From features to foundations

The Innovation Landscape for Software & SaaS Engineering Teams

Executive Summary

Software R&D in mid-market and scale-up SaaS is shifting from features to foundations. Where the last decade rewarded velocity (ship faster, iterate faster) 2026 rewards engineering work that holds up under regulation, scales economically under AI workloads, and is visible enough on the P&L to be governed. This trend analysis is written for engineering leaders, CTOs, VP Engineering, Engineering Directors, setting the technical direction in software and SaaS companies across the Netherlands and Germany.

Three forces are reshaping the roadmap:

1. **Regulation has become an architectural input.** The CRA, NIS2, AI Act, Data Act, DORA, and Germany's NIS2UmsuCG are setting design-stage decisions about data, models, identity, and supply chains.

2. **AI has stopped being a feature and started being infrastructure.** The interesting question is "what do we build underneath AI to make it reliable, auditable, and economical at scale?"

3. **The cost of software is on the boardroom agenda.** Cloud, GPU, licensing shocks, and inference carbon are pushing engineering leaders to treat infrastructure as a P&L item.

The strategic question for your team is no longer "what should we build?" but "what should we make foundational?"

Most of the foundational work in this analysis qualifies as non-dilutive capital under WBSO, FZulG, ZIM, Eurostars, Horizon Europe, or other European programmes. Your team has already earned most of it. The limiting factor is rarely eligibility. It is process.

Five engineering trends define 2026

01

AI-Native Software Engineering:

Agentic systems, RAG infrastructure, fine-tuning trade-offs, and the evaluation layer that makes AI features production-grade.

02

Cybersecurity as Engineering Infrastructure:

CRA reporting from 11 September 2026; NIS2 enforcement in Germany now active. Supply-chain security and SBOM tooling at the centre.

03

Data Engineering at Scale and Privacy by Design:

Streaming-first architecture, differential privacy, and AI Act technical compliance against the 2 August 2026 high-risk deadline.

04

Cloud Architecture and Sustainable Infrastructure:

Sovereign-by-design, edge AI inference, and the open-source licensing shock reshaping infrastructure roadmaps.

05

Platform Engineering and Developer Productivity:

Internal developer platforms, OpenTelemetry, and FinOps tooling embedded in developer workflows.

FIVE EMERGING THEMES are visible in where public innovation funding is flowing: quantum-safe software infrastructure, AI governance and auditability tooling, trusted data spaces, energy-aware computing, and digital twins at software scale.

TREND 1

AI-Native Software Engineering

From AI-assisted to AI-native

Teams that bolted AI on as a feature layer are rebuilding. The question is no longer whether to adopt AI tooling, but how deeply to embed it in the architecture of the product itself. In mid-market SaaS, where teams compete on product intelligence, shipping reliable, explainable, cost-effective AI is now as foundational as shipping secure, scalable infrastructure.

Key developments

Agentic systems in production.

Multi-step AI agents are moving from demo to production, with engineering investment in orchestration, memory, and reliable tool-use APIs. In practice, the novel work is rarely the model itself. It is the layer above it: routing queries to the right agent, maintaining conversation state across turns, preventing unwanted re-routing mid-session, and chaining tool outputs across sequential API calls. Auditable agentic workflows remain an open engineering problem.

RAG as core infrastructure.

Retrieval-Augmented Generation has moved past proof-of-concept; the open problems are retrieval quality at scale, hybrid search, and incremental indexing. Standard vector similarity also breaks down for synthesis tasks across multiple documents: context-aware chunking, semantic deduplication, and knowledge-base CI/CD are emerging as production patterns where simple RAG falls short.

Fine-tune, distill, or call the frontier.

Architectural decisions about when to fine-tune smaller open-weight models, when to call frontier APIs, and when distillation makes sense rarely converge to a single answer.

AI evaluation and observability.

As AI components become business-critical where commercial tooling is emerging but most teams still build significant custom infrastructure

From the field: Large-scale e-commerce platform

NL, 290 S&O engineers

After validating that a single generic LLM cannot adequately handle the full range of domain-specific customer queries, this company moved to a multi-agent architecture: specialised Virtual Agents each with their own knowledge base, a self-developed intent interpreter enriched with order and contact history, and a conversation state manager with lock mechanisms preventing unwanted re-routing mid-session. The engineering challenge is not the LLM itself: it is the orchestration infrastructure that makes multi-agent systems production-grade. A GPU-aware Kubernetes layer (NVIDIA MIG partitioning, MLflow-based predictive autoscaling) manages thousands of concurrent sessions. This is where the R&D novelty lies, and what WBSO funds.

"One generic LLM is not enough. The interesting R&D is in the layer above it: how you route, remember, evaluate and coordinate agents at production scale."

TREND 2

Cybersecurity as Engineering Infrastructure

Security shifts left, driven by a regulatory clock

Cybersecurity is no longer a checkbox at the end of the development cycle. The CRA, NIS2, and Germany's NIS2-Umsetzungsgesetz are forcing teams to rearchitect security from the ground up. The regulatory calendar is doing the heavy lifting:

CRA	entered into force 10 December 2024. Notification of conformity assessment bodies applies from 11 June 2026; vulnerability and incident reporting from 11 September 2026; main obligations from 11 December 2027.
NIS2	transposition due 17 October 2024. Germany's NIS2UmsuCG entered into force 6 December 2025; BSI portal registration began 6 January 2026. Sanctions reach €10m or 2% of global turnover for "particularly important entities".

For teams that have not started the architectural work, the runway is shorter than it looks.

Key developments

Software supply chain security.	SBOM generation, dependency vulnerability scanning, and provenance verification (SLSA). Hardened CI/CD and reproducible builds are baseline expectations from enterprise customers and a hard requirement under the CRA.
Shift-left security tooling.	Static and dynamic analysis, IaC scanning, and secrets detection in the development workflow. The engineering challenge is triage at scale without blocking velocity.
Zero-trust architecture.	Identity-and-context-based access control across authentication, service mesh, and network segmentation. SPIFFE/SPIRE, mutual TLS at scale, and workload identity management are active areas.
Post-quantum cryptography migration.	NIST finalised FIPS 203, 204 and 205 in August 2024; both the Dutch NCSC and the German BSI have issued migration guidance.

From the field: Physical security systems provider

International, 100–200 engineers

This group is re-architecting its core security platforms from monolithic C/legacy systems to containerised, cloud-hybrid microservices. There is one hard constraint: security systems must respond in milliseconds, doors must open in emergencies. The team implemented Zero Trust architecture for hardware-constrained edge devices, developed custom C++ SRTP modules with packet handling and key rotation from scratch (no suitable library existed), and is migrating legacy encryption to meet EU RED certification requirements. The CRA and NIS2 are not compliance checkboxes. They are driving a multi-year re-architecture program.

"Cybersecurity in 2026 is not a feature. It is the foundation. This company is building Zero Trust into hardware-constrained edge devices, implementing SRTP from scratch because no suitable library exists, and migrating legacy encryption to meet EU RED certification. The CRA and NIS2 are driving multi-year re-architecture programmes."

TREND 3

Data Engineering at Scale and Privacy by Design

Real-time data architectures and privacy-preserving computation

Most data platforms were built for batch, not streaming. Most were not designed for GDPR-compliant AI workloads at all. AI's appetite for high-quality training and inference data, the rise of streaming use cases, and tightening privacy regulation are now forcing a new wave of platform R&D. The EU Data Act has been in application since 12 September 2025, with rules on connected-product data access and data-processing-service switching phasing in through 2026 and early 2027.

Key developments

Streaming-first architectures.

Apache Kafka, Flink, RisingWave and Redpanda enabling sub-second latency. Migrating from batch to streaming requires significant rearchitecting.

Data mesh and federated ownership.

Larger organisations moving from centralised data lakes to domain-owned data products, with governance tooling, data contracts, and discovery infrastructure as multi-year investments.

Differential privacy and synthetic data.

Increasingly used to enable analytics and ML training on sensitive user data without exposing the raw record.

AI Act compliance engineering.

Explainability, model documentation, data quality, and human oversight for high-risk systems with the 2 August 2026 applicability date approaching. (The Commission's Digital Omnibus on AI, proposed November 2025, would defer this to 2 December 2027; until that is formally adopted, the original date stands.)

From the field: AI-native software builder

NL, 5 S&O engineers

This company builds bespoke AI solutions, including assistants, chatbots, and code generation tooling, delivered as both SaaS and on-premise. The core R&D challenge: a framework that is LLM-agnostic, AVG-compliant, and deployable in cloud and on-premise environments without sacrificing response quality. Engineering problems solved: context loss from random chunking (Load-Transform-Embed algorithm with source-type-aware ingestion), GDPR-incompatible external LLM processing (isolated private LLM layer with per-organisation token tracking), and sharing AI assistants between organisations without privacy leakage (pseudonymisation module retaining context without storing personal identifiers). Company size is not a predictor of R&D depth.

"Small AI-native companies are confronting the same architectural problems as large enterprises, often with fewer resources and higher urgency. LLM routing, privacy-preserving RAG, and cross-channel latency constraints are not edge cases. They are central engineering problems in 2026."

TREND 4

Cloud Architecture and Sustainable Infrastructure

Sovereignty, edge inference, and the open-source licensing shock

Cloud bills ballooned when AI inference hit production. Open-source tools started charging. Vendor lock-in moved from technical debt to boardroom risk. Infrastructure engineering is being reshaped by three forces: cloud cost efficiency and vendor independence, edge and hybrid deployments driven by AI inference and data sovereignty, and pressure to reduce the environmental footprint of software. In Germany in particular, sovereignty and energy efficiency have moved from preferences to procurement requirements.

The **open-source licensing shock** is a recurring planning input. HashiCorp moved Terraform from MPL 2.0 to the Business Source License on 10 August 2023, prompting the OpenTofu fork; HashiCorp ended Terraform OSS under BSL after July 2025. Redis and Elastic have made parallel moves. The migration projects these triggers (fork adoption, drop-in replacements, internal hardening) are now standard items on infrastructure roadmaps.

Key developments

Sovereign cloud deployments.

EU data residency requirements and customer demand for GDPR-compliant hosting drive engineering investment in sovereign architectures, encryption key management, and cross-border data governance. German teams navigate BSI IT-Grundschutz and BSI C5; Dutch teams the BIO baseline.

Multi-cloud and cloud-agnostic infrastructure.

Kubernetes-based abstractions, service mesh, and IaC modules to enable workload portability across AWS, Azure, GCP.

Edge AI inference.

Model compression, quantisation, and edge runtime optimisation for IoT, mobile, and CDN-side deployment.

Green software engineering.

Energy consumption dashboards, carbon-aware compute region selection, and energy-efficient algorithms are increasingly demanded by enterprise customers.

From the field: Industrial IoT & automation platform

NL, 250 S&O engineers

This company's IoT platform connects thousands of devices across industrial environments, including water utility infrastructure and highway bridge management systems. The shift to critical infrastructure customers in 2026 triggered a step-change in engineering requirements: NIS2-adjacent compliance, vendor-agnostic architecture (existing fleet management frameworks assume homogeneous equipment), and a permission arbitration service for autonomous terminal trucks making decisions in under 200ms. A centralised cloud mechanism was too slow, so the team built a lock-free spatial state manager. Sovereign and edge constraints are live requirements, not roadmap items.

"Critical infrastructure deployments fundamentally change the engineering constraints. Real-time anomaly detection, compliance-grade data integrity, and sub-200ms decision loops are not features. They are architectural foundations."

TREND 5

Platform Engineering and Developer Productivity

Internal platforms become a strategic engineering investment

Without a platform team, every product engineer carries ops overhead. Platform engineering has matured from trend to standard organisational pattern in mid-market and scale-up software companies. Dedicated teams building internal developer platforms (IDPs) reduce cognitive load on product engineers, compress deployment cycles, and create a sustainable foundation for technical scaling.

Key developments

Internal Developer Platforms (IDPs).

Self-service abstractions enabling product engineers to deploy and operate without deep ops expertise. Real engineering effort goes into customisation, workflow integration, and governance. For AI-heavy products, the IDP must also surface evaluation pipelines, per-feature inference cost, and latency observability. Without this, product engineers operate blind on the components that matter most.

FinOps embedded in developer workflows.

Per-feature, per-service cloud spend surfaced in real time and integrated with deployment pipelines are increasingly competitive differentiators for AI-heavy products.

OpenTelemetry standardisation.

Convergence as the standard for distributed traces, metrics, and logs. Custom instrumentation libraries and collector pipelines for proprietary systems are real engineering work.

WebAssembly server-side.

Plugin systems, edge computing, and multi-tenant isolation as early production use cases.

From the field: Industrial IoT & automation platform

(NL) Agentic workflow infrastructure

The same company encountered a platform engineering problem specific to AI-native products: standard Azure Functions are deterministic and linear, incompatible with the dynamic branching and event-driven decisions that agentic workflows require. The solution was a custom workflow graph compiler and a replay-safe scheduler that records agent decisions as explicit events in a versioned graph, making agent behaviour auditable, reproducible and debuggable. This is the IDP layer that AI-heavy products need, and that most teams are still building themselves.

"Agentic workflows in standard serverless functions don't work. Once you have agents making dynamic decisions, you need infrastructure that records those decisions, replays them, and makes them inspectable. That's not a nice-to-have: it's the difference between a prototype and a production system."

LOOKING AHEAD

Emerging themes shaped by where innovation funding is flowing

Public innovation funding signals where investment is expected to flow. Four themes are visibly accelerating into mainstream R&D priorities over the next 24 to 36 months.

Quantum-safe software infrastructure.

NIST finalized FIPS 203/204/205 in August 2024. The EU's Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography (June 2025) requires Member States to begin transition by end of 2026, complete it for critical use cases by 2030, and for all systems by 2035. The EU Quantum Act proposal is expected mid-2026. Funding flows via the Quantum Flagship (>€400m phase 2), Quantum Delta NL, and German BMBF programmes.

AI governance and auditability tooling.

The AI Act creates a compliance market without mature tooling: audit trails, model cards, fairness testing, human-in-the-loop oversight. The AI Office's role is guidelines and sandboxes; direct funding flows via the Apply AI Strategy (launched October 2025), the Digital Europe Program, and Horizon Europe Cluster 4.

Trusted data spaces and interoperability.

The EU Data Act and Gaia-X provide the regulatory and technical foundation. The BMWE has funded 11 Gaia-X projects via the Gaia-X Hub Germany since 2022; further calls continue under the Data Spaces theme of Horizon Europe and the Digital Europe Program.

Energy-aware and sustainable computing.

Horizon Europe Cluster 5 ("Climate, Energy and Mobility") has allocated €2.2bn for the 2026–2027 work programme; concrete software-related calls exist (e.g. 2ZERO's Energy-efficient software-defined EVs, HORIZON-CL5-2026-05-D5-02). Pure-SaaS green computing funding remains thinner, with the discipline currently driven by the Green Software Foundation and the major cloud providers.

Where the smart money is going — programmes and rates

Most of the engineering work in this analysis qualifies for public innovation funding under Dutch and German definitions. We track these programmes professionally; here is a current snapshot of the capital instruments most relevant to a software roadmap in 2026.

Programme	What it is	Rate (2026)	Cap
WBSO (NL)	R&D wage tax credit	36% on first €380k of qualifying personnel cost; 16% above; 50%for starters on first bracket	No formal cap
RDA (NL)	Additionalcorporate tax deduction on top of WBSO	Variable, depends on R&D structure	—
Forschungszulage/ FZulG (DE)	R&D tax credit	25% standard / 35% for SMEs; new from 2026: 20% flat rate for overheads in addition to personnel and contract research costs	€12 million R&D base per year (up from €10m for projects starting on or after 1 January 2026); max benefit €3m / €4.2m per year
ZIM (DE)	Innovation grant for SMEs	25–55% (R&D projects, max €690k) · 40–60% (international cooperation)	€690k per R&D project · €560k per company in cooperation projects · €280k for research institutions
Eurostars 3	Joint European programme for R&D-performing SMEs, well-suited for cross-border NL/DE projects	Up to 50% (NL and DE SMEs)	max €500k per project share; SME budget must be ≥50% of total project cost

WHERE EACH TREND IN THIS ANALYSIS TENDS TO FIT BEST

AI-NATIVE, COMPLIANCE-AS-PLATFORM:	WBSO + RDA (NL); FZulG, ZIM (DE); Horizon Europe / EIC
SECURITY:	WBSO (NL); FZulG, ZIM (DE); Horizon Europe
DATA & PRIVACY:	WBSO + RDA (NL); FZulG, ZIM (DE); Horizon Europe
CLOUD & SUSTAINABILITY:	WBSO (NL); FZulG, ZIM (DE); Horizon Europe Cluster 4 (“Digital, Industry and Space”) and Cluster 5 (“Climate, Energy and Mobility”).
PLATFORM ENGINEERING:	WBSO + RDA (NL); FZulG, ZIM (DE); Horizon Europe / EIC

For most engineering leaders, the limiting factor is not eligibility. Most software teams qualify for at least one major programme. The limiting factor is process: documentation that survives audit, claims that capture the full scope of qualifying work, and a roadmap-to-funding workflow that runs continuously rather than once a year. That is the part we run alongside engineering teams every day. The capital is forecastable. The question is whether you are claiming it.

Your roadmap has non-dilutive capital in it. Let us show you where.

We offer a focused 30-minute R&D Funding check for software companies. We take your current engineering roadmap or a recent sprint backlog and tell you, on the call, which work qualifies under WBSO, FZulG, ZIM, Eurostars or other programmes — what fits best for you.

No prep required from your side beyond a roadmap or backlog excerpt.

Ignite Group — 2025 in numbers

REALISED FUNDING	APPLICATION SUCCESS RATE	COMPLIANCE UNDER AUDIT
€500M	95%	99%

International full-service partner for innovation funding, working alongside engineering and finance teams to secure non-dilutive capital as a forecastable, recurring capital instrument, with audit-ready documentation and minimal internal overhead.

Rates, caps, and regulatory milestones in this paper are current as of May 2026 based on publicly available primary sources (RVO, BMF/BMWE, Eureka, EUR-Lex, NIST, BSI). Funding rates are reviewed annually and may change; please confirm current values with us before basing financial planning on the figures in this document.

Ignite Group
info@ignite-group.de
+49 6894 999560