ATTENDANCE PROPOSAL

APIcon Berlin 2026

Part of AI-Native Week Berlin

16–20 November 2026 · Maritim proArte Hotel, Friedrichstraße 151, Berlin

In Short

APIcon Berlin 2026 runs 16–20 November, co-located with AI-Native Week Berlin (alongside MLcon, VibeKode, and AI Engineering Summit). This edition's program is dense with a single shift: AI agents calling, orchestrating, and governing APIs — not just people. Governance, security, and design content are all built around that shift.

For me: I'll come back with concrete, retrieved-from-source patterns for building and governing APIs that AI agents can call directly — contract-first design for agent orchestration, the guardrails that keep agentic traffic from breaking production, and the design craft (type-safe contracts, media handling, endpoint structure) that still underpins all of it.

For the team: One trip returns a shared playbook for agent-ready API governance and security that the rest of the team can build on — details below.

The ask: Approve the two-day core conference (18–19 November); if budget allows, the AI-Native Week Full Week Pass covers the same trip plus MLcon, VibeKode, and AI Engineering Summit, and extends the recordings below to all four programs.

Topics at a glance:

- Lead theme — building APIs that AI agents can actually call, orchestrate, and trust
- Governing agentic systems as they move from pilot into production
- Defending LLM- and agent-driven systems against a newer class of attack
- Sharpening core API design and delivery craft that still underpins all of it

What the Team Gets

One attendee, whole-team return: a short recap session for the team on what's changed in agent-facing API design, an internal briefing translating the governance and security patterns into a checklist we can apply to our own APIs, and a lightweight adoption plan for whichever pattern fits our roadmap. Sessions are recorded, so anything missed live is still reachable afterward — and if the Full Week Pass is approved, that recording access extends across all four AI-Native Week programs, not just this one.

The Program

Four themes run through this program — each opens with why it matters this year, then the sessions that back it up.

1. Building for a Machine Caller

AI agents are no longer edge cases for an API — they're becoming primary consumers, calling endpoints autonomously, chaining services, and generating integrations no developer explicitly designed for. OpenAPI can describe what an endpoint does, but not how to sequence it into a business process, and an agent treated as "just a faster API client" breaks in ways a human integrator wouldn't. This year's program addresses that gap directly.

Session	What It Shows
Making Your APIs AI-Ready — Mike Amundsen	Five concrete shifts (intentionality, context, standardization, discovery, adaptability) to make an API ecosystem machine-ready.
What AI Agents Are Teaching Us About Integration — Dr. Erik Wilde	Why agents consume integration surfaces at runtime with no guide to read, and what that demands of API design.
Building Agent Interfaces: Designing for a User That Is a Machine — Michael Hablich	Treating the agent as its own kind of user, with its own strengths and failure modes, rather than a faster client.
Beyond OpenAPI: Workflows That AI Agents Can Actually Execute — Frank Kilcommins	Why endpoint-level specs aren't enough for multi-step agent workflows, and what orchestration needs instead.

> ***Bottom line:*** *If agents are already calling our APIs unsupervised, the design choices that make that safe are worth making on purpose, not discovering after an incident.*

2. Governing the Agentic Enterprise

As agent adoption moves from pilot to production, the autonomous, multi-step nature of agent-to-agent and agent-to-API interaction raises governance questions traditional API management never had to answer. Several sessions this year tackle that directly — from the operating principles of agent governance to what it takes to run agentic AI as a genuine enterprise capability.

Session	What It Shows
Governance in the Age of AI Agents — Frederick Elewaut	Core principles of how AI agents interact with each other and with applications, and where governance has to sit.
Beyond the Hype: How to Lead, Govern, and Scale AI Responsibly — David Roldán Martínez (Keynote)	Why weak AI governance turns impressive systems risky and opaque, and what scaling responsibly requires.
Transforming into an Agentic Enterprise — Matthias Biehl	Moving agentic AI from sporadic chatbot use to a deeply embedded, governed part of enterprise process.
Workshop: Architecting AI-Ready APIs with Agent Skills — Dr. Erik Wilde & Frank Kilcommins	A hands-on masterclass designing governed APIs for AI-driven integrations that weren't explicitly hand-built.

> ***Bottom line:*** *Governance built after agents are already in production costs a rewrite; governance designed alongside them costs a workshop day.*

3. AI & API Security: the New Attack Surface

LLMs and agents introduce risks traditional API security controls weren't built for: prompt injection, jailbreaks, data leakage, and agentic clients that quietly degrade performance or blow through rate limits at scale. The program pairs the defensive fundamentals with the AI-specific failure modes.

Session	What It Shows
Hack Me If You Can: Designing Unbreakable LLM Guardrails — Cansu Kavili-Örnek (Keynote)	Layered guardrails for GenAI systems — prompt-injection detection, content safety, and continuous red-teaming.
Workshop: API Security — From Attacks to APIOps Defense — Thomas Bayer	A hands-on turn as attacker and defender against real API vulnerabilities, then how to defend at scale.
Mind Your Manners: API Etiquette — Michael Dowden	Why poorly optimized clients — increasingly agentic ones — degrade performance and inflate cloud bills.

> ***Bottom line:*** *A single avoidable prompt-injection or rate-limit incident costs far more than the days it takes to build the guardrails that prevent it.*

4. API Design & Delivery Craft

None of the agentic-era content replaces the fundamentals — it depends on them. This year's craft sessions cover the recurring, unglamorous problems (fragmented ownership, type drift between frontend and backend, media handling, minimal endpoint design) that determine whether an API is solid enough to put an agent — or a person — in front of.

Session	What It Shows
Fixing API Chaos: A Practical, Open Method to Structure API Work — Marjukka Niinioja	Why most API problems are organizational, not technical, and a method to structure the work across teams.
Closing the Gap: Seamless Frontend–Backend APIs with tRPC — Mario Goller	A TypeScript-native, end-to-end type-safe approach that removes redundant schema synchronization.
Images in APIs: Databases, Blobs, and Practices — Dino Esposito	The real performance, cost, and maintainability trade-offs behind the "just use blob storage" default.
Hands-On Minimal API: Building Robust Endpoints — Dino Esposito	Packaging Minimal API endpoints for production, beyond the one-liner demo.

> ***Bottom line:*** *Agent-ready is built on top of well-designed — the craft sessions are what make the agentic-era patterns actually shippable.*

Also on the program: AI Trustworthy by Design: A Practical Path to ISO/IEC 42001 (David Roldán Martínez), From APIs to Outcomes: Building AI-Powered Data Ecosystems (Heather Xiao), APIs Are the New Product Surface (Alpesh Pawar), Workshop: Intentional API Design (Matthias Biehl), and the two-day Hands-On Bootcamp: Vibe Coding & GenAI Development (John Davies).

Sources

Every session below is drawn from the conference's own confirmed (Live) program schedule.

- Making Your APIs AI-Ready
- What AI Agents Are Teaching Us About Integration
- Building Agent Interfaces: Designing for a User That Is a Machine
- Beyond OpenAPI: Workflows That AI Agents Can Actually Execute
- Governance in the Age of AI Agents
- Beyond the Hype: How to Lead, Govern, and Scale AI Responsibly
- Transforming into an Agentic Enterprise
- Workshop: Architecting AI-Ready APIs with Agent Skills
- Hack Me If You Can: Designing Unbreakable LLM Guardrails
- Workshop: API Security — From Attacks to APIOps Defense
- Mind Your Manners: API Etiquette
- Fixing API Chaos: A Practical, Open Method to Structure API Work
- Closing the Gap: Seamless Frontend–Backend APIs with tRPC
- Images in APIs: Databases, Blobs, and Practices
- Hands-On Minimal API: Building Robust Endpoints