <number> description <text>
set service pppoe-server interface <port> vlan-id <number> mru <number>
set service pppoe-server interface <port> vlan-id <number> mtu <number>
set service pppoe-server interface <port> vlan-id <number> name <name>
set service pppoe-server interface <port> vlan-range <first-last> mru <number>
set service pppoe-server interface <port> vlan-range <first-last> mtu <number>
set service pppoe-server interface <port> vlan-range <first-last> name <name>
set service pppoe-server ip-pool address <network/mask> name <name> next <next-pool>
set service pppoe-server ip-pool gateway-address <ip>
set service pppoe-server ip-pool-name <name>
set service pppoe-server ipv6-dns-servers dns1 <ipv6>
set service pppoe-server ipv6-dns-servers dns2 <ipv6>
set service pppoe-server ipv6-pool address <ipv6-prefix> name <name> next <next-pool>
set service pppoe-server ipv6-pool delegated-address <ipv6-prefix> name <name> next <next-pool>
set service pppoe-server ipv6-pool gateway-address <ipv6>
set service pppoe-server ipv6-pool-delegate-name <name>
set service pppoe-server ipv6-pool-name <name>
set service pppoe-server local-user username <username> password <password> speed download
<value> upload <value> address <ip>
set service pppoe-server mru <1280-1500>
set service pppoe-server mtu <1280-1500>
```

```
set service pppoe-server pado-delay <number> sessions <number>
set service pppoe-server service-name <name> interface <port> vlan-id <number>
set service pppoe-server service-name anyname interface <port> vlan-id <number>
set service pppoe-server service-name blank disabled
set service pppoe-server service-name blank enabled
set service pppoe-server service-name blank interface <port> vlan-id <number>
delete service pppoe-server dns-servers dns1 <ip>
delete service pppoe-server dns-servers dns2 <ip>
delete service pppoe-server fastpado login-fail-ban [<setting>]
delete service pppoe-server fastpado whitelist <mac>
delete service pppoe-server subscriber-ping [source]
delete service pppoe-server interface <port> description
delete service pppoe-server interface <port> mru
delete service pppoe-server interface <port> mtu
delete service pppoe-server interface <port> vlan-id <number> description <text>
delete service pppoe-server interface <port> vlan-id <number> mru
delete service pppoe-server interface <port> vlan-id <number> mtu
delete service pppoe-server interface <port> vlan-id <number> name <name>
delete service pppoe-server interface <port> vlan-range
delete service pppoe-server interface <port> vlan-range <first-last> mru
delete service pppoe-server interface <port> vlan-range <first-last> mtu
delete service pppoe-server ip-pool address <network/mask> name <name> next <next-pool>
delete service pppoe-server ip-pool gateway-address <ip>
delete service pppoe-server ip-pool-name <pool-name>
delete service pppoe-server ipv6-dns-servers dns1 <ipv6>
delete service pppoe-server ipv6-dns-servers dns2 <ipv6>
delete service pppoe-server ipv6-pool address <ipv6-prefix> name <name> next <next-pool>
delete service pppoe-server ipv6-pool delegated-address <ipv6-prefix> name <name> next <next-
pool>
delete service pppoe-server ipv6-pool gateway-address <ipv6>
delete service pppoe-server ipv6-pool-delegate-name <pool-name>
delete service pppoe-server ipv6-pool-name <pool-name>
delete service pppoe-server local-user username <username>
delete service pppoe-server pado-delay <number> sessions <number>
delete service pppoe-server service-name <name> interface <port> vlan-id <number>
delete service pppoe-server service-name anyname interface <port> vlan-id <number>
delete service pppoe-server service-name blank interface <port> vlan-id <number>
show service pppoe-server
show service pppoe-server access-concentrator
show service pppoe-server authentication radius dae-server
show service pppoe-server authentication radius server
show service pppoe-server dns-servers
show service pppoe-server fastpado banned
show service pppoe-server fastpado login-fail-ban
show service pppoe-server fastpado whitelist
show service pppoe-server subscriber-ping
show service pppoe-server interface                                    (VLANs and the packet-size
table)
show service pppoe-server interface <port>
show service pppoe-server ip-pool
show service pppoe-server ip-pool-name
show service pppoe-server ipv6-dns-servers
show service pppoe-server ipv6-pool
show service pppoe-server ipv6-pool-delegate-name
show service pppoe-server ipv6-pool-name
show service pppoe-server local-user
```

```
show service pppoe-server pado-delay
show service pppoe-server service-name interface <port>
```

IPoE server (details)

```
set service ipoe-server idle-timeout <value>
set service ipoe-server interface <port> vlan-id <number> description <text>
set service ipoe-server interface <port> vlan-id <number> name <name>
set service ipoe-server interface <port> vlan-range <first-last> name <name>
set service ipoe-server network <x.x.x.x/x> gateway <ip> interface <port> vlan-id <number>
delete service ipoe-server idle-timeout <value>
delete service ipoe-server interface <port> vlan-id <number> description <text>
delete service ipoe-server interface <port> vlan-id <number> name <name>
delete service ipoe-server interface <port> vlan-range
delete service ipoe-server network <x.x.x.x/x> gateway <ip> interface <port> vlan-id <number>
show service ipoe-server interface <port>
```

NAT (SNAT / DNAT) (details)

```
set system dnat destination <ip> source ip <ip> protocol <tcp|udp> port <number> to-port <number>
set system dnat destination <ip> source pool <x.x.x.x/x> protocol <tcp|udp> port <number> to-port
<number>
set system dnat destination any source ip <ip> protocol <tcp|udp> port <number> to-port <number>
set system dnat destination any source pool <x.x.x.x/x> protocol <tcp|udp> port <number> to-port
<number>
set system snat destination-address-pool <x.x.x.x/x> translation ip <ip> out-interface <port>
set system snat destination-address-pool <x.x.x.x/x> translation pool <ip> out-interface <port>
set system snat nat-session-per-user limit <number>
set system snat source-address-pool <x.x.x.x/x> translation ip <ip> out-interface <port>
set system snat source-address-pool <x.x.x.x/x> translation pool <ip> out-interface <port>
delete system dnat destination <ip> source ip <ip> protocol <tcp|udp> port <number> to-port
<number>
delete system dnat destination <ip> source pool <x.x.x.x/x> protocol <tcp|udp> port <number> to-
port <number>
delete system dnat destination any source ip <ip> protocol <tcp|udp> port <number> to-port
<number>
delete system dnat destination any source pool <x.x.x.x/x> protocol <tcp|udp> port <number> to-
port <number>
delete system snat destination-address-pool <x.x.x.x/x> translation ip <ip> out-interface <port>
delete system snat destination-address-pool <x.x.x.x/x> translation pool <ip> out-interface
<port>
delete system snat nat-session-per-user
delete system snat source-address-pool <x.x.x.x/x> translation ip <ip> out-interface <port>
delete system snat source-address-pool <x.x.x.x/x> translation pool <ip> out-interface <port>
show system dnat
show system snat
restart system dnat
restart system snat
```

Log server (details)

```
set system log-server address <ip> port <number> ipfix all
set system log-server address <ip> port <number> netflowv9 all
set system log-server address <ip> port <number> syslog
set system log-server source-ip-pool <x.x.x.x/x>
set system log-server source-ipv6-pool <ipv6-prefix>
delete system log-server address <ip> port <number> ipfix all
delete system log-server address <ip> port <number> netflowv9 all
delete system log-server address <ip> port <number> syslog
delete system log-server source-ip-pool <x.x.x.x/x>
delete system log-server source-ipv6-pool <ipv6-prefix>
```

DNS cache and DNS redirect (details)

```
set system dns address <ip>
set system dns allow-pool <x.x.x.x/x>
set system dns disable
set system dns enable
set system dns forwarder <ip>
set system dns port <number>
set system dns threads <value>
set system dns-redirect state disabled
set system dns-redirect state enabled
set system dns-redirect target <ip>
delete system dns address <ip>
delete system dns allow-pool <x.x.x.x/x>
delete system dns forwarder <ip>
delete system dns port <number>
delete system dns threads <value>
delete system dns-redirect
delete system dns-redirect state
delete system dns-redirect target
show system dns-redirect state
show system dns-redirect target
restart system dns
```

SNMP (details)

```
set system snmp allow <x.x.x.x/x>
set system snmp community <text>
set system snmp contact <text>
set system snmp disable
set system snmp enable
set system snmp listen-address <ip>
set system snmp location <text>
set system snmp port <number>
set system snmp trap-target <ip> community <value>
set system snmp v3 auth md5 <password>
set system snmp v3 auth sha <password>
set system snmp v3 priv aes <password>
set system snmp v3 priv des <password>
set system snmp v3 user <name>
set system snmp version 1
set system snmp version 2c
set system snmp version 3
delete system snmp allow <x.x.x.x/x>
delete system snmp community <text>
delete system snmp contact <text>
delete system snmp listen-address
delete system snmp location <text>
delete system snmp port
delete system snmp trap-target <ip>
delete system snmp v3
delete system snmp version
show system snmp
restart system snmp
```

Web panel, API and CLI users (details)

```
set system api address <ip> port <number>
set system api login user <user> authentication password <password>
set system login user <user> authentication password <password>
set system login user <user> level admin
set system login user <user> level superadmin
set system login user <user> level user
set system web address <ip> port <number>
set system web login user <user> authentication password <password>
set system web login user <user> interface <port> vlan-id <number>
set system web login user <user> level viewer | operator | superuser
delete system api address <ip> port <number>
delete system api login user <user>
delete system login user <user>
delete system web address <ip> port <number>
delete system web login user <user> interface <port> vlan-id <number>
delete system web login user <user> level
show system login
show system web login
restart system web
```

Speedtest, license, tuning and other system settings (details)

```
set system auto-tune apply
set system auto-tune irq-spread <auto|on|off>
set system auto-tune napi-defer <on|off>
set system auto-tune idle <c1|poll>
set system auto-tune off
set system auto-tune status
set system host-name <name>
set system license key <key-string>
set system routing <bgp|ospf|isis>
set system speedtest domain <name>
set system speedtest server disable
set system speedtest server enable
set system time-zone <Region/City>
set system upload-mode htb
set system upload-mode police
delete system host-name
delete system license key
delete system license temp-key
delete system speedtest domain
show system cpu
show system health
show system host-name
show system license
show system status
show system upload-mode
```

High availability (details)

```
set system high-availability link <port>
set system high-availability role <primary|secondary>
set system high-availability key <text>
delete system high-availability
show high-availability
high-availability switchover
high-availability switchover force
high-availability takeover force
```

Routes, IX policy and QoS groups (details)

```
set ip route <x.x.x.x/x> gateway <ip> out-interface <port> vlan-id <number>
set policy qos group <name> network <x.x.x.x/x>
set policy qos group <name> port <1-65535>
set policy qos group <name> priority <1-63>
set policy qos group <name> speed all download <rate> upload <rate> [time any | time <HH:MM> to
<HH:MM>]
set policy qos group <name> speed multiply download <factor> upload <factor> [time any | time
<HH:MM> to <HH:MM>]
set policy qos group <name> speed plus download <rate> upload <rate> [time any | time <HH:MM> to
<HH:MM>]
set policy route ix address <ip>
set policy route ix address6 <ipv6>
set policy route ix network <x.x.x.x/x>
set policy route ix network6 <ipv6-prefix>
set policy route ix qos cbq all <rate>
set policy route ix qos cbq multiply <factor>
set policy route ix qos cbq plus <rate>
delete ip route <x.x.x.x/x> gateway <ip>
delete policy qos
delete policy qos group <name>
delete policy qos group <name> network <x.x.x.x/x>
delete policy qos group <name> port <1-65535>
delete policy qos group <name> priority
delete policy route ix address <ip>
delete policy route ix address6 <ipv6>
delete policy route ix network <x.x.x.x/x>
delete policy route ix network6 <ipv6-prefix>
delete policy route ix qos cbq
show ip route <x.x.x.x/x>
show policy qos
show policy qos group <name>
show policy route ix
show policy set
```

Other show commands (details)

```
show arp interface <port>
show arp ip <ip>
show configuration
show configuration set
show connections
show connections <username>
show connections log
show date
show interfaces
show ip neighbour interface <port>
show system
show traffic interface <port> bandwidth
```

Operations (take effect at once, no commit) (details)

```
delete arp ip <ip>
commit
commit confirm
commit discard
disconnect ifname <pppN>
disconnect user <username>
reboot system
restart service pppoe-server
restart service pppoe-server interface <port>
restart service pppoe-server interface <port> vlan-id <number>
shutdown system
```

Backup / Restore (details)

```
backup create [note]   |   backup list   |   backup status
backup show <name|number>   |   backup export <name|number>
backup import [name]   |   backup delete <name|number>
backup restore <name|number>   |   backup restore <name|number> logins
backup keep   |   backup return
```

Tools

```
ping <host>   |   ping <host> count <number>
ping6 <host>   |   ping6 <host> count <number>
traceroute <host>   |   traceroute <host> max-hops <number>
traceroute6 <host>   |   traceroute6 <host> max-hops <number>
ssh <host>   |   telnet <host>
clear   |   help   |   exit
```

Troubleshooting

Symptom	Check
Users log in but have no internet	`show system health`: "SNAT rule" and "un-NATed private source". The SNAT out-interface must be the WAN port of this box and the private pool must match the ip-pool. Users in a pool without a NAT rule (expired pool) are blocked by design.
Some users cannot connect	`show subscriber <username|mac|ip>` gives the reason and what to do; `show subscriber failures` lists who failed and why in the last hour. Behind it, `show connections log` for "authentication failed" (wrong password or account), RADIUS reply messages, or "Service-Name mismatch" (the line names the service name, MAC, port and VLAN: add that service name, or the VLAN, to the configuration); check RADIUS reachability in the health check.
A device gets no answer while FastPADO is on	`show service pppoe-server fastpado banned` and `... login-fail-ban`: a flooding or repeatedly rejected device is held back for a few minutes. Fix its account or password, or `set service pppoe-server fastpado whitelist <mac>`. Make sure its VLAN is in the configuration: a VLAN that is not configured is never answered.
Some sites or large downloads hang for the users of one VLAN (small pages and ping work)	The access link under that VLAN carries less than 1500 bytes (radio link, tunnel, extra tags). Set its real size: `set service pppoe-server interface <port> vlan-id <number> mtu <number>` (and `mru`), `commit`, then `show service pppoe-server interface`. The MTU column of the web Connections page shows what each user has; a user gets the new size at the next login. See PPPoE packet sizes.
The BNG cannot ping a user who is online	`show service pppoe-server subscriber-ping` must be enabled with a public source. If it is, the user's router itself blocks ping. To check by hand, `tcpdump -ni ppp<N> icmp` while pinging: only "echo request" lines and no reply means the router is not answering.
Upload is not limited	`show system health`: "upload mark program" must be pinned and "session accounting" must show policers = sessions. Sessions started before a fix pick the new limit at re-login.
Packet loss at peak on the PPPoE port	PHYSICAL PORTS card / health check "receive ring": one receive queue is full. Turn flow control off on that port, split subscribers over more ports, or use X710/XL710 with DDP.
Monitoring says SNMP does not answer	`show system snmp`: is the NMS address in the allow list? Does "requests received" grow while it polls?
One user slows everyone	Stat tab, DDoS WATCH and Top 20 NAT sessions: a device at the connection limit or flooding; the limit protects the table, disconnect or inform the subscriber.
Web password lost (no web user can log in)	CLI, as admin: `set system web login user super authentication password <new password>` and `commit`. Then log in to the web as `super` and repair the other users in Config > Users.
No Config tab in the web panel	The Config tab is shown only to web users with the superuser level. `show system web login` lists every user and level; `set system web login user <name> level superuser` and commit, then log out and in again.
Web login refused although the password is right	Five failed logins from one address lock that address for five minutes; wait and try again. A password set before V7.3.2 still works; `show system web login` marks such accounts "old-hash" until the password is set again.
Disconnect or Unban answers "Not allowed"	Those actions need the operator level or higher: `set system web login user <name> level operator`.
Config page: "configuration service is not running"	The web hands its changes to the `dandy-confd` service. `systemctl restart dandy-confd`, then reload the page. Nothing is lost: uncommitted lines are simply dropped.

Symptom	Check
Config page: changes are held by another user	Only one superuser holds pending changes at a time; the page names that user. Ask them to commit or discard, or press Discard to drop their lines. Pending changes left alone for 15 minutes expire by themselves.
A NAT or DNAT line is refused	The page shows the reason (overlapping pool, private public pool, inside host outside every subscriber pool, port already used by the box). Correct the value; the same line typed in the CLI is accepted only when it is valid there too.
Web panel not reachable	`restart system web`, or set `set system web address ... port ...` again and commit. The panel restarts for about a minute after a WAN address change.
Clock wrong, "time sync" WARN	Make sure the box can reach NTP servers from the WAN; set the time zone.

DandyBNG V7.3.3 user guide. Firewall/DPI filtering is not part of this release.