On the Feasibility of Matrix as a Messaging Substrate for a Tactical Radio MANET

Richard Neuschulz
`rneuschulz@gmail.com`

June 2026

Abstract

This work asks whether the Matrix protocol can serve as the messaging substrate for a tactical radio mobile ad-hoc network (MANET): a mesh of vehicle-borne homeservers joined by intermittent, very-low-bandwidth federation links. Matrix federation performs no multi-hop routing, so cross-node delivery happens only opportunistically through room-directed acyclic graph (DAG) backfill, which makes the room graph an implicit store-carry-forward (SCF) layer. A single-host Kubernetes (k8s), single node testbed shapes only the server-to-server links and drives a measurement campaign across a bandwidth sweep and a fleet-size sweep ($N = 3$ to 70). Two findings emerge. First, room formation at scale is a deployment requirement, not a protocol barrier: naive concurrent joins fork the DAG through resident-server rate limiting, but sequential pre-establishment behind a convergence gate brings all $N = 70$ servers to a consistent membership that then delivers near 100 %. Second, the binding barrier is bandwidth and fan-out. Each message fans out into $(N - 1)$ point-to-point transactions, each carrying a measured, fleet-size-independent 762 B signed envelope through one shared radio, so the bandwidth required for full delivery grows super-linearly: from 10 kbit/s at $N = 3$ to 459 kbit/s at $N = 70$. The latter sits above the 162 kbit/s to 250 kbit/s headline rate of fielded wideband combat-net radios, so a 70-vehicle operational room is infeasible at tactical link rates.

1 Introduction

Tactical radio networks for vehicle-borne units are the archetypal challenged network: links are slow, asymmetric, and frequently partitioned as vehicles move in and out of range. Messaging in such an environment is a delay-tolerant, SCF problem rather than an end-to-end-path problem. Matrix is an attractive candidate substrate: it is an open, federated, eventually-consistent messaging protocol with a mature reference homeserver (Synapse), a per-room event model that is already a replicated DAG, and an extensive set of federation-tuning controls.

Matrix federation, however, does not route. A homeserver synchronises a room's event graph only with the servers it can currently reach; it does not forward another server's events on its behalf. When a server reconnects after a partition and notices it is behind, it pulls the missing history via backfill. The room DAG is therefore an implicit SCF mechanism, and whether that mechanism is good enough under tactical link rates is an empirical question. This work contributes: (1) a faithful single-host testbed that shapes only the federation links; (2) a corrected scaling analysis distinguishing a deployment requirement (room pre-establishment) from the real binding barrier (bandwidth and fan-out); and (3) a measured required-bandwidth law spanning the full $N = 3$ to 70 design range, interpreted against fielded tactical-radio capabilities.

2 Tactical radio context

The test design is anchored to the capabilities of fielded and emerging tactical communication systems. Table 1 summarises four representative systems from public vendor datasheets and defense references. Two observations frame the experiment. First, none of these systems publishes a headline rate as low as 1 kbit/s to 9 kbit/s; that regime represents the realistic *per-user* data rate in a shared, contested narrowband net, where a 64 kbit/s combat-net radio (CNR) channel is divided across a platoon-sized group and voice has priority. Second, tactical IP data on wideband waveforms sits at 64 kbit/s to 250 kbit/s nominally (Thales F@stnet at 64 kbit/s, F@stnet HD GeoMux at 162 kbit/s, CNR-HD at 250 kbit/s) [12, 13]. These two tiers, a constrained per-user tier and a wideband IP tier, bound the bandwidth axis swept in section 5; the constrained tier is the stress case for Matrix synchronisation.

Table 1: Representative tactical communication systems (public datasheet values). N/A denotes a category that does not apply to a software or server product.

System	Max data rate	Frequency band	Voice	IP	Network flattening
blackned Tactical Core (middleware) [7, 8]	Inherits bearer rate; wideband and narrowband	N/A (hardware-agnostic)	Yes	Yes	Self-organising, self-healing mesh
R&S SOVERON VR / VR5000 [9, 10, 11]	High-rate IP waveforms; up to 4000 addressable radios	VHF/UHF 30 MHz to 512 MHz	Yes	Yes	Integrated self-forming, self-healing MANET
Thales PR4G / F@stnet (HD) [12, 13, 14]	64 kbit/s (FH); HD up to 250 kbit/s	VHF 30 MHz to 88 MHz	Yes	Yes	Frequency-hopping CNR; SDR
KommServer (KNDS + Tactical Core) [15, 8]	Aggregates attached bearer rates	N/A (server)	Yes	Yes	Gateway collapsing legacy nets into one network

3 Background and theory

3.1 Matrix federation and the room DAG

In Matrix, homeservers exchange room history over the Server-Server (federation) API [1]. Each room's history is a DAG of events whose partial order encodes causal ordering; the graph is synchronised between participating servers with eventual consistency, and each server validates incoming events against signature, hash, and authorisation checks. A server exchanges events only with servers it can currently contact and recovers missing history through backfill rather than through any in-network relaying. Synapse [2] is the reference homeserver used here, chosen for its federation retry controls and Admin API. The delay-/disruption-tolerant networking (DTN) literature formalises this regime: networks where a contemporaneous end-to-end path may never exist, so messages are stored and forwarded on later contacts [3, 4, 5]. Classical MANET routing protocols [6] instead assume a maintainable path and degrade under chronic partition, which is why the SCF model is the right lens here.

3.2 The fan-out and overhead model

Matrix delivers a room message by emitting one point-to-point federation transaction per remote participant. In an N-server room a sender therefore transmits $(N-1)$ copies through its single radio. Let r be the per-node send rate, S the application payload, and O the per-message federation envelope. The fluid byte-budget required at one radio for full delivery is

$$B_{\text{req}} = r\,(N-1)\,(S+O). \tag{1}$$

Equation (1) fixes the *shape* of the scaling (linear in $N-1$); the measured requirement (section 5) sits well above it because federation is request/response- and round-trip-bound rather than throughput-bound.

A decisive protocol property concerns O. The signed federation protocol data unit (PDU) carrying a single room message contains an origin-only signature (exactly one), a fixed set of `auth_events` (three: create, power-levels, membership), a content hash, and event metadata; only the `prev_events` reference list grows, and only under concurrent sends. Consequently O is independent of fleet size: the per-message envelope does not grow with N; what grows is the *number* of transactions, $(N-1)$. Direct serialisation of the signed PDU confirms this, giving

$$O = 762\,\text{B} \tag{2}$$

measured identically from $N = 3$ to $N = 70$ (fig. 7). This value is the Matrix-PDU layer only and excludes lower-layer transport layer security (TLS)/TCP/IP framing, so it is a lower bound on bytes-on-wire. With a representative 220 B tactical payload, the useful-payload ratio $S/(S+O)$ is only ≈ 0.22.

4 Method

4.1 Mapping and testbed

One node equals one vehicle equals one Synapse homeserver. Users are modelled as clients permanently attached to their own vehicle's homeserver over a reliable intra-vehicle link that is never impaired; the constrained, dynamic radio links are the server-to-server (federation) edges. The mesh runs on a single host on purpose: co-locating all servers makes every federation edge fully synthetic and controllable, avoiding real-network variance on the links under study. The stack is Kubernetes (k8s), single node with Calico providing enforced policy as the container network interface (CNI), Chaos Mesh for link emulation and partitioning, Synapse homeservers backed by one PostgreSQL cluster (one database per node), and a single async orchestrator that holds one client session per node, generates randomised traffic, and timestamps sends and receipts against one clock (no inter-node skew). Pinned versions are listed in table 4.

4.2 Link shaping and Synapse tuning

Two layers gate server-to-server traffic and are kept distinct: Calico policy defines the static allowed-edge graph (which servers may peer at all), while Chaos Mesh applies the dynamic per-edge bandwidth shaping and scheduled partition/heal windows. Shaping is applied only to the federation port (8448); Postgres and client links are local and unthrottled. Default Synapse is built for the public internet and would silently mis-measure a tactical link, so the homeserver template was tuned against the constraints of section 2, verified against the Synapse configuration schema:

- **Federation retry backoff.** Defaults (10 min minimum, multiplier 5, ∼1 week maximum) are far too slow for up-windows of seconds to minutes; they were lowered (minimum ∼30 s, multiplier 2, low maximum), and the topology operator calls the Admin API to reset a destination's backoff on every link heal, removing the confound at the moments that matter.
- **Chatter suppression.** Presence is disabled (its federation ephemeral data units (EDUs) are the worst offender on a 7 kbit/s link), the orchestrator never emits typing or read-receipt EDUs, and federation profile and device-name lookups are disabled.

- **Closed-mesh corrections.** On an air-gapped mesh, trusted key servers must be emptied so each node fetches signing keys directly from the origin over the shaped link rather than hanging on a public notary; the anti-SSRF IP filter must whitelist the in-cluster pod range (classless inter-domain routing (CIDR) `10.42.0.0/16`), otherwise `make_join` fails despite working DNS and TLS; the Chaos Mesh bandwidth `rate` field is bytes per second, so target rates are converted as kbit/s × 125; and one stable ed25519 signing key per node is pre-minted so identity survives restarts. Each correction would otherwise masquerade as a delivery failure.
- **Footprint.** Cache factors and event-cache size are trimmed in monolith mode, holding idle per-Synapse resident memory near 100 MiB.

4.3 Convergence-gate setup

Forming a large room is itself a measured step. Joining $N = 70$ servers concurrently overruns the resident server's rate limiter, fails a fraction of joins, and forks the room DAG, so membership never converges. The orchestrator instead performs robust sequential joins (retry with backoff, $\sim$0.4 s spacing) behind a *convergence gate* that blocks the measured phase until every server observes all N members. This models forming the operational room while the fleet is connected, before dispersal, which is realistic for a tactical deployment.

4.4 Scenarios and metrics

Scenario s0 (full mesh) shapes all federation edges at a single bandwidth and isolates the pure low-rate effect. Scenario s1 (two-convoy partition/heal) splits nodes into two convoys whose intra-convoy links stay up while the inter-convoy link flaps on a staggered schedule, so cross-convoy delivery is partition-gated and can progress only during heal windows: the canonical SCF test. Per intended recipient the orchestrator records delivery ratio, end-to-end latency (receipt minus send, on the single clock) including the p95 and maximum tail, and a non-convergence flag (delivery below 50 % at run end). Runs are serial on the one box; every measured phase is gated on Chaos Mesh confirming the throttle is injected, and room membership is reset between runs. Per-run CPU-saturation is logged and saturated runs are discarded. The campaign comprises 237 runs, of which 228 are valid after discarding nine CPU-saturated runs.

5 Results

5.1 Delivery, latency, and goodput

Within a converged room, link bandwidth is the binding constraint and delivery is a clean monotone function of it. Figure 1 shows delivery ratio versus federation bandwidth: at $N = 10$ delivery is below 20 % across the entire 1 kbit/s to 9 kbit/s tactical band, crosses 50 % near 16 kbit/s, and reaches 100 % only from 30 kbit/s upward. Among delivered messages, latency stays in the tens of seconds throughout the throttled band and collapses to a few seconds only once the link clears (fig. 2); the latency distribution carries a long right tail at low bandwidth as backlog drains slowly (fig. 3). Goodput follows eq. (2): the useful-payload ratio is ≈ 0.22 at the 220 B representative payload and reaches 0.5 only once the payload itself equals the 762 B envelope (fig. 4).

5.2 Required-bandwidth scaling

Sweeping fleet size at a fixed per-node rate ($\approx$ 0.05/s per node) recovers B_{100}, the bandwidth at which a converged room reaches full delivery, as a function of N (table 2, fig. 5). B_{100} grows super-linearly, from 10 kbit/s at $N = 3$ to 459 kbit/s at $N = 70$. The fluid prediction of eq. (1) captures the shape but under-predicts the magnitude by a factor $k = B_{100}/B_{\text{fluid}}$ that ranges from

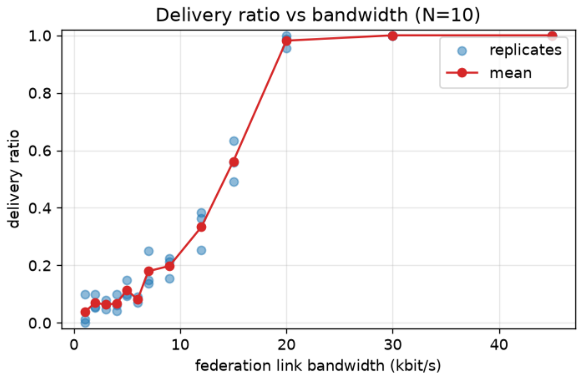

Figure 1: Delivery ratio versus federation bandwidth (s0, $N = 10$). Delivery is monotone in link rate and stays below 20 % across the tactical band.

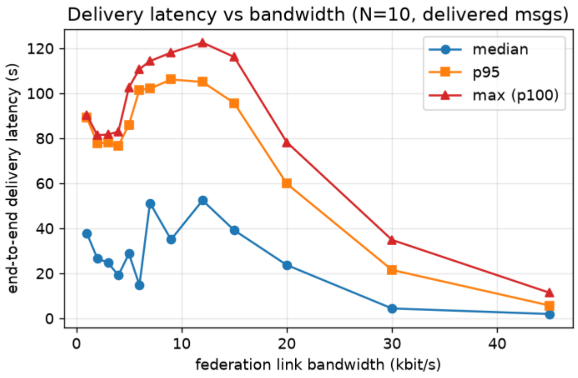

Figure 2: Median, p95, and maximum end-to-end latency versus bandwidth (s0, $N = 10$, delivered messages).

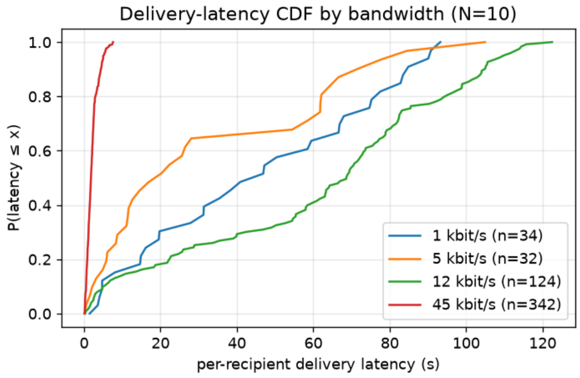

Figure 3: End-to-end latency CDF per bandwidth (s0, delivered messages). The low-bandwidth tail stretches out as backlog drains slowly.

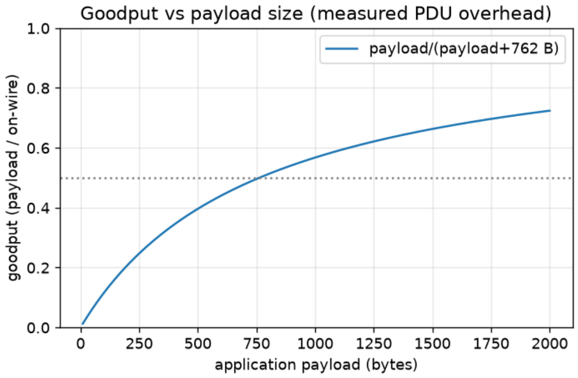

Figure 4: Goodput $S/(S + O)$ versus payload size using the measured 762 B envelope; per-message payloads overlaid.

≈ 8 to ≈ 17: this is the protocol's round-trip and request/response tax, which compounds with the number of peers. A second cut holds $N = 10$ fixed and varies the per-node rate; B_{100} rises with offered load (fig. 6), confirming the law is driven by aggregate offered transactions. The full delivery surface over the $N\times$ bandwidth grid is shown in fig. 8: the full-delivery frontier marches rightward as N grows.

5.3 Room formation and partition behaviour

Room formation at scale is a deployment requirement, not a protocol barrier. With the convergence-gate setup, membership converges to the full fleet at every N tested (to 100 % for $N \leq 40$ and at $N = 70$, and to 98.4 % at $N = 50$); a converged $N = 70$ room then delivers near 100 % once bandwidth is sufficient. An earlier apparent "convergence ceiling" near 37 % at large N was traced to naive concurrent joins forking the DAG under resident-server rate limiting, and is removed entirely by sequential pre-establishment. In the two-convoy scenario (fig. 9), intra-convoy delivery flows while cross-convoy delivery is gated by the partition and can progress only during heal windows, demonstrating the room DAG acting as an opportunistic SCF layer; at low bandwidth the heal windows do not drain the cross-convoy backlog.

Table 2: Required bandwidth for full delivery B_{100} versus fleet size at fixed per-node rate, with the fluid prediction of eq. (1) and the inefficiency factor k. At $N = 50$ delivery peaked at 98.4 % within the tested range, so B_{100} is not pinned.

N	B_{100} (kbit/s)	B_{fluid} (kbit/s)	k
3	10	0.7	12.9
5	20	1.6	12.5
10	28	3.5	7.9
20	67	7.5	9.0
30	189	11.4	16.6
40	210	15.3	13.7
70	459	27.1	17.0

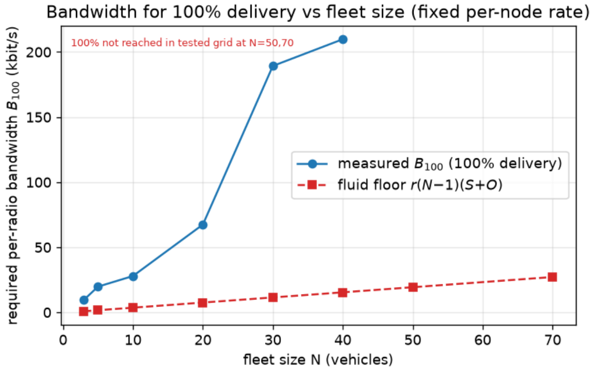

Figure 5: B_{100} versus fleet size N at fixed per-node rate, against the fluid floor. The gap is the fan-out and round-trip tax.

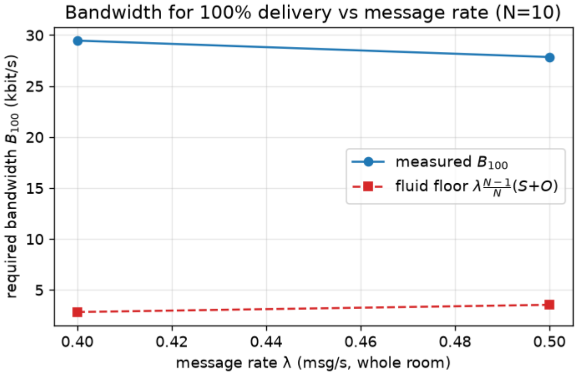

Figure 6: B_{100} versus per-node message rate ($N = 10$ fixed); required bandwidth rises with offered load.

6 Discussion

The corrected narrative has two parts. Room formation at scale is a procedure, not a limit: a single room of 70 servers does converge, provided it is pre-established while the fleet is connected rather than assembled by concurrent joins under rate limiting. Once that procedure is in place, the binding barrier is bandwidth and fan-out. Because a sender emits $(N-1)$ copies of a fleet-size-independent 762 B envelope (eqs. (1) and (2)) through one shared radio, the per-radio bandwidth required for full delivery grows super-linearly, reaching 459 kbit/s at $N = 70$ (table 2). Read against section 2, this is decisive: 459 kbit/s exceeds the 162 kbit/s to 250 kbit/s headline rate of fielded wideband combat-net radios such as F@stnet HD, and is two orders of magnitude above the 1 kbit/s to 9 kbit/s realistic per-user rate on a shared narrowband net. A 70-vehicle operational room over a single Matrix room is therefore infeasible at tactical link rates, and the measured k of 8 to 17 shows that a naive byte count understates the requirement because federation is round-trip-bound. The practical implication is to avoid one large room: fleet partitioning into many small rooms, a thinner-footprint homeserver, or an explicit SCF overlay each warrant evaluation.

7 Limitations

- **Replication.** Two replicates per s0 cell and a small s1 sweep; convergence latency is noisy, so confidence intervals are wide and trends are indicative rather than statistically settled.
- **Overhead measured at the PDU layer.** The 762 B envelope excludes TLS/TCP/IP framing and backfill/key-fetch round-trips, so the goodput figures are an upper bound on

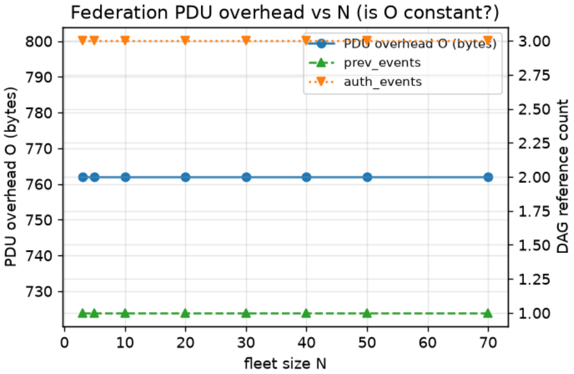

Figure 7: Measured signed-PDU envelope versus fleet size: flat at 762 B from $N = 3$ to $N = 70$ (eq. (2)).

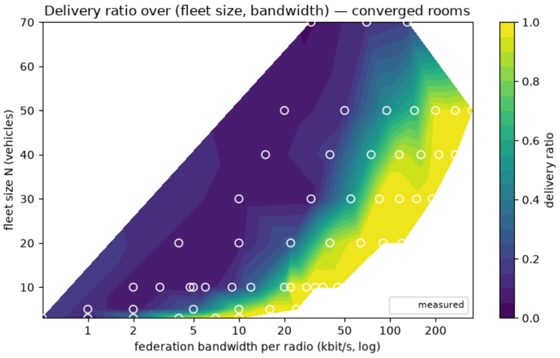

Figure 8: Delivery-ratio surface over the $N\times$ bandwidth grid (s0). The full-delivery frontier moves rightward as N grows.

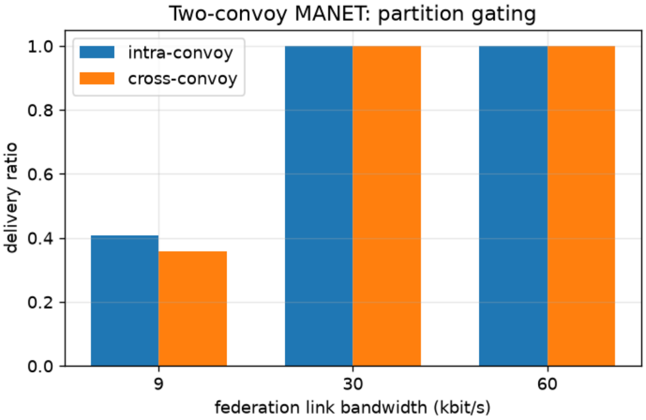

Figure 9: Two-convoy partition/heal (s1): intra-convoy delivery (links always up) versus partition-gated cross-convoy delivery, which advances only during the staggered heal windows.

real goodput; a packet capture would tighten them downward.

- **Synthetic links on one host.** Bandwidth, delay, and partitions are emulated via Chaos Mesh on a single host. This is deliberate (control over realism), but it omits real radio effects (fading, MAC contention, asymmetric/jittery rates, bit-level corruption) and shares one kernel network stack.
- **Shared resources.** All servers, Postgres, and the orchestrator contend for one 16-core host; despite the CPU-saturation guard, residual contention could perturb tail latency, and a synchronised backfill burst at large N stresses the box more than at the validated $N = 10$ scale.

A Per-run summary (fixed per-node rate series)

Table 3 aggregates the fixed-rate scaling series (mean over replicates per cell).

B Reproducibility and versions

Glossary of acronyms

CIDR classless inter-domain routing

CNI container network interface

CNR combat-net radio

DAG directed acyclic graph

DTN delay-/disruption-tolerant networking

EDU ephemeral data unit

MANET mobile ad-hoc network

PDU protocol data unit

SCF store-carry-forward

SDR software-defined radio

TLS transport layer security

References

[1] The Matrix.org Foundation. *Matrix Specification: Server-Server (Federation) API.* spec.matrix.org.

[2] Element / The Matrix.org Foundation. *Synapse: a Matrix homeserver.* GitHub: element-hq/synapse.

[3] K. Fall. A Delay-Tolerant Network Architecture for Challenged Internets. *Proc. ACM SIGCOMM*, 2003, pp. 27–34.

[4] V. Cerf et al. *Delay-Tolerant Networking Architecture.* IETF RFC 4838, April 2007.

[5] A. Vahdat and D. Becker. *Epidemic Routing for Partially-Connected Ad Hoc Networks.* Tech. Report CS-2000-06, Duke University, 2000.

[6] E. M. Royer and C.-K. Toh. A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks. *IEEE Personal Communications*, 6(2):46–55, 1999.

[7] blackned GmbH. *Products: Tactical Core (RIDUX, XONITOR, MeshFlow, PRA).* blackned.de/products, accessed 2026-06-27.

[8] ESUT. *Tactical Core als Schaltstelle fuer Datenkanaele.* esut.de, 2023, accessed 2026-06-27.

[9] Rohde & Schwarz. *SOVERON VR product page (*30 MHz *to* 512 MHz*; self-forming MANET).* rohde-schwarz.com, accessed 2026-06-27.

[10] Rohde & Schwarz. *R&S SDTR VR5000 Vehicular Tactical Radio, Product Brochure v02.01.* rohde-schwarz.com, accessed 2026-06-27.

[11] Rohde & Schwarz. *SOVERON WAVE High-Data-Rate Waveform Suite, Product Brochure v02.01.* rohde-schwarz.com, accessed 2026-06-27.

[12] Thales. *F@stnet HD VHF tactical radio (GeoMux HD up to* 162 kbit/s*; CNR-HD up to* 250 kbit/s*).* thalesgroup.com, accessed 2026-06-27.

[13] Army Technology. *Thales PR4G F@stnet waveform (up to* 64 kbit/s *ECCM; integrated IP router).* army-technology.com, accessed 2026-06-27.

[14] Defense Update. *Thales PR4G (VHF* 30 MHz *to* 88 MHz*; secure voice; simultaneous voice and data).* defense-update.com, accessed 2026-06-27.

[15] KNDS. *KNDS Deutschland Mission Electronics: KommServer central communications node.* knds.com, accessed 2026-06-27.

Table 3: Aggregated per-cell results for the fixed per-node rate series (≈ 0.05/s per node). Delivery and median latency are means over the replicates in each cell.

N	Bandwidth (kbit/s)	Reps	Delivery (%)	Median lat. (s)	Non-conv.
3	0.5	4	19.5	17.6	yes
3	1	4	9.4	39.8	yes
3	2	4	35.9	34.7	yes
3	4	4	82.0	42.0	no
3	7	4	99.2	9.0	no
3	10	4	100.0	1.1	no
5	1	4	12.4	25.8	yes
5	2	4	6.7	31.4	yes
5	5	4	28.6	33.4	yes
5	10	4	80.0	34.7	no
5	16	4	99.8	3.3	no
5	24	4	100.0	1.4	no
10	2	2	6.3	26.9	yes
10	3	4	7.9	26.9	yes
10	5	4	10.0	31.4	yes
10	9	4	14.7	57.0	yes
10	12	4	30.2	62.5	yes
10	20	2	72.6	38.0	no
10	28	2	99.9	13.1	no
10	32	2	100.0	5.9	no
10	45	2	100.0	2.8	no
20	4	2	6.1	60.0	yes
20	10	2	7.9	28.9	yes
20	22	2	27.9	82.8	yes
20	40	2	87.1	41.8	no
20	64	2	99.9	12.2	no
20	90	2	100.0	3.0	no
20	120	2	100.0	1.8	no
30	10	2	6.1	50.1	yes
30	30	2	7.5	62.2	yes
30	55	2	50.1	43.5	yes
30	85	2	99.4	18.9	no
30	150	2	96.4	5.6	no
30	190	2	100.0	2.7	no
40	15	2	6.2	61.5	yes
40	40	2	17.0	48.4	yes
40	75	2	56.8	34.6	no
40	115	2	91.8	19.0	no
40	210	2	100.0	6.0	no
40	270	2	99.9	3.6	no
50	20	2	6.6	53.9	yes
50	50	2	13.7	34.6	yes
50	95	2	48.4	30.6	yes
50	145	2	63.5	25.7	yes
50	200	2	96.1	10.1	no
50	350	2	98.4	3.9	no
70	30	2	6.8	57.7	yes
70	70	2	14.3	32.8	yes
70	130	2	19.3	36.8	yes
70	200	2	73.3	26.6	no
70	360	2	86.0	15.0	no
70	460	2	100.0	6.6	no

Table 4: Pinned environment and software versions.

Component	Version
Host	x86-64, 16 physical cores, 125 GiB RAM, NVMe SSD
OS	Debian 13 (trixie)
Kubernetes (k8s), single node	k3s v1.35.5+k3s1 (flannel and built-in netpol disabled)
CNI	Calico v3.29.1 (pod CIDR `10.42.0.0/16`)
Link emulation	Chaos Mesh 2.7.0
Package manager	Helm v3.21.2
Database	PostgreSQL 16
Homeserver	Synapse (element-hq image)