

INDAGO: Test Report TRL 4 / TRL 5

Sovereign local search and analysis engine · verification and validation evidence

DOCUMENT CONTROL	
Document ID	VEC-INDAGO-TRL-2026-01
Version / date	1.1 · 07.10.2026
Issuer	Vectorio GmbH · Weserstr. 15 · 12047 Berlin
Author	Marcel Alexander Mayr, CEO & Key Architect
System under test	INDAGO engine and service · NativeAOT build
Assessment type	Self-assessment · internal test records traceable to version-controlled builds · no independent replication
Scale	EU TRL scale, Horizon Europe General Annexes, Annex B
Classification	Company confidential · for evaluation

1 Result

LEVEL	DEFINITION	BASIS	VERDICT
TRL 4	Technology validated in lab	Integrated system · public benchmarks · controlled sets with controls · 21,726 automated tests	MET
TRL 5	Technology validated in relevant environment	Real corpora > 1M units · CPU-only single machine · production binary · fault injection · continuous operation	MET · CIVIL
TRL 6	Technology demonstrated in relevant environment	Defence-representative environment: air-gap, defence corpora, edge hardware	PLANNED

Key measurements

METRIC	VALUE	CONDITION	REF
Corpus indexed and sealed	1,571,226 chunks	96,929 EPO full-text files · 174.9 GB	E-10
Ingest throughput	6,304 chunks/s	≈ 22.7M chunks/h · ≈ 1.4M patent docs/h	E-10
Query latency, warm p50 / p90 / p99	53.9 / 105.5 / 167.9 ms	1,168,739 chunks · REST API	E-11
Resident memory, idle after seal	17.7 GB	1.57M chunks	E-12
Data loss after SIGKILL during ingest	0	Resume from last commit	E-13
BEIR SciFact nDCG@10	0.6864	BM25 reference 0.662	E-02
Hardware	CPU only	No GPU · no cloud · no outbound connection	\$4

2 Definitions

TERM	MEANING
Chunk	Indexed text unit (section or passage of a document)
Seal	Conversion of ingested data into the immutable, memory-mapped serving format
nDCG@k	Normalised discounted cumulative gain over the top k results · 0–1, higher is better
R@1	Share of queries with the correct item at rank 1
p50 / p90 / p99	Latency percentiles

TERM	MEANING
SIGKILL	Hard process termination with no clean shutdown
NativeAOT	Ahead-of-time compiled native binary · no runtime install, no JIT

3 System under test

CONFIGURATION	
Function	Ingest → index → rank → analyse → verify · every result links to its source · deterministic
Inputs	PDF, Office, text · source code (C#, Python, Java, Go, Rust, C++, TypeScript, ...) · images (OCR, visual retrieval) · mail · structured exports · authorised web and literature sources
Analysis functions	Ranked retrieval · gap discovery · contradiction detection · novelty scoring · cross-domain bridges · isolated topics · unusual pairs · code and security scanners
Interfaces	REST API · MCP server · CLI · desktop app · HTML / PDF / Markdown reports
Connectors	230 source connectors · 18 bundles
Build	.NET NativeAOT single binary · ARM64 / x86-64 release targets
Code base	≈ 411k LOC C# · 40 projects · 17 test projects · 7,798 commits
Builds referenced	288cf0b52 (repo, 07.10.2026) · cbbc62d (production) · 0d2786a (scale run)

4 Test environment

ENVIRONMENT	
Hardware	1 workstation · Apple M5 Max · 18 CPU cores · 128 GB RAM · internal SSD + external data volume
Accelerators	None used (no GPU for indexing, ranking or analysis)
OS	macOS · arm64
Isolation	Scale runs: dedicated service instance (own home, port, registry) · load state recorded per run
Network	Not required for ingest, search or analysis · test data from local storage
Interfaces exercised	REST API · CLI · MCP

5 Test matrix

ID	TEST	DATA	EVALUATION CRITERION	TRL	VERDICT
E-01	Regression suite	21,726 test methods	Automated, repeatable	4	PASS
E-02	Prose retrieval	BEIR ×4	nDCG@10 ≥ BM25 ref. – 0.02 on every set	4	PASS
E-03	Code retrieval	RepoQA · 600 needles	File-level R@1	4	PASS
E-04	Visual retrieval	ViDoRe ×4 · 699 q · 799 pages	nDCG@5 vs single-vector and OCR+BM25	4	PASS
E-05	Gap discovery	Controlled corpus · temporal hold-out	P@30 separated from shuffled-date control	4	PASS
E-06	Contradiction / novelty	Constructed sets · 16–60 cases	Recall, FPR, separation	4	PASS
E-07	Security scanning	DVWA · NodeGoat	Fires on vulnerable, silent on secure variants	4	PASS
E-10	Full-scale ingest + seal	EPO full text · 174.9 GB	> 1M chunks sealed and served	5	PASS
E-11	Latency at scale	1,168,739 chunks	Warm p50 < 100 ms	5	PASS
E-12	Memory / disk	1.57M chunks	Fits one 128 GB machine, seal included	5	PASS

ID	TEST	DATA	EVALUATION CRITERION	TRL	VERDICT
E-13	Fault injection	Live ingest	No committed-data loss · consistent remount	5	PASS
E-14	Corpus-wide analysis	1.17M–1.57M chunks	Completes within time budget	5	PARTIAL
E-15	Second domain	Wikipedia · 70,887 articles	Ingest without adaptation	5	PASS
E-16	Continuous operation	Production instance	Daily use · live updates	5	PASS

6 Results: lab validation (TRL 4)

E-01 Regression suite

ITEM	VALUE
Test methods	21,726
Test projects	17
Coverage areas	ingest · index format · ranking · analysis · interfaces · crash recovery · golden outputs

E-02 Prose retrieval · BEIR · nDCG@10 · 06.10.2026

DATASET	DOMAIN	INDAGO	BM25 REF.	Δ
SciFact	Scientific claims	0.6864	0.662 ¹	+0.024
NFCorpus	Biomedical	0.3153	0.325 ²	-0.010
SciDocs	Citation prediction	0.1537	0.158 ²	-0.004
FiQA-2018	Financial QA	0.2368	0.236 ²	+0.001

¹ BM25 run in the same harness. ² Published BM25 (Thakur et al., BEIR, NeurIPS 2021). Weights fitted on the SciFact and NFCorpus train splits and evaluated on the test splits. FiQA was not used in fitting.

E-03 Code retrieval · RepoQA · R@1

SCOPE	R@1	BUILD
All languages · 600 needles · file level	0.628	06.10.2026
Rust / TypeScript / Python	0.61 / 0.54 / 0.47	29.09.2026
Go / Java / C++	0.42 / 0.42 / 0.41	29.09.2026

E-04 Visual document retrieval · ViDoRe · nDCG@5

SUBSET	ARXIVQA	DOCVQA	INFOVQA	SHIFTPROJECT	MACRO
INDAGO multi-vector · 256M-param model · CPU	0.76	0.58	0.85	0.81	≈ 0.75

699 queries · 799 page images · 2.5–2.7× the single-vector baseline · above OCR+BM25 on every subset.

E-05 Gap discovery · temporal hold-out · 06.09.2026

METRIC	VALUE
Precision@30	0.80 (24/30) · CI95 0.667–0.933
Breakthrough recall	1.0
Median lead time	659.5 days
Negative control (shuffled dates), mean P	0.0275 · CI upper 0.069

Controlled deterministic corpus · no external citation corpus.

E-06 Contradiction and novelty detection

FUNCTION	RESULT	SET
Contradiction detection	Recall 1.00 · FPR 0.00	Constructed pairs
Novelty scoring	Separation +88.9 pp · P@1 = 1	Constructed

Small constructed sets (16–60 cases) · show the mechanism works · not a measure of how well it generalises.

E-07 Security scanning · 31.07.2026

TARGET	CASE	RESULT
DVWA	SQL injection · command injection · unrestricted upload	DETECTED / SECURE SILENT
NodeGoat	NoSQL \$where injection	DETECTED

Recall on known-vulnerable code · no numeric precision is claimed for large real code bases.

7 Results: relevant environment (TRL 5)

E-10 Full-scale ingest and seal · EPO full text · 28.09.2026

METRIC	VALUE	CONDITION
Source	96,929 files · 174.9 GB	EPO full text, weeks 1–15/2026 · multilingual
Chunks sealed and served	1,571,226	1,615,586 slots mounted
Ingest throughput	6,304 chunks/s	Fresh process · quiet machine
Document rate (derived)	≈ 1.4M docs/h	16.2 chunks/doc in this corpus · shorter docs give a higher rate
Build	NativeAOT 0d2786a	Isolated instance

E-11 Query latency at scale · 07.09.2026

METRIC	VALUE	CONDITION
Warm p50	53.9 ms	1,168,739 chunks · 20 distinct uncached queries · REST API
Warm p90	105.5 ms	
Warm p99 (= max)	167.9 ms	
Warm min	28.1 ms	
Cold first query	2.36 s	Index paging in after remount · 4.92 s wall

E-12 Memory and storage · 28.09.2026

METRIC	VALUE	CONDITION
RSS idle after seal	17.7 GB	1.57M chunks · memory-mapped sealed data
Managed heap per chunk	87 B	1.57M chunks
Peak RSS during seal	≈ 40 GB	One-off seal phase
Index on disk	31 GB · ≈ 20 KB/chunk	From 174.9 GB compressed source

E-13 Fault injection

FAULT	OBSERVED	VERDICT
SIGKILL at 807,894 chunks during ingest	Resume from commit point 802,812 · uncommitted tail re-ingested · 0 loss	PASS
Forced shutdown with queued work	Consistent remount · ≈ 3.2K uncommitted chunks re-ingested on next run	PASS
Crash during seal	Phase resumes from last completed range	PASS
Cold remount, 1.6M slots	≈ 2 s	PASS

E-14 Corpus-wide analysis at scale

FUNCTION	OBSERVED	RATING
Contradiction mining	1,168,739 scanned · 969,561 evaluated · 8,979 pairs scored · 232 s (background job)	TRL 5
Isolated topics	12.7 s at 1.57M · identical top 25 after restart	TRL 5
Ranked retrieval	See E-11	TRL 5
Unusual pairs	Time budget exceeded at 1.57M · calibration control failed	TRL 4
Similarity expansion	Time budget exceeded at 1.17M (07.09.2026)	TRL 4

E-15 Second domain · Wikipedia · 25.09.2026

METRIC	VALUE
Input	70,887 articles (1% sample)
Chunks · time	165,821 · ≈ 2 min
Index size	1.34 GB
RSS peak / steady	15.4 GB / 6.0 GB

E-16 Continuous operation · 07.10.2026

METRIC	VALUE
Deployment	Background production service · daily use
Registered workspaces	216
Largest workspace	366,667 chunks · 18,719 files · live updates
Code query p50 (n = 60)	25.5 ms
Build	cbbc62d

Spot check taken just after a cold restart, with tests running in parallel · upper bound, not steady state.

8 Deviations and limitations

ID	DEVIATION	IMPACT ON CLAIM	DISPOSITION
D-1	Single test machine · no independent lab	All values platform-specific	Replication in V3
D-2	Largest measured scale ≈ 1.6M chunks	No claim above 1.6M	Scale step in V4
D-3	Unusual pairs and similarity expansion over time budget	Rated TRL 4 · excluded from TRL 5	Open defect
D-4	Wikipedia: template false positives (contradictions) · 0 gaps found	Domain-specific tuning needed	Open defect
D-5	Prose ranking ≈ BM25 on 3 of 4 BEIR sets · dense re-rank shows no significant gain	Quality parity, not superiority, on general prose	Ranking work ongoing
D-6	E-05 and E-06 use controlled or constructed sets	Precision on real corpora not shown	V2
D-7	Not yet validated: physical air-gap · Linux x86-64 / ARM64 · edge devices · multi-user access control · patent prior-art benchmarks	Outside TRL 5 claim	V1, V3

9 TRL compliance matrix

LEVEL	CRITERION	EVIDENCE	STATUS
4a	Components integrated into one system	E-01, E-1x	MET

LEVEL	CRITERION	EVIDENCE	STATUS
4b	Benchmarks with references or controls	E-02 ... E-07	MET
4c	Automated, repeatable procedures	E-01	MET
4d	Results recorded with build and date	Annex A	MET
5a	Real data at operational volume, > 1M units	E-10, E-11, E-15	MET
5b	Target-class hardware: CPU only, single machine, no cloud	E-10 ... E-13	MET
5c	Production binary through real interfaces	E-10, E-11, E-16	MET
5d	Fault conditions	E-13	MET
5e	Continuous operation	E-16	MET
5f	Deviations recorded	§8	MET
6	Defence-representative environment	V1 ... V5	PLANNED

10 Validation plan to TRL 6

ID	STEP	ACCEPTANCE CRITERION (TO BE AGREED WITH END USER)
V1	Physically air-gapped deployment	Full ingest, search, analysis and reporting with no network interface · 0 outbound connection attempts
V2	Defence-representative corpus (unclassified or synthetic: doctrine, lessons learned, OSINT, imagery, multilingual)	nDCG@10 and R@k on the end user's question set vs an agreed baseline · complete evidence path per result
V3	Hardware: x86-64 server · ARM64 server · low-resource edge device	Latency and memory envelope per hardware class · graceful degradation on edge
V4	Scale to end-user volume (≥ 10M chunks)	Bounded ingest memory · p50/p95 recorded · fault-injection tests repeated
V5	Analyst evaluation with end-user group	Time to evidence and completeness vs the current process · independent observer

A Evidence register

REF	INTERNAL RECORD	BUILD	DATE
E-01	Repository test inventory	288cf0b52	07.10.2026
E-02, E-03	Benchmark ledger: prose ranking, RepoQA	base 501200fc5	06.10.2026
E-04	Benchmark ledger: visual retrieval	sweep record	2026
E-05	Temporal hold-out result record	validation record	06.09.2026
E-07	Corpus security scan report	scan record	31.07.2026
E-10, E-12, E-13	Scale campaign report, 1.57M chunks	0d2786a	28.09.2026
E-11, E-13, E-14	Scale KPI report, 1.17M chunks	scale-run build	07.09.2026
E-15	Wikipedia benchmark report	bench instance	25.09.2026
E-16	Production health and latency check	cbbc62d	07.10.2026

Records available for review on request.

ROLE	NAME	DATE	SIGNATURE
Prepared / approved	Marcel Alexander Mayr · CEO & Key Architect		