



Pinecone Expands in Europe with New Frankfurt Cloud Region, Delivering the Knowledge Infrastructure for AI to Central European Enterprises

Descrizione

COMUNICATO STAMPA - CONTENUTO PROMOZIONALE

Alongside new availability on AWS eu-central-1, Pinecone announces Nexus knowledge engine, KnowQL query language, Marketplace, Builder tier, and native full-text search

FRANKFURT, Germany, May 5, 2026 /PRNewswire/ - Pinecone, the knowledge infrastructure for AI at scale, today announced its expansion into the AWS Europe (Frankfurt) Region (eu-central-1), bringing its full serverless vector database and knowledge infrastructure to central Europe for the first time. The new region enables organizations across Germany, Austria, Switzerland, and the broader European market to run AI workloads with low-latency performance and local data residency.

The Frankfurt region launch coincides with a series of major product announcements from Pinecone, including the introduction of Pinecone Nexus, a knowledge engine for AI agents; KnowQL, a declarative query language for agentic retrieval; the Pinecone Marketplace with more than 90 production-ready knowledge applications; a new Builder tier at \$20 per month; native full-text search in the core Pinecone database; and Dedicated Read Nodes for high-throughput, low-latency production workloads delivering up to 97% cost reduction at scale.

"The best knowledge infrastructure should be accessible to every builder, in every region," said Ash Ashutosh, CEO of Pinecone. "Expanding to Frankfurt gives European organizations the ability to build on the same infrastructure that more than 9,000 customers worldwide rely on - with the data residency, low latency, and regulatory alignment that enterprises in the region require. Combined with today's launches of Nexus, KnowQL, Marketplace, and our new Builder tier, we are delivering the most comprehensive knowledge infrastructure for AI at a scale and price point that removes every barrier to building."

Pinecone Nexus: The Knowledge Engine for Agents

As AI shifts from assistive applications to autonomous agents, the demands on knowledge infrastructure are changing fundamentally. Roughly 85% of an agent's effort today is spent on context retrieval, resulting in task completion rates of just 50-60%, unpredictable completion times, and runaway token costs. Pinecone Nexus addresses this by moving reasoning upstream, from retrieval to knowledge compilation. Its context compiler transforms raw data into task-optimized artifacts that agents consume directly, while its composable retriever serves those artifacts with low latency, per-field citations, and deterministic conflict resolution. Early results show up to 90% reduction in token usage, task completion rates above 90%, and 30x faster time-to-completion.

At the core of Nexus is KnowQL, a declarative query language that gives agents a universal interface for knowledge access, replacing custom tool definitions and bespoke glue code with a single call that specifies output format, citation requirements, and latency budgets, and returns trusted, structured, grounded knowledge.

Pinecone Marketplace

The Pinecone Marketplace is a catalog of fully working, production-ready knowledge applications that users can deploy, customize, and run immediately. It launches with more than 90 production-ready solutions across sales and revenue, insurance, real estate, legal and compliance, people and HR, customer support, and more categories. Built by Pinecone and partners, these applications require no infrastructure assembly. The Marketplace is free at launch, with partner-built commercial solutions coming soon.

New Pricing and Infrastructure Updates

Pinecone's new Builder tier provides full access to its production-grade infrastructure for \$20 per month and free support. Dedicated Read Nodes offer provisioned read capacity with warm data paths and fixed hourly pricing, delivering 77-97% cost reduction at scale for sustained high-throughput workloads. Bring Your Own Cloud (BYOC) deploys Pinecone fully managed within a customer's own cloud environment for organizations with data residency or regulatory requirements. Additionally, Pinecone has integrated native full-text search into its core database, now in public preview, enabling hybrid retrieval that combines semantic breadth with exact-match precision for agent-scale query volume.

The Frankfurt region is available now. It joins Pinecone's existing serverless regions on AWS in the United States, Ireland, and Asia, with additional regions planned.

Get started here. For more detailed information on today's announcements, dive into the blog.

About Pinecone

Pinecone is the knowledge infrastructure for AI at scale. Its leading vector database and knowledge engine, Pinecone Nexus, power accurate, performant AI applications for more than 9,000 customers and 800,000 developers worldwide. Pinecone's mission is to make AI knowledgeable. For more information, visit pinecone.io.

Media Contact: Mike Sefanov mike.s@pinecone.io Sr. Director, Communications

Logo: https://mma.prnewswire.com/media/2418074/Pinecone_Systems_Inc_Logo.jpg

View original content: <https://www.prnewswire.co.uk/news-releases/pinecone-expands-in-europe-with-new-frankfurt-cloud-region-delivering-the-knowledge-infrastructure-for-ai-to-central-european-enterprises-302761619.html>

Copyright 2026 PR Newswire. All Rights Reserved.

COMUNICATO STAMPA â?? CONTENUTO PROMOZIONALE: Immediapress Ã? un servizio di diffusione di comunicati stampa in testo originale redatto direttamente dallâ??ente che lo emette. Lâ??Adnkronos e Immediapress non sono responsabili per i contenuti dei comunicati trasmessi

â??

[immediapress/pr-newswire](https://www.immediapress/pr-newswire)

Categoria

1. Comunicati

Tag

1. ImmediaPress

Data di creazione

Maggio 5, 2026

Autore

redazione

default watermark