

zapier.georgelambert.org · Markdown indexes

VERAE TIME x ZAPIER

Proxmox 3-node NATS JetStream cluster

<site/packages/verae-nats-cluster/README.md>

verae-nats-cluster

Three **distinct Proxmox LXC** guests running **NATS JetStream** as a cluster, on the private bridge only. This is the test stand before the same layout moves to separate hardware.

Forgejo: <https://git.georgelambert.org/marchon/verae-nats-cluster>

SSH: <ssh://git@git.georgelambert.org:2223/marchon/verae-nats-cluster.git>

Guest	VMID	<code>vmbr1</code> IP	Client	Cluster	Monitor
nats-a	511	10.10.10.21	:4222	:6222	127.0.0.1:8222
nats-b	512	10.10.10.22	:4222	:6222	127.0.0.1:8222
nats-c	513	10.10.10.23	:4222	:6222	127.0.0.1:8222

Cluster name: `verae` . Client URL:

```
nats://10.10.10.21:4222,nats://10.10.10.22:4222,nats://10.10.10.23:4222
```

Cluster routes (:6222) and hub leaf (:7422) stay on `vmbr1` . A token-gated public client edge on NS1 forwards to that hub:

```
nats://<token>@70.88.205.138:4222
wss://nats.zapier.georgelambert.org          # WebSocket (CONNECT with the same token)
```

Token file on NS1: `/etc/nats/public-edge.token` . Enable/recreate: `bash scripts/expose-public.sh` . Do **not** publish `:6222` or `:8222` . Internal clients keep using the three `10.10.10.21-23` URLs (no token).

```
# on NS1 (the Proxmox host)
git clone ssh://git@git.georgelambert.org:2223/marchon/verae-nats-cluster.git
cd verae-nats-cluster
bash scripts/create-cluster.sh
bash scripts/status.sh
bash scripts/test.sh
```

Test-env cut-over (keep/fleet/middleware on NS1):

```
# 1) create product streams (replicas=3)
bash scripts/ensure-streams.sh
# 2) fleet overlay nats.url + services/*.json already list the three URLs
# 3) restart fleet and keep with NATS_URL from client.env
```

`JETSTREAM_REPLICAS=3` on middleware. Host `127.0.0.1:4222` can stay up unused. Zapier cloud never talks to NATS.

nkeys/mTLS: `verae-nats-accounts` still has the INTERNAL/LEAF sketch. Do **not** enable accounts on this cluster until every client passes credentials in `NATS_URL`. Private `vmbr1` is the current isolation.

Message speed (throughput + delay)

From a **client that is not a nats-* server** (LXC 510 `verae-px-worker`):

```
# on NS1 (hypervisor). Bench client is LXC 510.
bash scripts/bench.sh

# Full study that must run ON NS1 (orchestrator + charts + HTML/PDF):
bash scripts/study-on-ns1.sh

# Maximize cores/RAM + tmpfs JetStream, same ladder, restore ZFS after:
bash scripts/maximize-ns1-study.sh

# Exhaustive factorial (r=1/r=3, file/memory, reconnect, UDP, MQTT) + combined report:
bash scripts/exhaustive-ns1-study.sh
```

That writes `results/<utc>/` and copies a summary to [BENCH.md](#). Full write-up with charts (HTML + PDF) is in the progress repo: [nats-cluster-bench](#). Core NATS is fire-and-forget; JetStream rows use **file** storage and **replicas=3**. Ping delay is sequential RTT; flood delay is queueing under burst.

Latest run (`20260912T045131Z`):

Kind	Load	Result
Core 1p1s 50k×128 B	pub/sub	~1.20M msgs/s aggregate
Core 8p8s 200k×128 B	pub/sub	~2.01M msgs/s aggregate
Core 4p4s 50k×1 KiB	pub/sub	~630k msgs/s (~616 MB/s aggregate)
JetStream r=3 file	1–4 pubs, 128 B	~16k durable pubs/s
JetStream r=3 file	4 pubs, 1 KiB	~13.5k durable pubs/s
Ping 1k×128 B	sequential RTT	avg 0.307 ms, p99 0.734 ms
Flood 8p 20k×128 B	burst queueing	avg 454 ms

Hardware move: same three configs, three boxes, private NIC only — change IPs in `cluster.env` and `conf/nats.conf.tpl`.