REFERENCE NODE R1 - 180-DAY RESEARCH PILOT - PARTNERS SOUGHT

Sovereign AI Security Node R1

Partner brief for on-premises sovereign AI, secure code analysis and agentic-AI assurance

SKOPION is planning a professional workstation target architecture for reproducible local AI, software-security and private cybersecurity workloads. Reference Node R1 is a controlled 180-day research pilot with measurable gates and human validation.

R1 is a target architecture, not a system currently owned. SKOPION does not claim ownership of the named components or a partnership with or endorsement by their manufacturers. Equivalent professional platforms remain in scope.

Why on-premises infrastructure

Sensitive source code, incident data and evidence do not always belong in external AI services. Large local models, long contexts, deterministic security tools and agentic test environments need a robust, isolatable platform. Model opinion does not replace evidence.

Reference Node R1 - target architecture

- CPU - AMD Ryzen Threadripper PRO 9985WX, 64-core class, or an equivalent professional workstation platform.
- GPU - 1 x NVIDIA RTX PRO 6000 Blackwell, 96 GB ECC VRAM; 96 GB Max-Q as a power-efficient alternative.
- RAM - 256 GB DDR5 ECC RDIMM.
- Storage - 2 TB NVMe system plus 8 TB NVMe AI/data, or an equivalent securely segmented layout.
- Network - 10 GbE target; controlled data paths and no uncontrolled remote access.
- OS - Ubuntu 24.04 LTS; Windows 11 Pro / WSL2 only where tooling requires it.

180-day pilot

- Phase 1 - Baseline & Hardening: deployment, isolation and reproducibility.
- Phase 2 - Local AI: LLM/VLM inference, context, memory and data locality.
- Phase 3 - Software Security: secure code analysis, tool correlation and triage.
- Phase 4 - Agentic AI: authorised tool-use and multi-agent assurance.
- Phase 5 - Optimisation: performance, power, thermals and stability.
- Phase 6 - Report: benchmarks, architecture conclusions and a safe demonstrator.

Workloads and measurable gates

- Local AI: 70B-class where architecture and quantisation permit; long context, embeddings, reranking and private RAG.
- Secure Software Analysis and Fuzzing Support: CPU compiles/tests; AI assists clustering, harness analysis and triage.
- Agentic AI Assurance: prompt injection, excessive agency, RAG and tool-use safety.
- OSINT/EASM/Incident: public sources, classification and evidence structuring.
- LoRA/QLoRA, domain embeddings, rerankers and classifiers where technically appropriate.
- Measure reproducibility, latency, throughput, resources, stability, power, data locality and false positives.

Technology partnership

- Models: evaluation workstation, GPU/CPU/platform, 180-day loan, component contribution or project pricing.
- SKOPION can deliver reproducible benchmarks plus compatibility, stability and architecture feedback.
- Optional: safe demonstrator and a jointly approved anonymised case study.
- Public acknowledgement only after written approval.
- No customer data, credentials, unpublished vulnerabilities or uncontrolled remote access.

Explicit boundaries

No unauthorised or production-disrupting testing. No promises of zero-days, finding counts or payouts. No funding, programme, ownership or partnership claim without published confirmation.

Discuss a technology partnership: https://skopion.de/en?topic=research-partnership#contact