INNOVATION & RESEARCHING

sdg

Epoch #7
Innovation Radar

orbitae

Epoch #7
Innovation Radar

EXECUTIVE SUMMARY

What are the most relevant trends we are witnessing right now?

AI Fueled Living Software

Agents that self-heal in production detecting bugs via logs, reasoning about fixes, and deploying corrections autonomously while humans design the strategy.

Product Agents & Low-Code Platforms

Platform vendors offer pre-built agents and low-code builders with instant time-to-value, but flexibility, cross-platform scope, and complex orchestration remain limited.

Foundation Models for Tabular Data

Pre-trained time-series models (Chronos, TimeGPT, TiRex) enable zero-shot forecasting with dramatically lower development effort and a compelling cost-to-accuracy tradeoff.

Content Generation for Marketing

AI content shifts from volume to data-driven relevance and personalization; content must now be optimized for AI-powered discovery channels (GEO) where purchases increasingly begin.

THE INNOVATION QUADRANTS

Discover how each segment is evolving

Architecture | Data | AI | Business

INNOVATION RADAR EPOCH #7

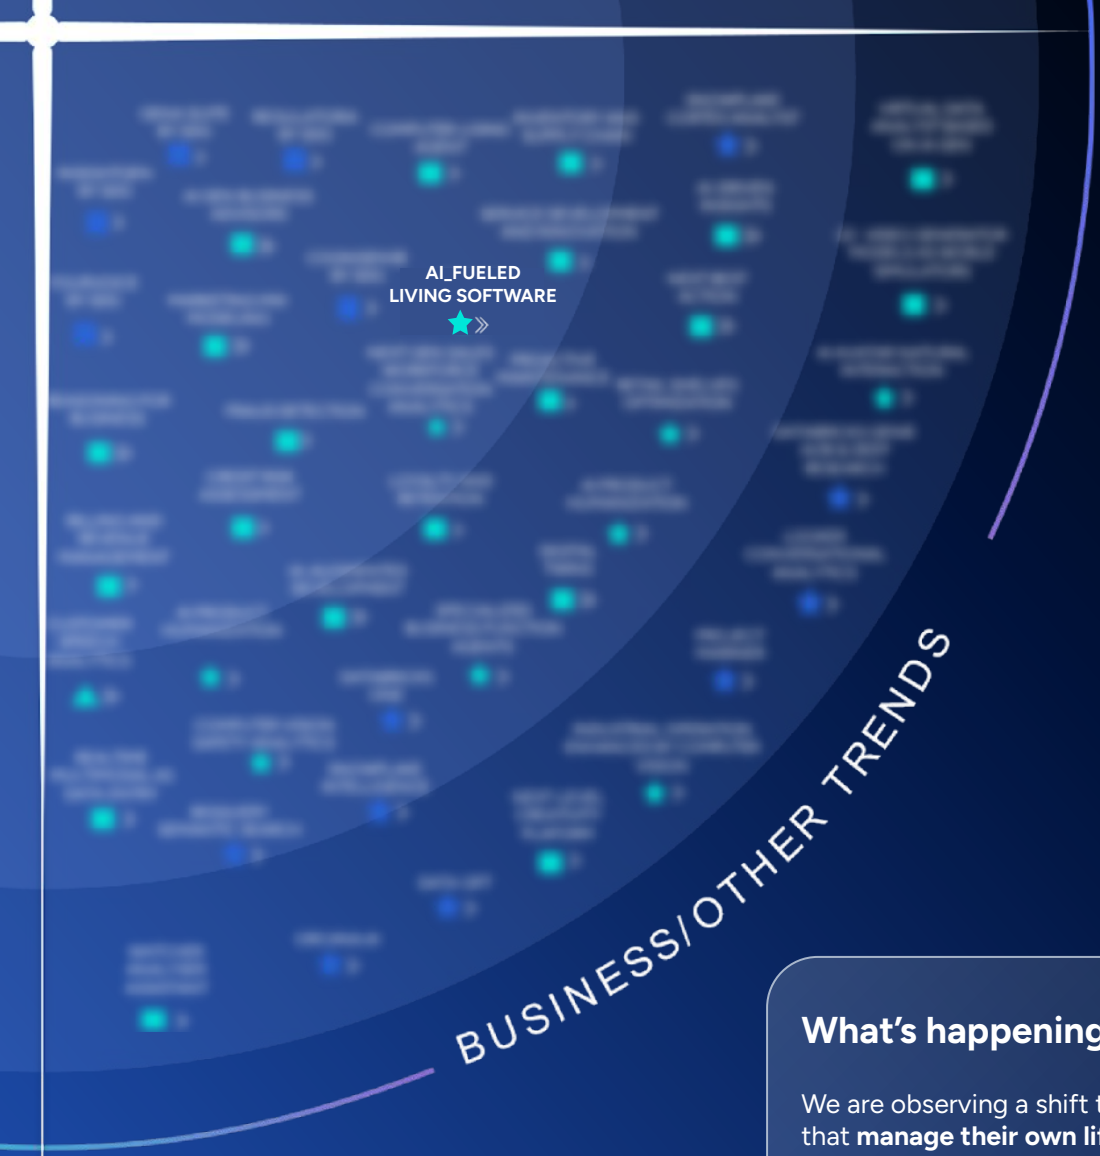

INNOVATION RADAR EPOCH #7

What's happening?

We are observing a shift toward systems that **manage their own lifecycle autonomously**.

Instead of traditional monitoring, these agents use **high-quality log data to detect anomalies** in production environments.

By applying advanced reasoning to identify root causes, they can **deploy corrections and patches without human help**.

This evolution moves development from a human-heavy labor model to a compute-based reasoning model.

While **humans remain the designers of high-level strategy**, the AI manages the operational intricacies, ensuring enterprise software remains resilient and "alive."

Why it matters?

This trend represents a fundamental shift in software reliability and economics.

By **enabling systems to self-heal,** it becomes possible to **minimize downtime** and transition the technical focus from manual firefighting to high-level strategic design, significantly **lowering long-term operational costs**.

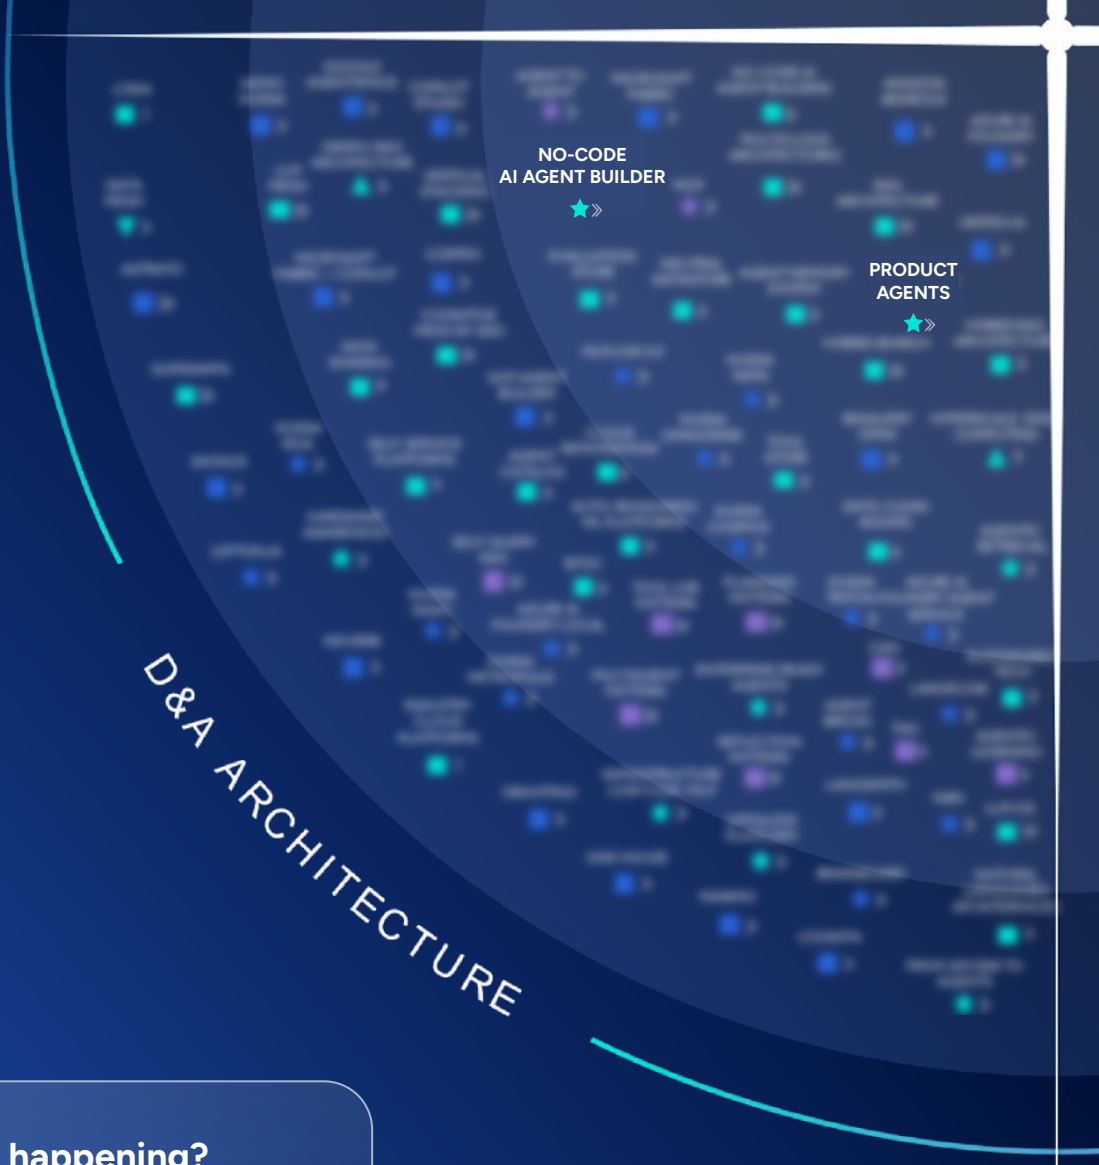

What's happening?

The market is aggressively democratizing AI through two tracks: **low-code agent platforms** and **pre-built product agents** embedded in data ecosystems.

Low-code builders allow users to **configure complex RAG pipelines in a few clicks**, while product agents offer **immediate, zero-configuration deployment** with natively inherited governance.

However, we recognize a dual reality: while these tools deliver instant value, they currently face limitations regarding architectural flexibility and the orchestration of highly complex, cross-platform enterprise logic.

Why it matters?

Speed and democratization are the primary drivers here.

These platforms allow organizations to deploy AI capabilities with **instant time-to-value**, **empowering non-technical users** to build complex workflows while reducing the burden on specialized development teams.

What's happening?

Just as GPT models revolutionized text, **Time Series Foundation Models (TSFMs) are transforming traditional forecasting**.

Pre-trained on billions of data points, they learn complex seasonality and trends to generate predictions on completely unseen data.

Zero-shot capabilities mean **development effort drops from weeks** of manual feature engineering **to mere seconds**, simplifying scalability and reducing deployment costs.

The model landscape is expanding: **Chronos** and **TimeGPT** lead in zero-shot performance, while **TiRex** uses an LSTM architecture to specialize in high-frequency data.

Why it matters?

The foundation model paradigm is arriving in traditional ML territory. The adoption of pre-trained time-series models unlocks zero-shot prediction capabilities.

This drastically **reduces development effort** and delivers a fundamentally **superior cost-to-accuracy tradeoff** for enterprise forecasting.

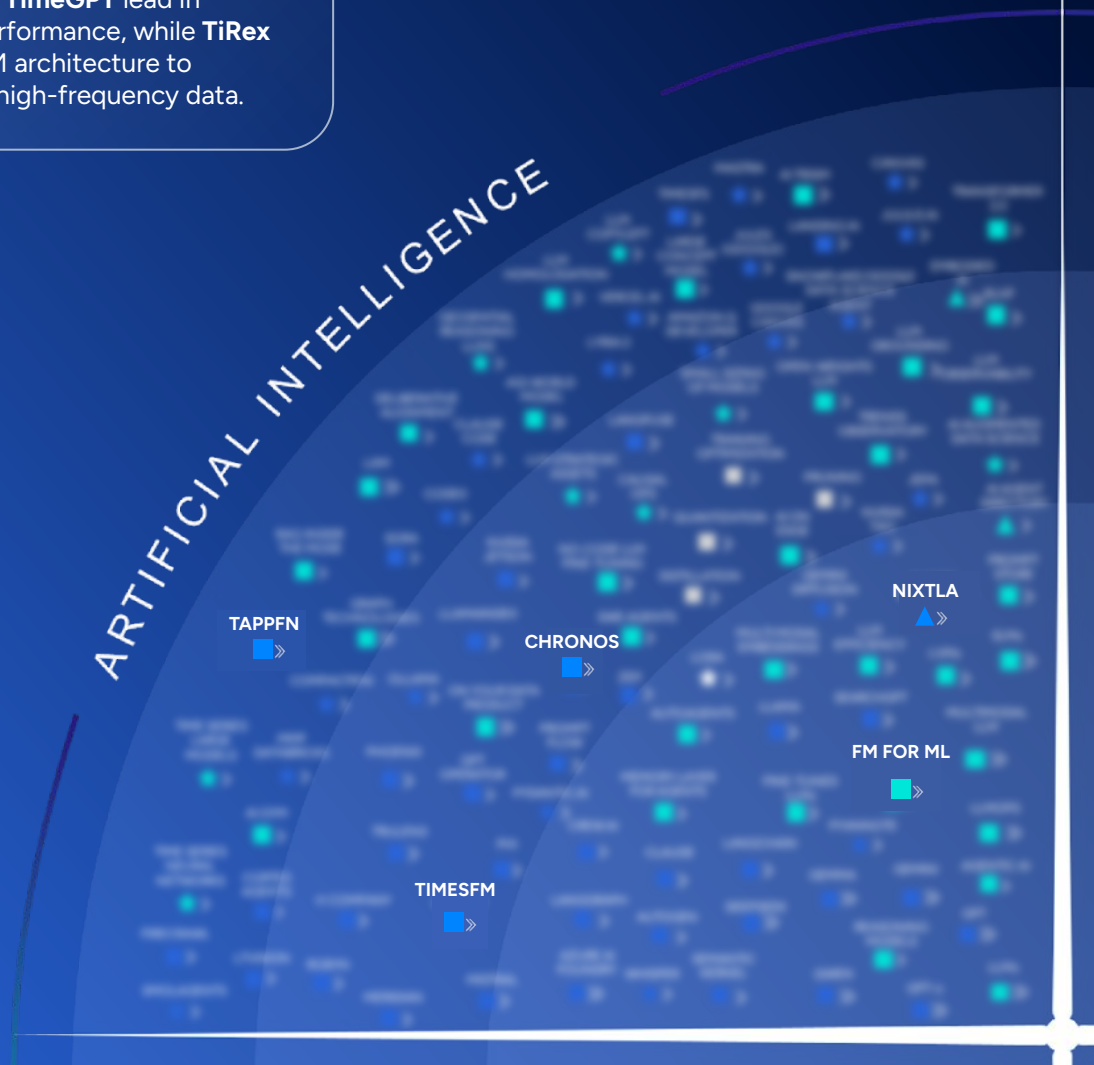

INNOVATION RADAR EPOCH #7

What's happening?

Marketing AI is evolving beyond speed and scale. **Content creation is now deeply connected to behavioral histories and propensity scores,** adapting message, tone, and timing with high precision.

For highly regulated sectors like pharma, mature solutions with clear supervision layers are finally accelerating.

Furthermore, as consumers increasingly use AI to inform purchase decisions, content must be **specifically designed for generative engines** to understand, interpret, and prioritize.

Strategies must now focus on ensuring brand presence in this new, AI-constructed discovery space.

Why it matters?

AI content generation is transitioning from pure volume to an **intelligence play** driven by behavioral data and customer signals.

Simultaneously, AI is becoming a **primary discovery channel** where purchasing decisions begin, making it crucial to optimize content for AI interpretation and prioritization.

sdg

A **Deep Dive** into the **hottest** *Futures*

INNOVATION RADAR EPOCH #7

INNOVATION RADAR EPOCH #7

What's happening?

The ability for AI to **autonomously navigate the web** and **operate computer interfaces** is transitioning into an architecturally viable reality.

Computer Using Agents (CUA) can execute tasks exactly as a human would, from navigating complex forms to extracting data from legacy dashboards.

Key drivers include **Playwright MCP**, which allows LLMs to interact via accessibility snapshots, and **Amazon Nova Act**, providing the fleet management needed for production.

Furthermore, the **WebMCP** standard proposes a future where websites expose tools directly to AI agents.

Why it matters?

CUA technology bridges the gap between disconnected legacy systems and modern automation.

By allowing agents to operate UIs like humans, we can **automate** massive, manual **workflows** across web interfaces and dashboards **without** the need for **traditional API** integrations.

INNOVATION RADAR EPOCH #7

What's happening?

We observe a recurring pattern: **multi-agent systems** work locally but face unexpected API bills and zero visibility in production.

To operate reliably, they require unattractive infrastructure like state persistence, per-agent token accounting, circuit breakers, and deadlock detection.

The industry is solving this by **building dedicated execution environments** rather than just improving orchestration frameworks.

These new environments provide **granular control** over invoked binaries and network endpoints, ensuring **security through active enforcement** before execution rather than simple logging.

Why it matters?

As multi-agent systems move into production, the real bottleneck isn't the AI models, but the **lack of operational infrastructure**.

Building this "unglamorous plumbing" is essential to **make autonomous systems reliable**, **observable**, and **economically viable** at an enterprise scale.

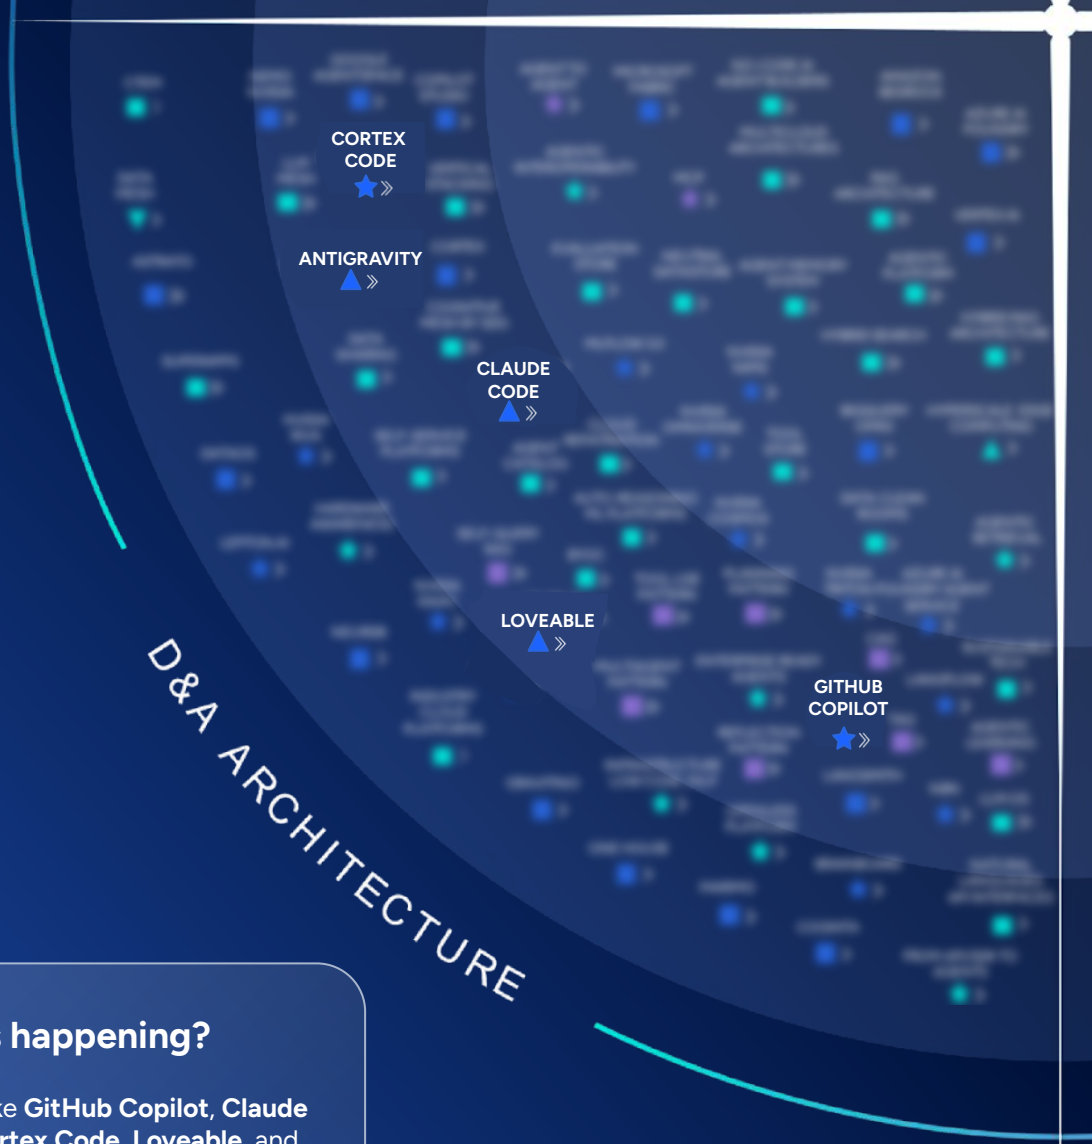

INNOVATION RADAR EPOCH #7

What's happening?

Agents like **GitHub Copilot**, **Claude Code**, **Cortex Code**, **Loveable**, and **Antigravity** now autonomously **read codebases**, **run tests**, and **commit code**.

The market battle focuses on integrated experiences, raising the risk of workflow lock-in. To mitigate this, we champion **portable layers**.

Plain-text files like CLAUDE.md and AGENTS.md define **transferable coding conventions** in natural language.

Additionally, the **Model Context Protocol (MCP)** provides a partially portable layer for agnostic tools.

While MCP interoperability requires strict platform adherence, it ensures **custom workflows remain adaptable**.

Why it matters?

AI developer tools are moving far beyond simple autocomplete into **full-lifecycle autonomy**.

The strategic priority for engineering teams is no longer deciding whether to adopt these capabilities, but rather how to deploy them while successfully **avoiding severe vendor lock-in** from day one.

Why it matters?

In an era of strict privacy regulations, **synthetic data** is the default solution for **secure innovation**.

It allows us to **eliminate GDPR risks** in non-production environments and **overcome data scarcity**, **accelerating AI training** without exposing sensitive customer information.

What's happening?

Copying production data into testing environments is becoming a risky anti-pattern.

Synthetic data generation is now the regulator-backed solution for eliminating privacy exposure.

To be valuable, these **algorithmically generated datasets** must maintain **strict referential integrity** and deliver **high semantic fidelity**.

While single-table scenarios are common, we are pioneering complex multi-table solutions.

This ensures that the data tells a credible business story and respects physical structures, allowing organizations to **innovate safely** and **at scale**.

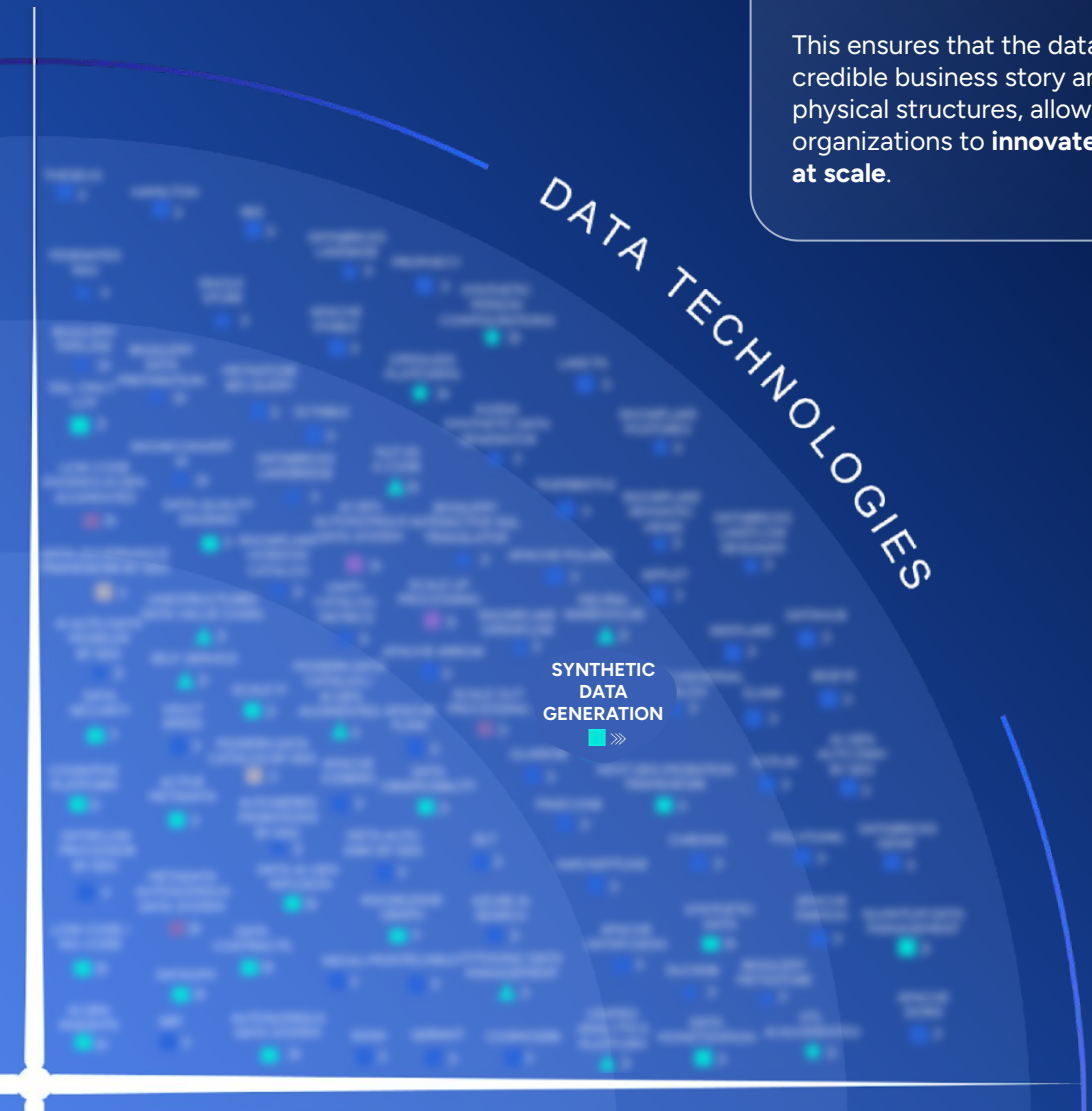

Why it matters?

The **semantic layer** is the critical bridge connecting structured tabular data with LLMs.

It solves the naive approach of letting AI directly query databases, ensuring that business teams get **precise**, **accurate reports autonomously**.

What's happening?

Historically, LLMs have struggled to translate business concepts into precise database queries without a structured semantic layer.

We are now seeing a significant evolution where platform vendors **bundle** the LLM, agent, semantic layer, and application into a **unified stack**.

A major breakthrough is **Snowflake's Semantic Autopilot**, which automatically generates an initial semantic layer by analyzing existing query histories and dashboards.

This capability drastically **reduces the effort** required to bootstrap AI Ready Data, **accelerating enterprise adoption** and **enabling true data self-service**.

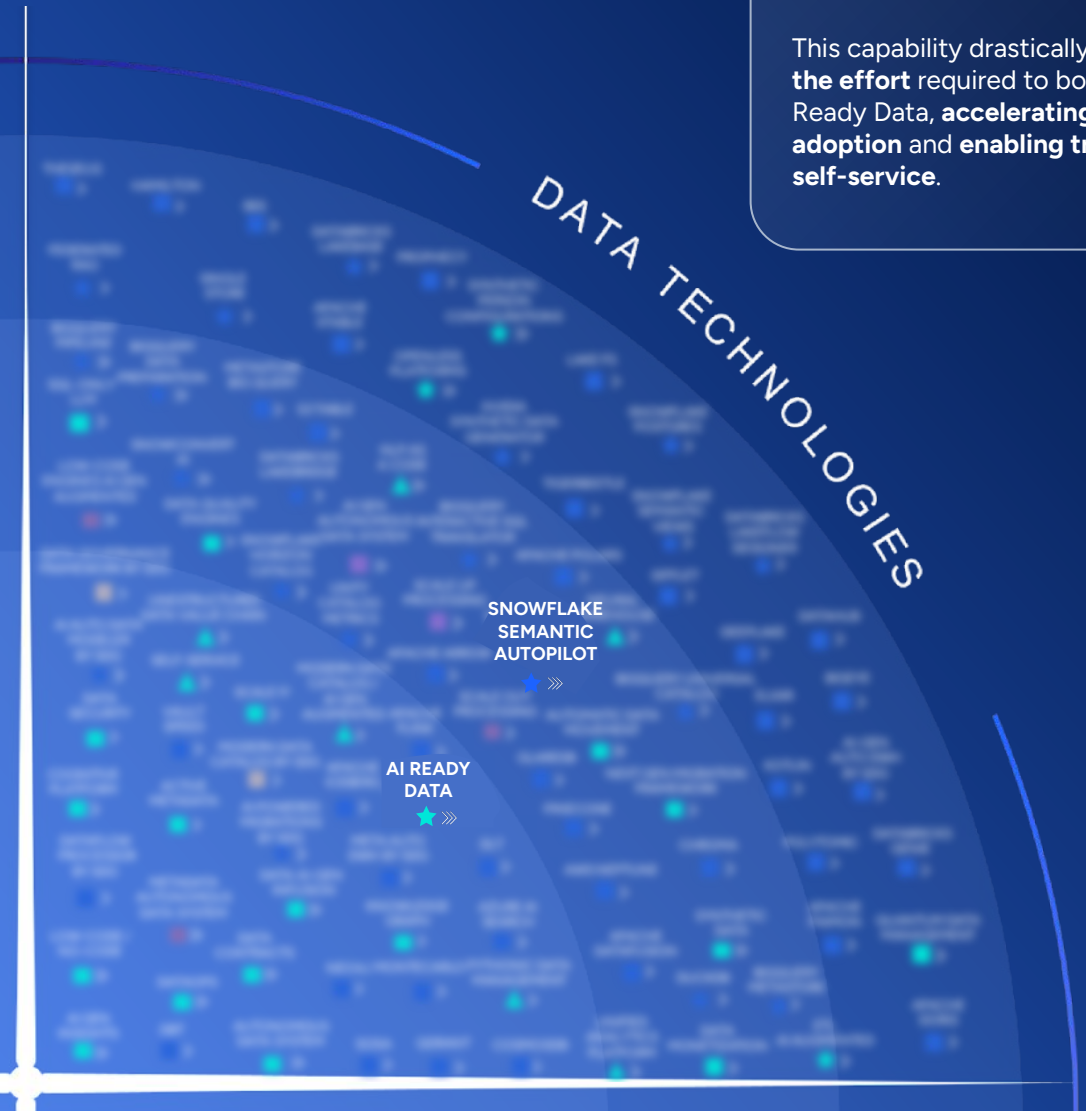

What's happening?

Generative audio is moving toward highly specialized architectures, unlocking improvements in latency and real-time support.

The **NVIDIA NeMo ecosystem** is capable of deploying specific capabilities: **PersonaPlex** enables full-duplex, real-time human-AI interaction; **Parakeet** ensures exceptional transcription precision; and **Canary** performs autonomous multilingual translation in a single pass.

Additionally, multimodal innovations like **Meta's SAM Audio** are expanding the boundaries of audio isolation, driving true ROI for enterprise speech analytics.

Why it matters?

Real-time, natural audio interaction unlocks a **new dimension** of **customer experience** and **operational efficiency**.

Specialized architectures are making **human-AI conversations** viable for enterprise use by drastically **reducing latency** and operational **costs**.

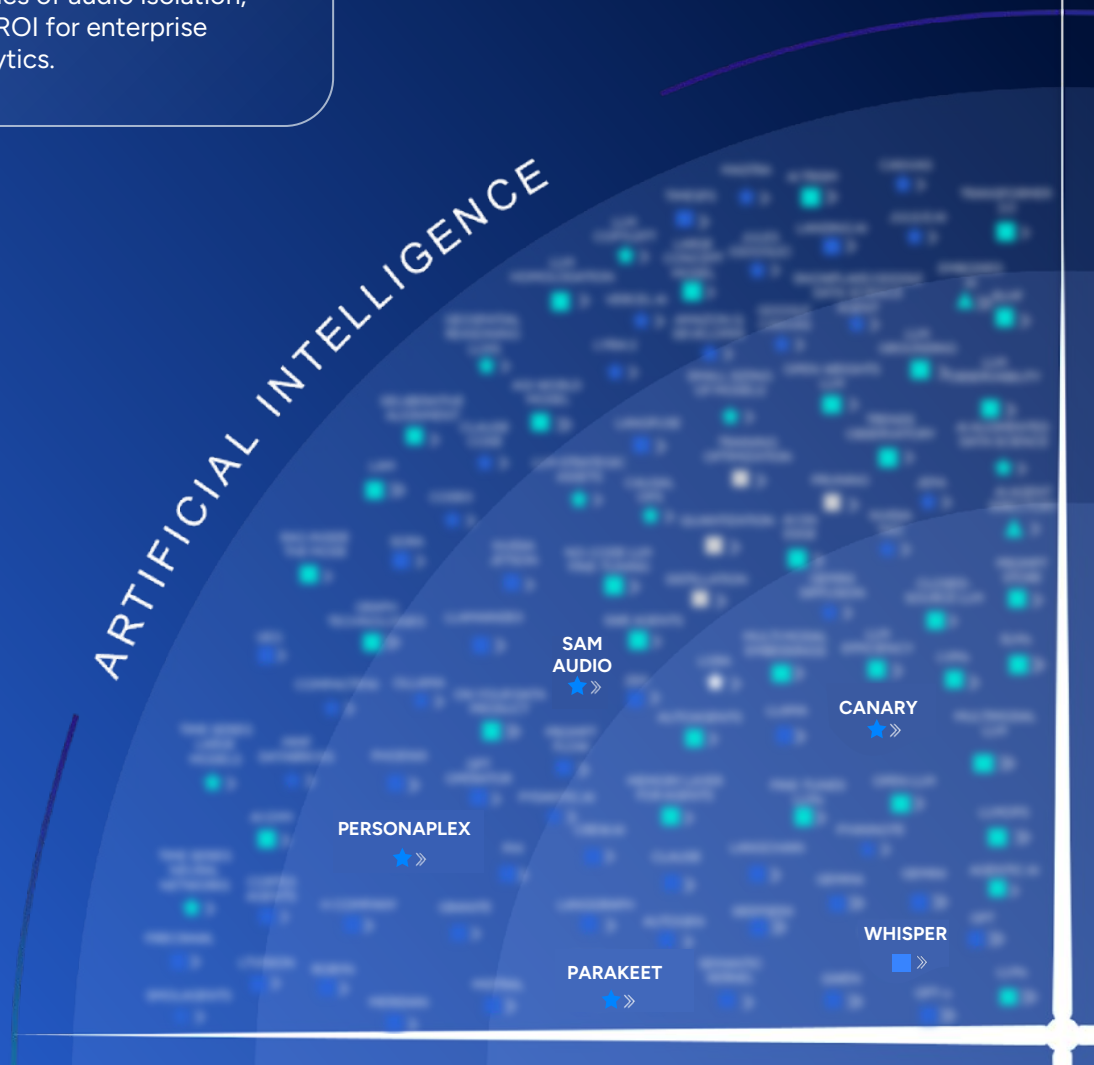

What's happening?

The industry is shifting from multi-agent systems to a "**Skill-Driven Development**" paradigm.

Skills and Agentic Prompt Blocks act as structured **layers between agents and tools**. Each skill contains discovery metadata, detailed Markdown instructions for specialization, and scripts that execute in controlled environments.

We leverage this approach to **load extended context** only when strictly necessary, **preventing prompt saturation**.

This provides agents with targeted domain knowledge, combining autonomous reasoning with explicit teaching to improve both **generation quality** and **cost efficiency**.

Why it matters?

As agentic AI architectures mature, teams face challenges like degraded model precision from excessively long prompts and escalating costs from unbounded reasoning loops.

Skills provide a necessary abstraction layer for modular, on-demand specialization.

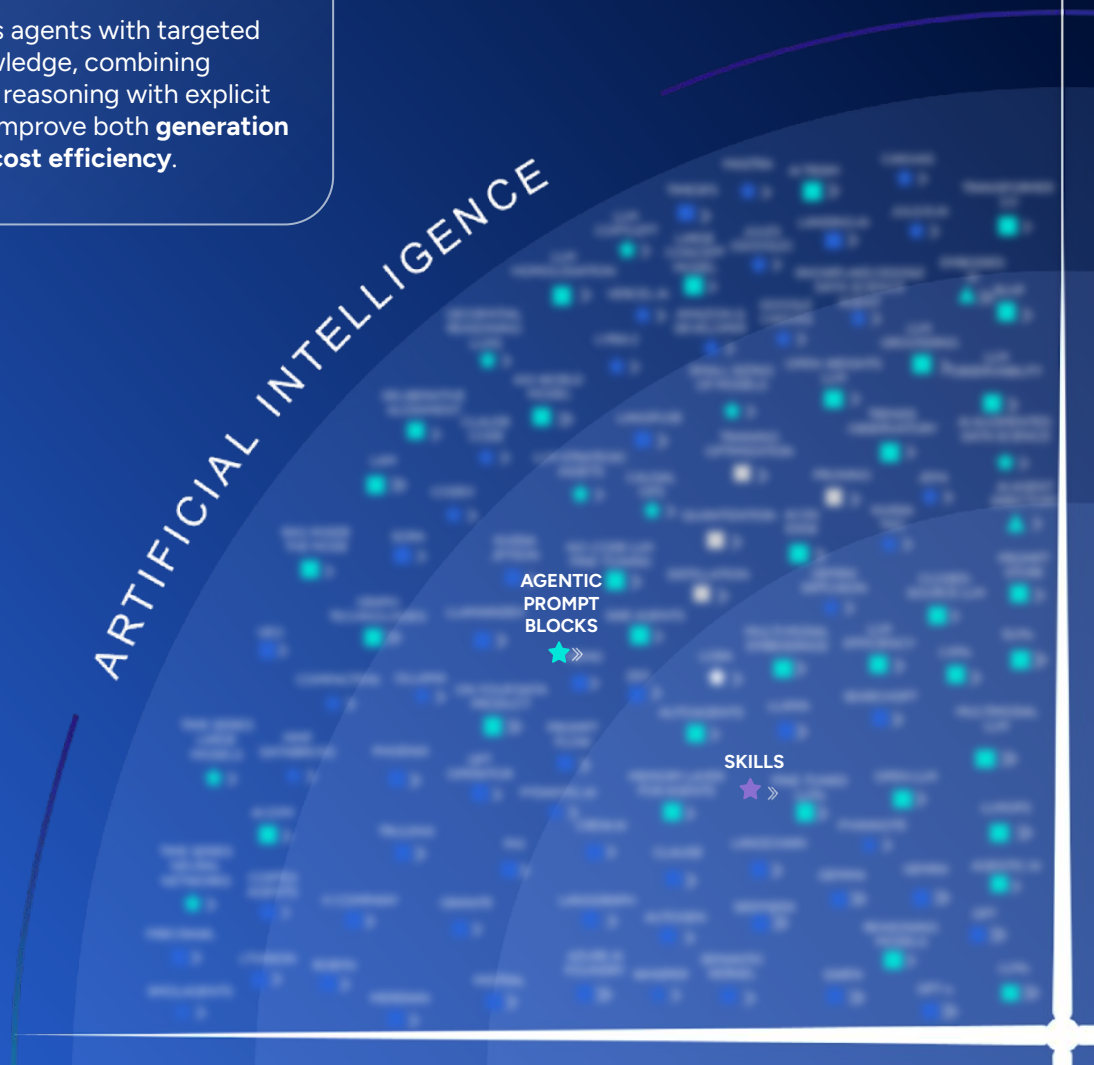

INNOVATION RADAR EPOCH #7

What's happening?

The convergence of local infrastructure ownership and agentic AI is creating **long-running agents that interact deeply** with a user's entire system over hours or days.

Originating from open-source experiments like **OpenClaw**, these local agents can **manage calendars**, **draft emails**, and **summarize files** via messaging channels.

We are seeing major players rapidly shape these into enterprise-ready solutions, introducing tools like **Claude Cowork** and the **AI Browser**.

As organizations adopt them, new challenges are rising, such as **infrastructure costs**, **context degradation**, and secure **error recovery**.

Why it matters?

We are shifting from reactive chatbots to **proactive**, **autonomous assistants** that receive high-level objectives rather than step-by-step instructions.

This paradigm shift significantly **boosts individual productivity** by allowing agents to execute complex tasks in the background.

INNOVATION RADAR EPOCH #7

What's happening?

The automation of individual tasks is rapidly scaling into a **fully autonomous Agentic Economy**.

In this paradigm, customer-side and merchant-side agents negotiate and execute transactions **without human intervention**.

We are monitoring the essential foundations making this possible, including the **Universal Commerce Protocol (UCP)** for interoperability and the **Agent Payments Protocol (AP2)** to ensure financial security.

This shift requires **new governance frameworks** as machines begin to autonomously manage budgets and fulfill commercial orders in a frictionless market.

Why it matters?

The **Agentic Economy** transforms AI from a research tool into an economic actor.

By enabling agents to execute transactions autonomously, we remove friction from commerce and create a world where **machines negotiate** and trade on behalf of humans.

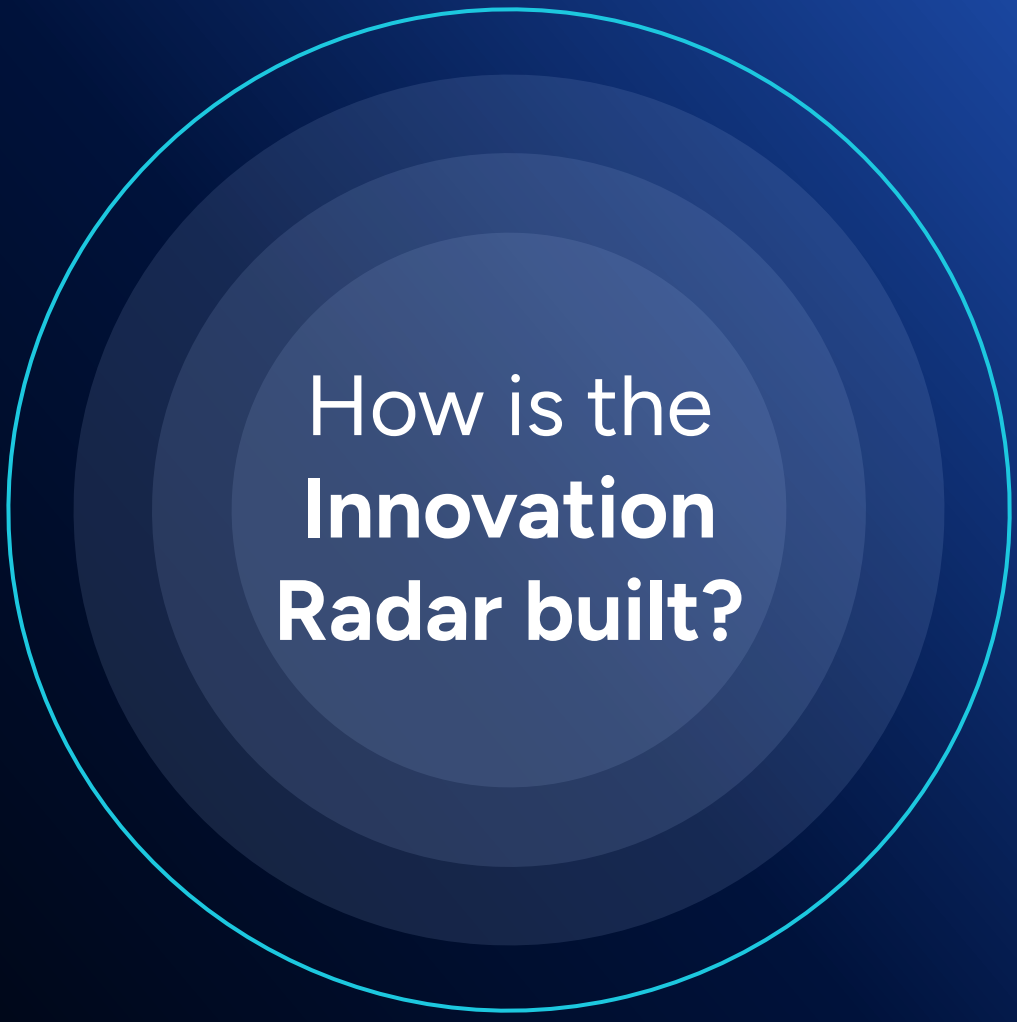

How is the **Innovation Radar built?**

INNOVATION RADAR EPOCH #7

What Elements Will You Find in the Radar?

The **Innovation Radar** is a publication from SDG Group that focuses on **emerging trends** and **technologies**. It aims to conceptualize and recommend **trends that are gaining traction** and providing **insights** into recent developments in the technology and business landscape.

***Futures**: The types of elements that you will be able to find inside the radar are:*

- **TREND**: Trends signal evolutions or emerging changes in practices, technologies, or theories that affect data architectures from the point of view of infrastructure, data engineering, advanced analytics, and business needs. These trends may arise from technological advances, changes in data regulation, or new business strategies that demand different ways of handling information.
- **TOOL**: Software applications, platforms, or utilities that facilitate the performance of specific tasks within the process of building data & AI architectures.
- **PATTERN**: Design patterns are reusable solutions for common problems in the design and implementation of data & analytics systems. These patterns provide a framework for creating robust and efficient data architectures.
- **FRAMEWORKS**: A structured set of concepts, practices, guidelines, and tools that provide a systematic way to build and manage the architecture for a specific case.

Quadrants: The sections of the Radar where trends are grouped under a common theme:

- **Data & AI Architecture**
- **Data Technologies**
- **Artificial Intelligence**
- **Business / Other Trends**

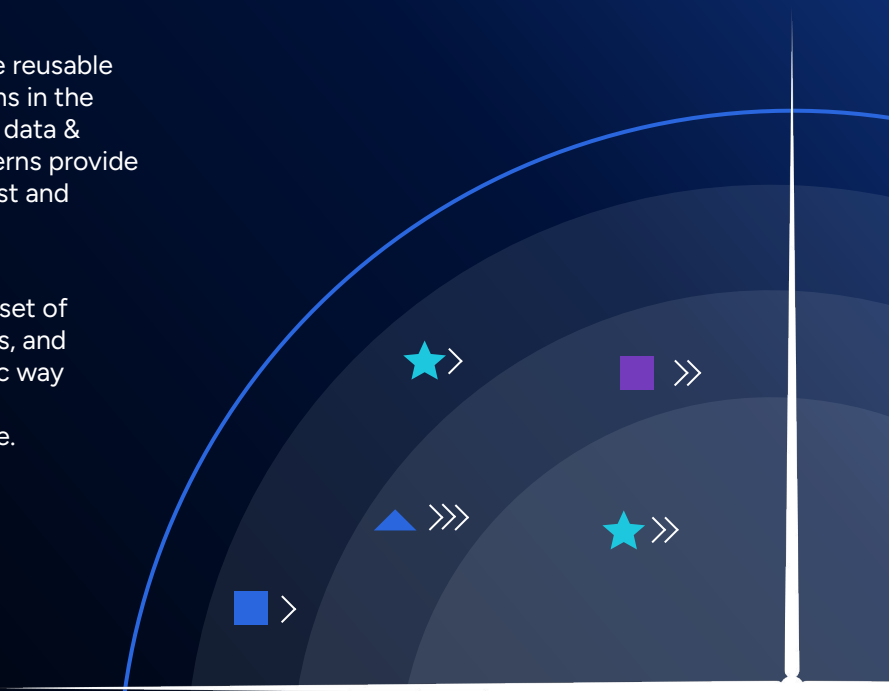

Impact: This dimension of analysis will refer to the impact that the new trend will have on data & AI architectures in the future:

- **HIGH**: The element will have a very positive impact on the architectures that integrate it due to the innovative capabilities it brings, leading to an evolution in most use cases.
- **MEDIUM**: The element will have an impact in certain areas of the architecture. Although its influence is considerable, it manifests itself in a more limited spectrum of situations.
- **LOW**: The element will make new capabilities available to users of the architecture, but will not represent a significant change in the architecture or will only be applicable in very limited use cases.

Relevance: This dimension refers to the priority of application based on differential capabilities that the element will provide, and the subsequent competitive advantage. This dimension will not be represented in the form of color or an associated symbol, but with the position of the trend within the circles that form the radar. **The closer to the center, the more transcendence it will have in the short term.**

Evolution: This will be an indicator of how the trend is evolving within the different radar volumes. The possible evolutions are:

- **NEW** : An element that has appeared on the radar for the first time.
- **INCREASE** : An element that has increased in its implementation priority.
- **STAND** : An element that has not undergone any change.
- **DECREASE** : An element that has lowered in its prioritization level.

How are Futures evaluated?

EVALUATION OF EXISTING AND NEW FUTURES FOR EVERY EPOCH

Each Epoch every future is evaluated on the different dimensions (relevance, impact) although futures have meaning on its own and description of its own, they are usually grouped into families (click on a trend in the radar in our webpage to see all the futures of the same family).

For every family & future:

- **Impact:** Has the impact of the future changed since the last epoch? If so, change its Impact in the radar database.
- **Relevance:** Has the future matured enough or increase/decrease in its short-term relevance?.
- **Change Description**: If the future has evolved the description will be updated based on the latest news
- **Disappearance:** If a future of a family of Futures has been decreasing in its relevance or Impact for several epochs it will disappear on the next epoch.

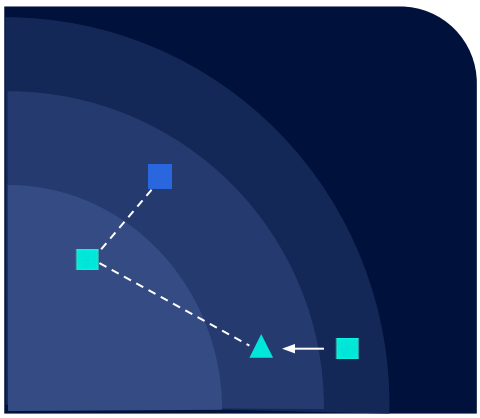

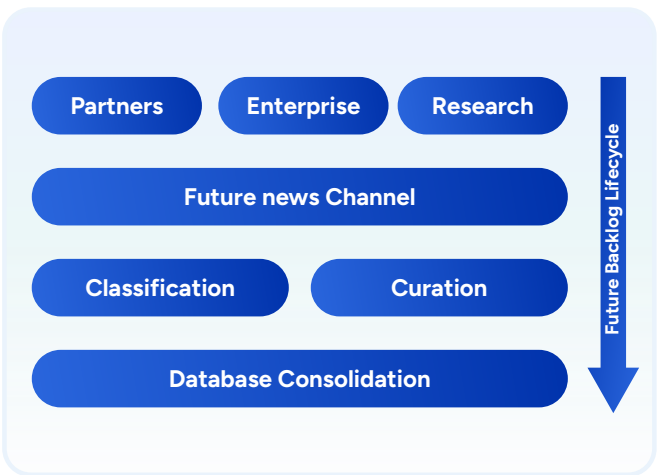

EVALUATION OF EXISTING AND NEW *FUTURES* FOR EVERY EPOCH

Future selection is a multi step process where news from different sources are gathered from multiple origins (SDG Partners (AWS, Azure, GCP; NVIDIA; Snowflake, ...), Associated Universities and Research Hubs (Huggingface ...) as well as on going projects and conversation with our wide client pool across the different business sectors are gathered into a single information channel.

Each of those news is classified, grouped and labeled to define if those news are relevant to an existing future or if they constitute a new future entirely. When that step has finalized, we add those *futures* candidates to existing families or create a new family. After that process is done the candidates go into a final curation process before being added to the database.

sdg

Thank you!

To make it easier to understand the radar, we hold a live session on our YouTube Channel where we talk about the trends that shape our radar to provide more information and help with understanding. Additionally, this live session is available on-demand at any time on the same channel.

Our Innovation department keeps its doors open to organize informative and enrichment sessions about our corporate data management and treatment strategy, aimed at other interested companies.

To participate in these sessions, you can contact us at **innovation@sdggroup.com**

orbitae