Fontys > FOR SOCIETY

GENAI ENGINEERING

Software Engineering in the AI Era: from ML to LLMs and Agentic AI

Petra Heck – Associate Professor Fontys ICT - p.heck@fontys.nl

Program

14.00 – 14.30	**What is GenAI Engineering?** Petra Heck (Fontys)
14.30 – 15.00	**GenAI Engineering in Practice** Bart Scheffer (Moonly)
15.00 – 15.30	**How AI is Reshaping Software Engineering** Pascal Widdershoven (Kabisa)
15.30 – 17.00	Drinks and time for discussion

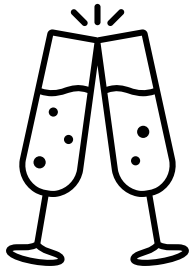

Fontys

About Me

- **Software engineer** (MSc Computer Science) since 2002
- **Software quality** consultant since 2004
- Lecturer **Fontys ICT** since 2012
  - 2016 PhD “Quality of JIT Requirements” (Computer Science, AI4RE)
  - 2019 Senior researcher AI Engineering (SE4AI)
  - **2023 DEMAND: Data Engineering & Data Management (SE4DATA)**
- Associate Professor AI & Software Engineering (SE4AI & AI4SE) since 2025

=> Practice-oriented researcher with a Software Engineering view on AI & Data

Fontys

2019 till 2026: From AI to GenAI

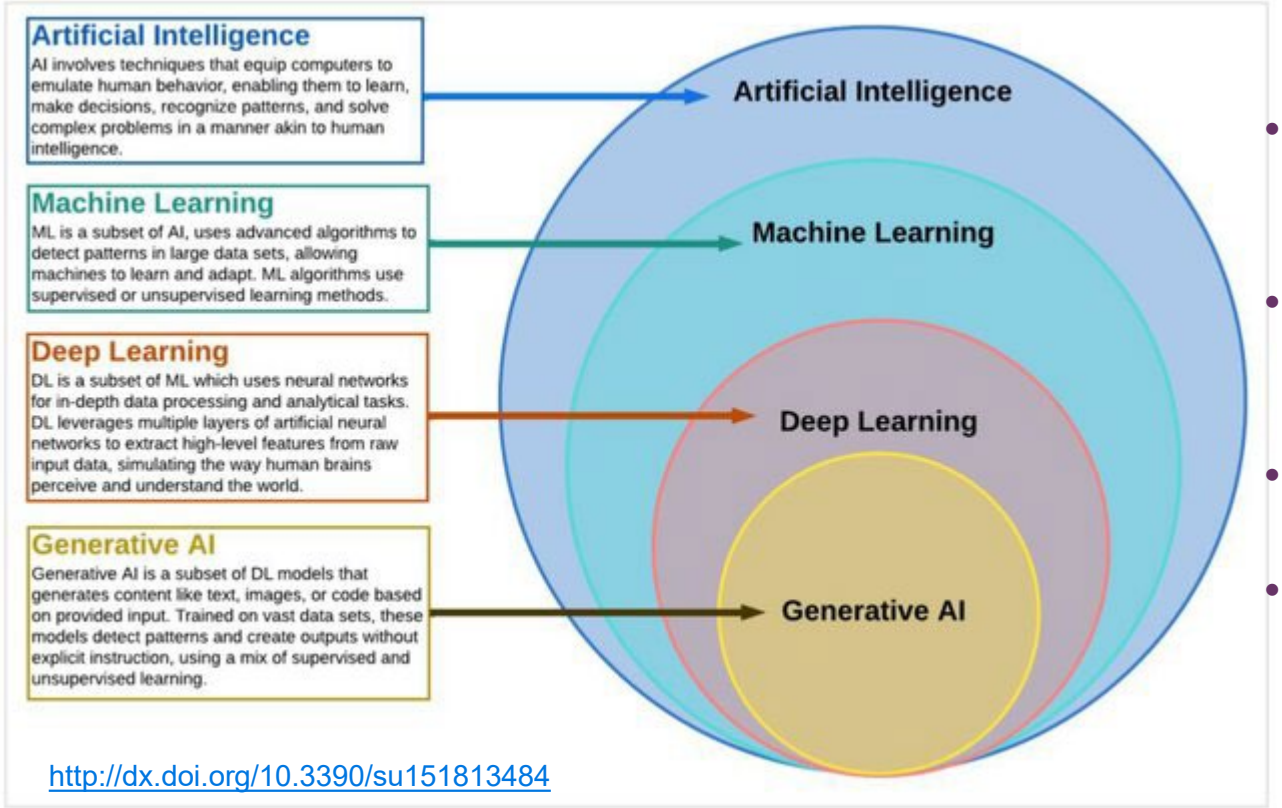

http://dx.doi.org/10.3390/su151813484

- Emulate human behavior
- Pattern detection algorithms
- Neural networks
- Generate content

Fontys

From AI Engineering to AI-augmented Engineering

“How to build production-ready AI systems?” (Petra Heck, 2019)

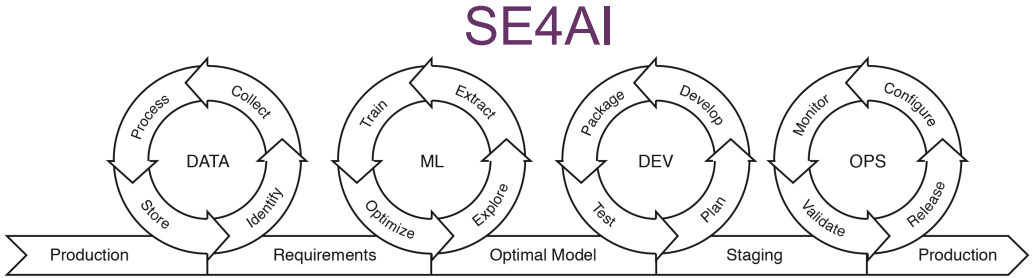

“How to use AI to build production-ready systems?” (Petra Heck, 2024)

Fontys

GenAI Engineering = LLM Engineering =

SE4LLM: How to build production-ready LLM-enabled systems?

+

LLM4SE: How to use LLM-enabled tools to build production-ready systems?

Fontys

Fontys > FOR SOCIETY

SE4LLM: BUILDING LLM-ENABLED SYSTEMS

GenAI Landscape

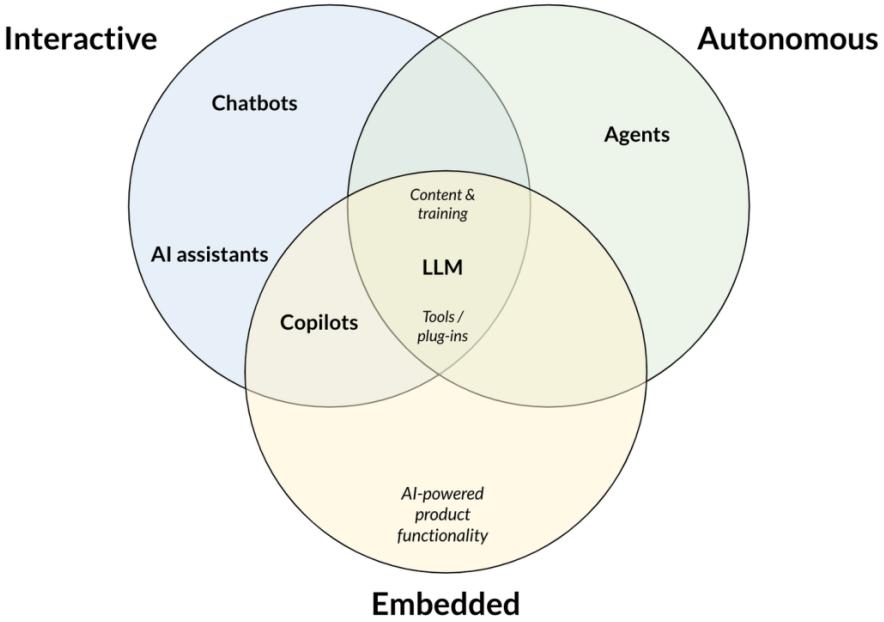

A visual guide to today's GenAI landscape - MERL Tech

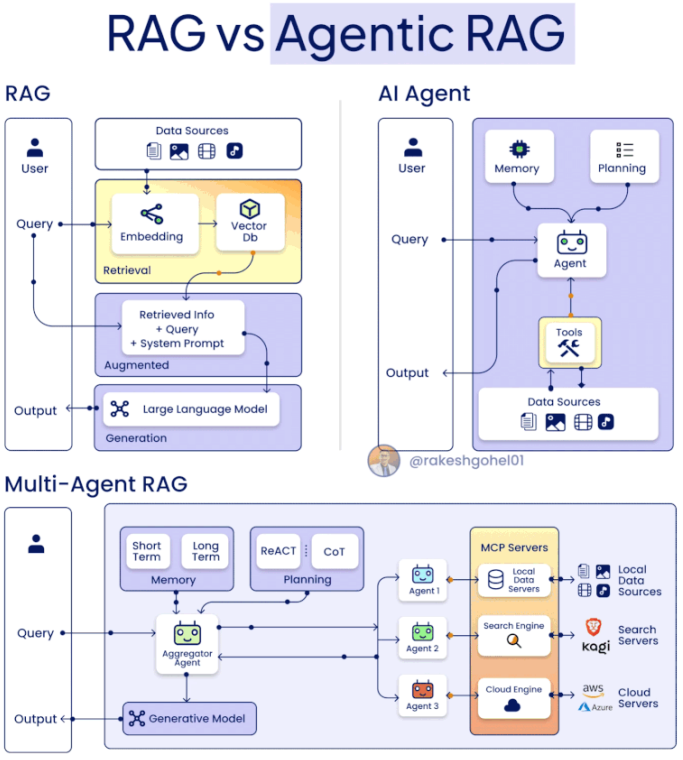

(9) Post | LinkedIn

From DevOps to MLOps to LLMOps

LLMops adapts the MLops tech stack for generative AI use cases

INSIGHT PARTNERS

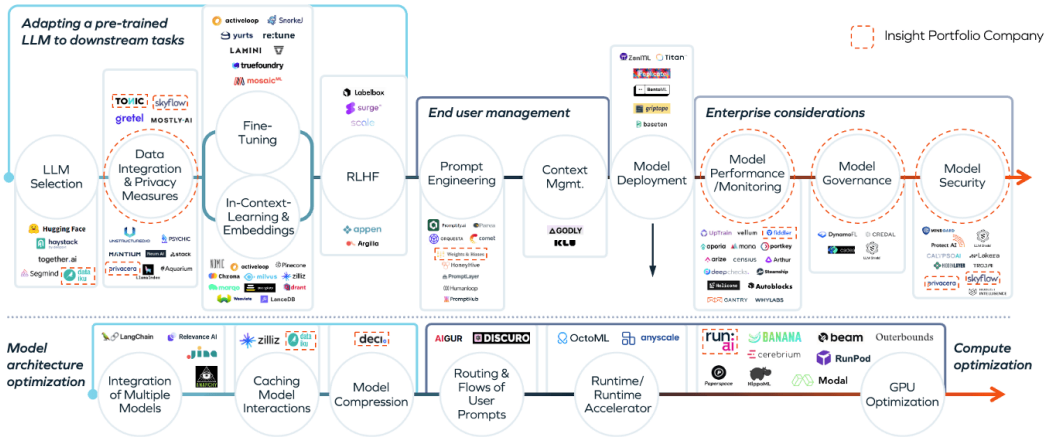

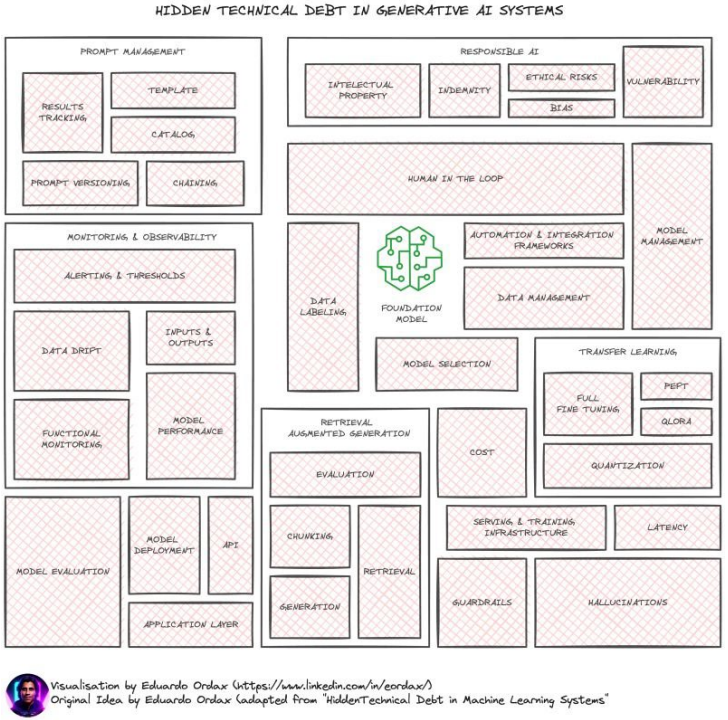

[Navigating the LLMops landscape: What you need to know | Insight Partners](https://fontysblogt.nl/llmops-engineering-trustworthy-llm-systems/)

https://fontysblogt.nl/llmops-engineering-trustworthy-llm-systems/

Fontys

Agentic AI

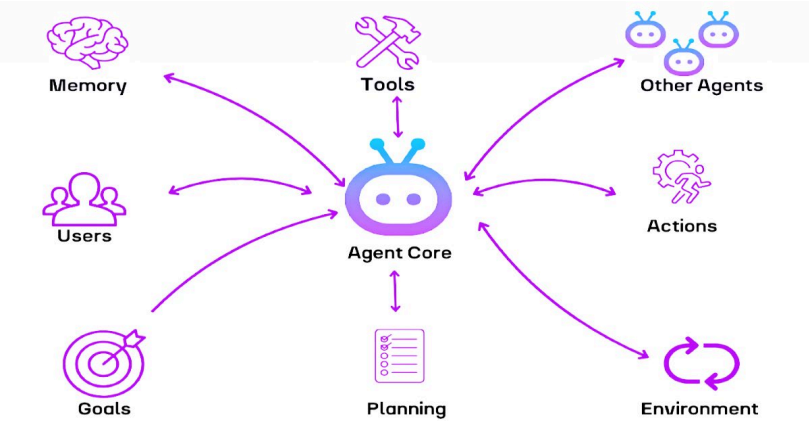

What is Agentic AI?

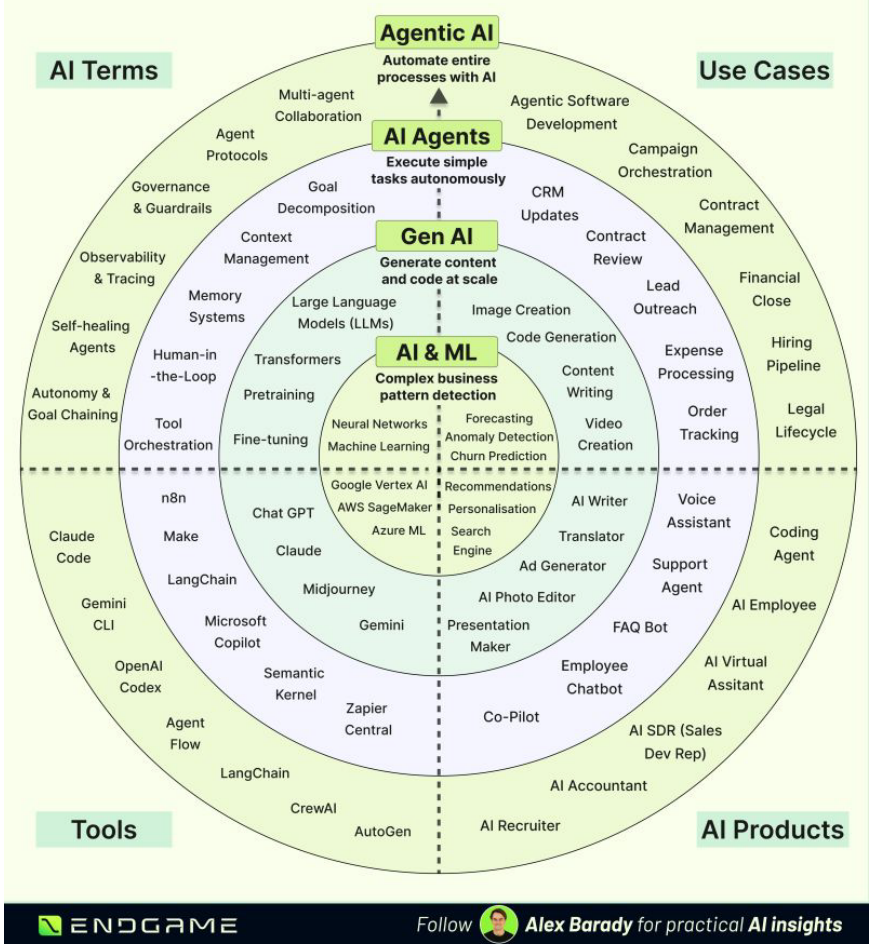

(16) Post | LinkedIn

Agentic AI – Terminology

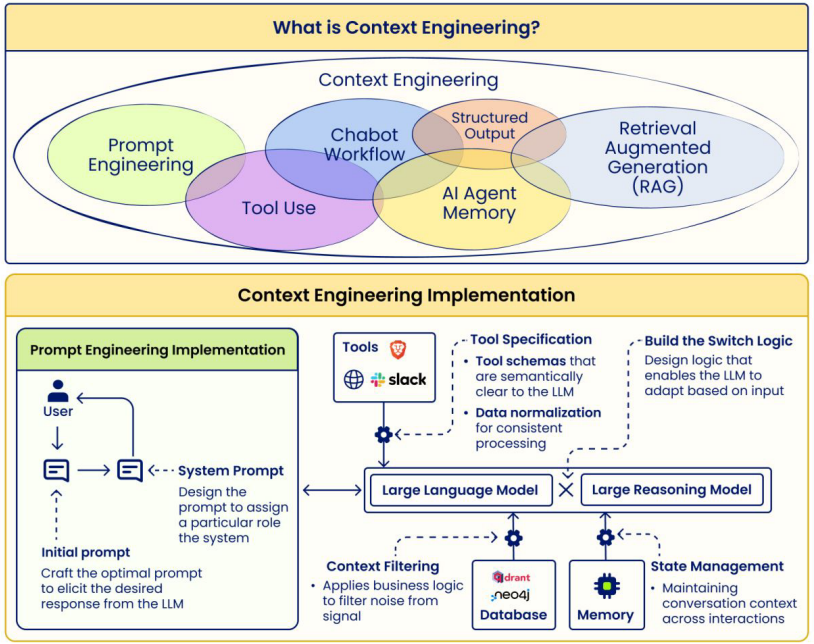

[(9) Post | LinkedIn]

Fontys

GenAI Tech Stack

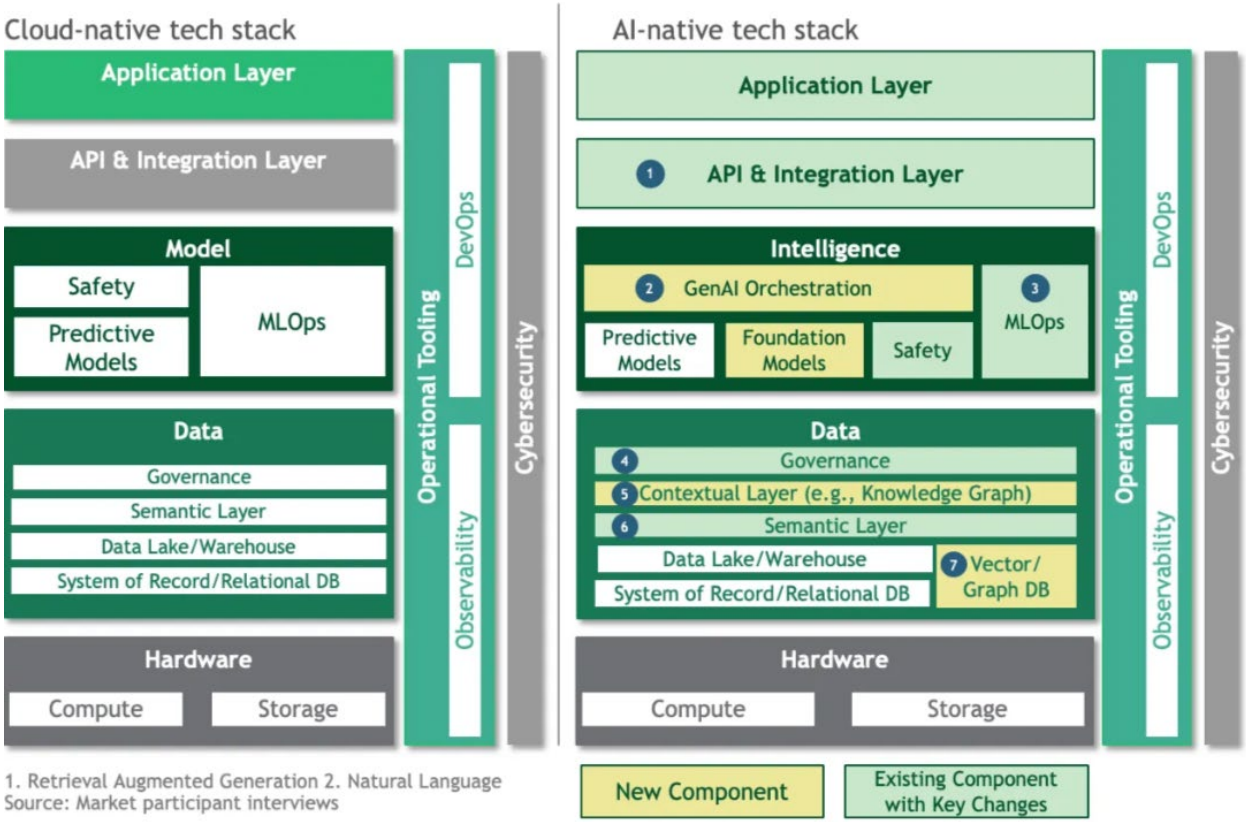

1. Retrieval Augmented Generation 2. Natural Language
Source: Market participant interviews

From Apps to Agents: How the AI-Native Tech Stack Is Transforming Software | BCGonTech

Fontys

Patterns for Building LLM-based Agents

Gartner

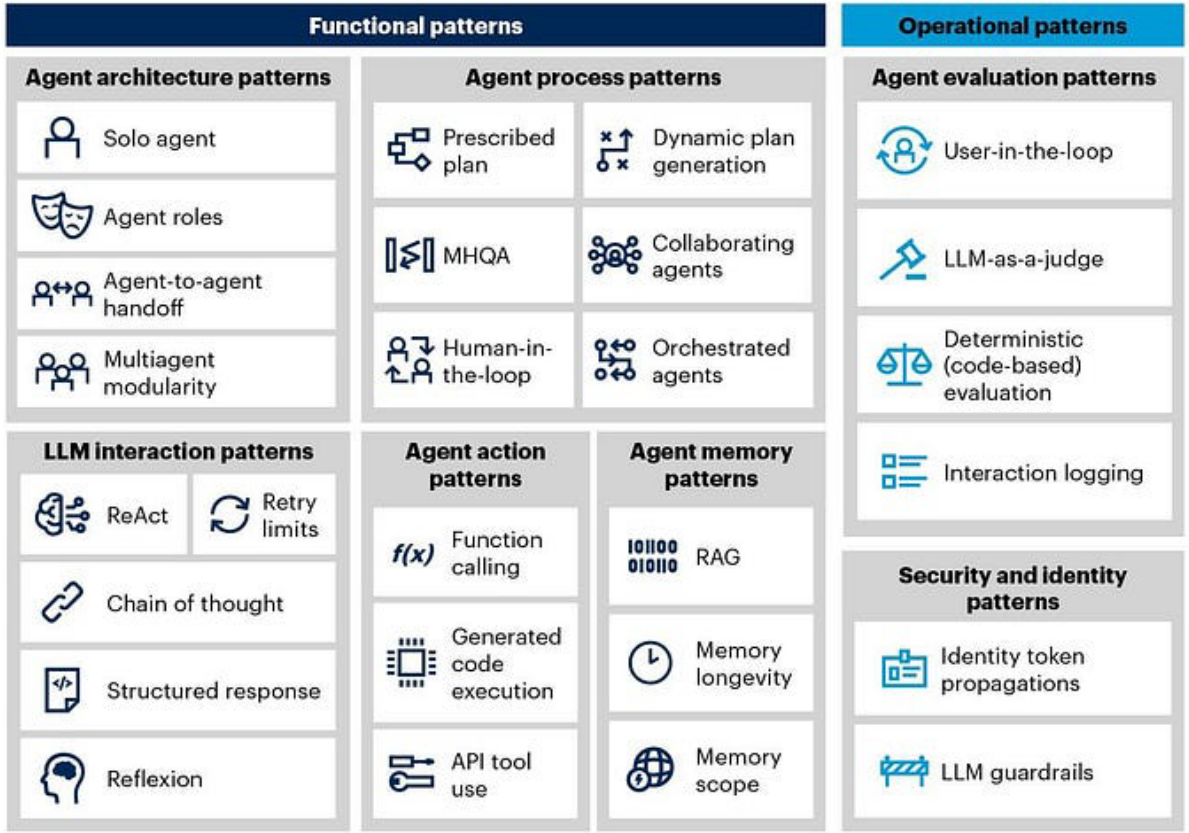

Emerging Patterns for Building LLM-Based AI Agents

Fontys

From LLMOps to AgentOps

AgentOps Patterns by Category

Category	Pattern	Description
Artefact Management	Agent Artefact Registry	Versioned store for behavioural, eval, and metadata artefacts.
Execution	Context Engine	Assembles just-enough context within budget.
	Runtime Workflow Simulator	Shadow/sandbox execution to vet plans before/during exe.
Safety & Enforcement	Guardrail Enforcement Points	Policy checks at all critical agent action points.
Monitoring & Observability	Trace Logger	Structured spans for replay and audit.
	Cumulative Risk Ledger	Append-only per-step risk signals; shows accumulation.
Evaluation-Driven Learning	Dynamic Test Case Factory	Harvests operational insights into versioned tests.
	Marginal Risk Assessor	Compares candidate vs. baseline risk.
	Agent-as-a-Judge	Evaluator agent scoring runs against contracts.
Interaction	Interactive ChatOps	Conversational queries over ops data; safe controls.
	Trust Calibration Controller	Calibrates user cues to measured reliability.
Assurance & Compliance	Assurance Evidence Factory	Sealed bundles linking artefacts to bounded claims.

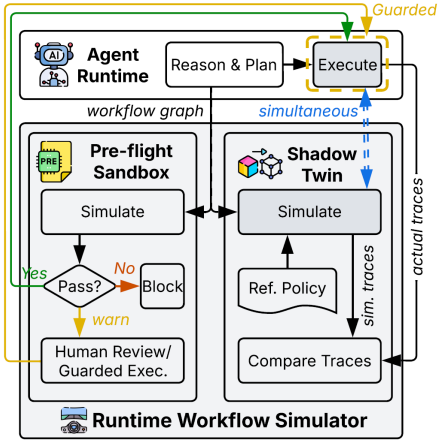

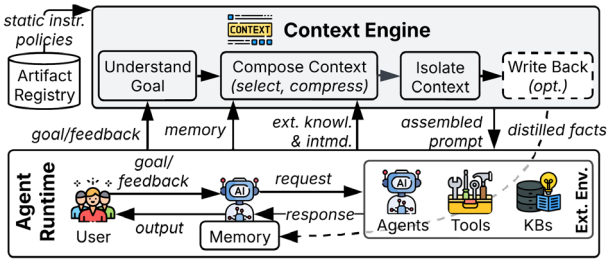

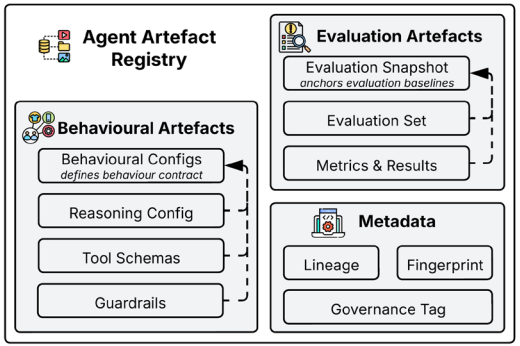

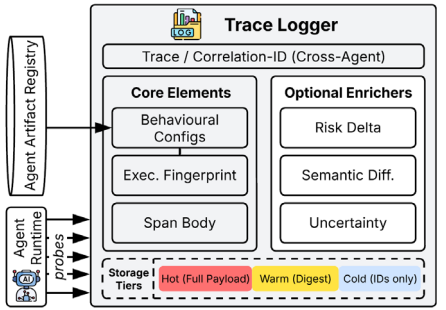

Fontys

Xia, Boming and Xia, Boming and Liu, Yue and Lu, Qinghua and Zhu, Liming and Sejdinovic, Dino and Sejdinovic, Dino,
AgentOps Pattern Catalogue: Architectural Patterns for Safe and Observable Operations of Foundation Model-Based Agents (September 25, 2025).

Fontys > FOR SOCIETY

LLM4SE: AI-AUGMENTED DATA ENGINEERING

From MLOps to DataOps

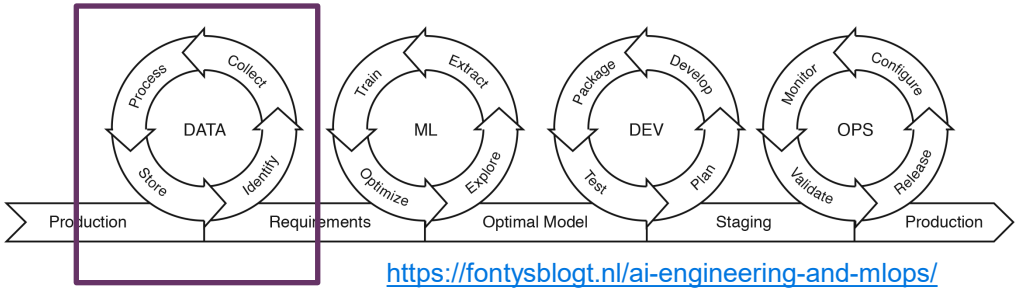

https://fontysblogt.nl/ai-engineering-and-mlops/

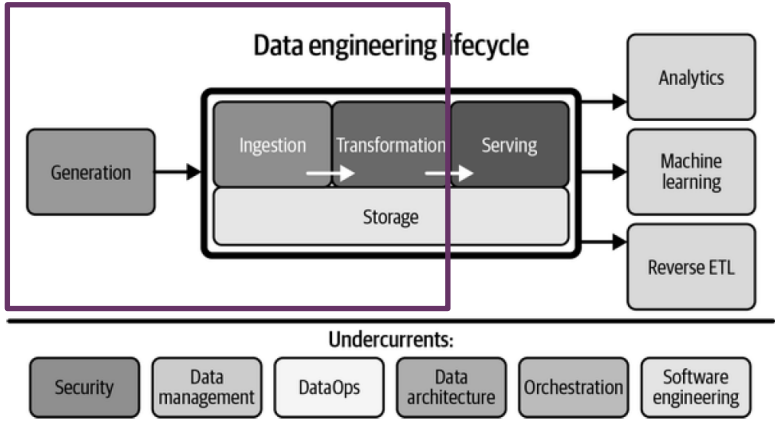

Fundamentals of Data Engineering by Joe Reis, Matt Housley.

"How to build production-ready AI systems?" (Petra Heck, 2019)

"How to build (and maintain!) production-ready data-driven (AI) systems?" (Petra Heck, 2024)

Fontys

DataOps

"an approach that accelerates the delivery of high-quality results by *automation and orchestration* of data life cycle stages."

DataOps adopts the best practices, processes, tools and technologies from Agile software engineering and DevOps thereby promoting the culture of *collaboration and continuous improvement*.

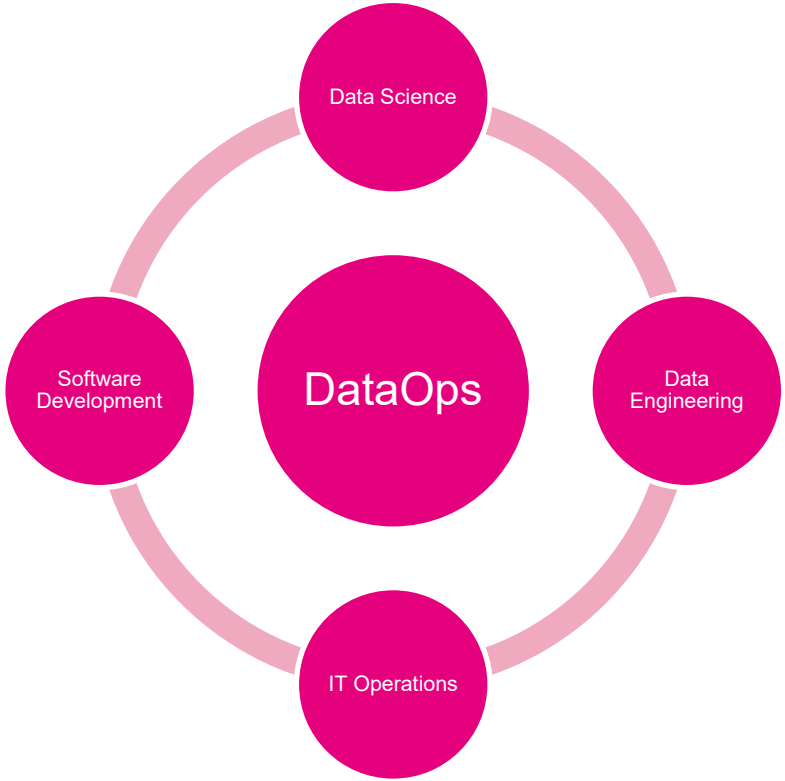

Fontys

Aiswarya Raj Munappy, David Issa Mattos, Jan Bosch, Helena Holmström Olsson, and Anas Dakkak. 2020.
From ad-hoc data analytics to dataops.
In Proceedings of the International Conference on Software and System Processes. 165–174.

Framework for Unified DataOps Strategy

http://dx.doi.org/10.1016/j.jjimei.2025.100321

Comprehensive DataOps Strategy

Structured Approaches

- End-to-end DataOps lifecycle integration.
- Workflow orchestration to streamline processes.
- Team collaboration.
- Automate routine tasks & reduce manual effort.
- Integrate data governance throughout.

Policy and Governance

- Ensure data privacy & security with strong governance policies.
- Define clear data ownership and sovereignty responsibilities.
- Establish ethical data management guidelines.
- Incorporate continuous monitoring and testing.
- Use cross-organizational governance frameworks.

Innovations and Tools

- Choose appropriate tools for each DataOps stage.
- Leverage scalable cloud-based solutions.
- Adopt real-time data processing tools (e.g., Kafka, Spark).
- Use machine learning for data quality management (e.g., DQSOps).
- Automate versioning & deployment using CI/CD pipelines.

Generative AI for DataOps

- Synthetic Data Generation
- Anomaly Detection
- Code Automation
- Data Quality Assurance
- Documentation Automation
- Intelligent Monitoring

Fontys

GenAI for Dashboards

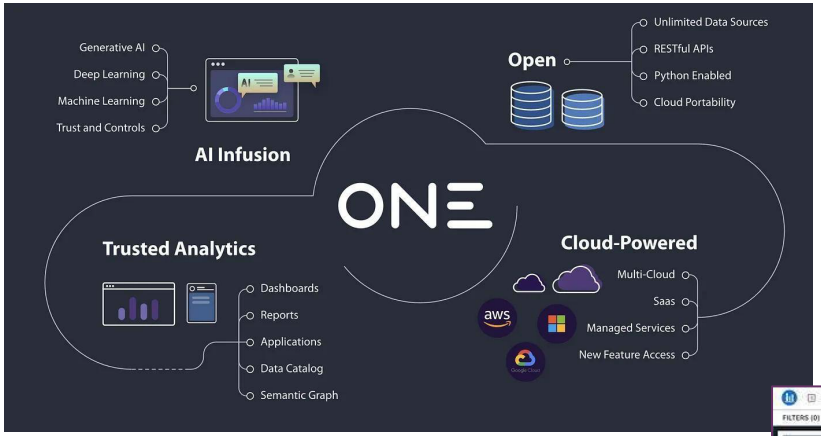

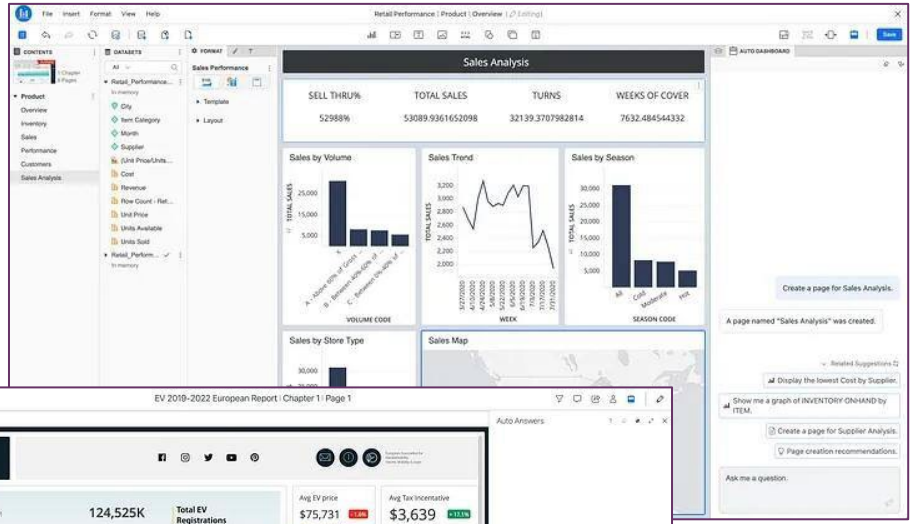

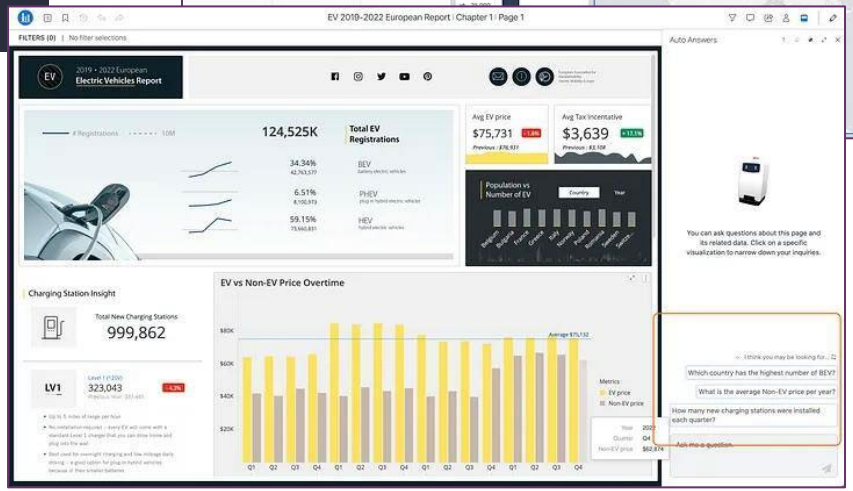

[AI en Design Thinking in de Dashboardmeesterschap van MicroStrategy]

Fontys

GenAI for Self-Service Analytics

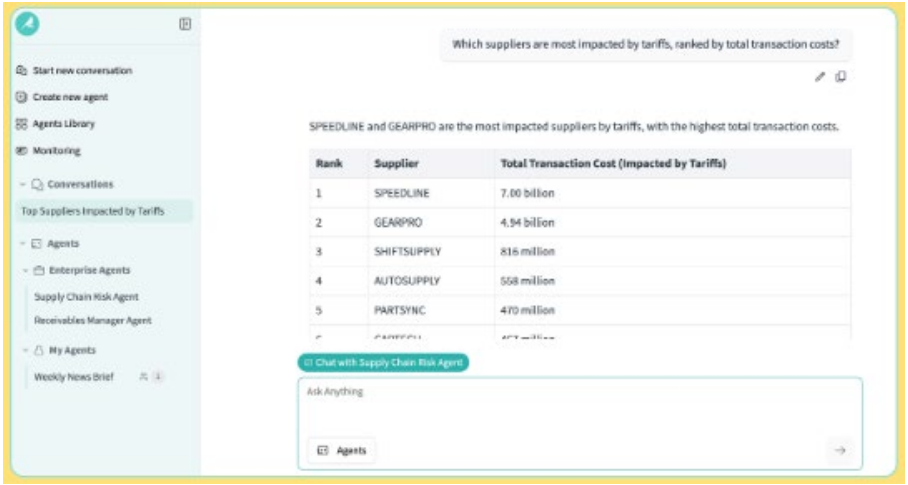

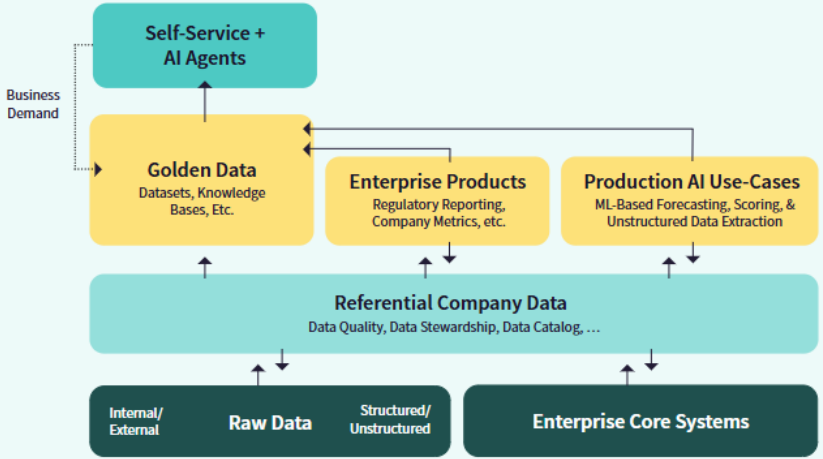

Fontys

[Self-Service Analytics In The Age Of AI Agents]

Agentic AI for Data Lakehouse

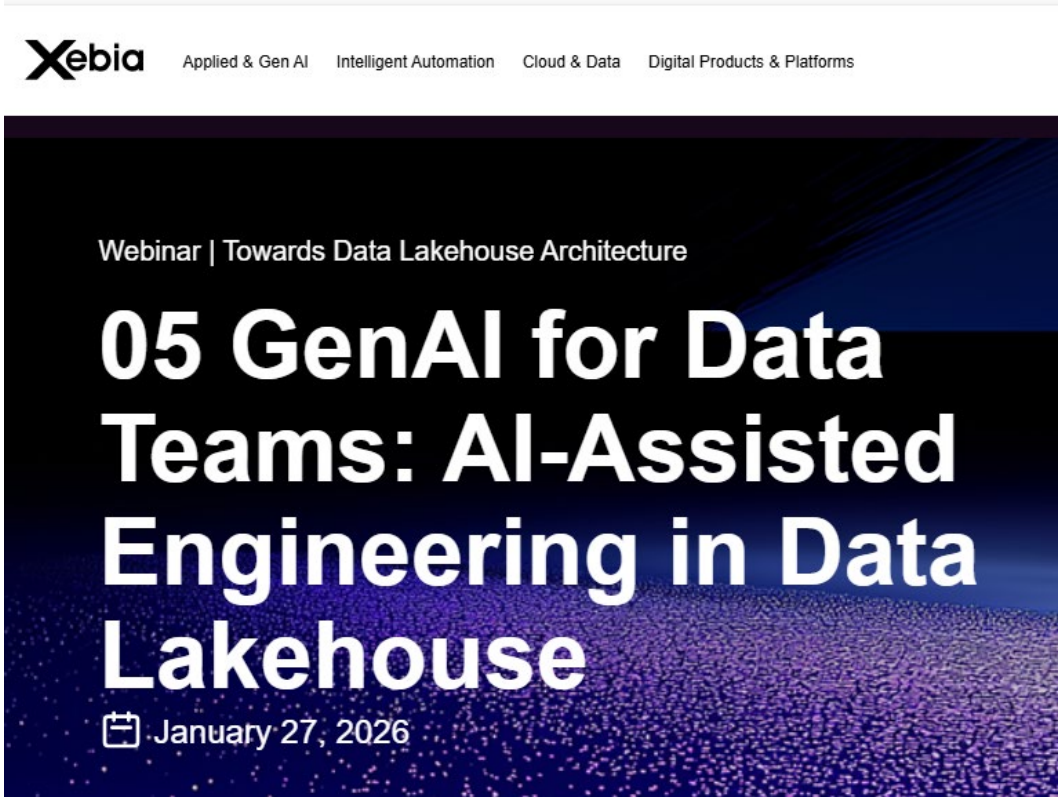

“how **Agentic AI** shifts data teams from hand-built pipelines to **outcome-driven data products** guided by clear intent and intelligent assistants.

You’ll learn how **MCP integrations** enable AI assistants to tap your data catalogs, databases and transformation tools like dbt - dramatically accelerating development while preserving governance and trust.

We’ll show how **spec-driven development** turns business needs and technical requirements into unambiguous, testable data products - boosting delivery speed and confidence.

Finally, we’ll explore **conversational data interfaces** that allow both technical and business users to interact with data products using natural language.”

https://events.xebia.com/5-towards-data-lakehouse-architecture-webinar-ai-assisted-engineering

Agentic DataOps

Agentic DataOps

Proposal: Agentic data stack that can autonomously resolve data asset and insight requests from a stakeholder (including AI itself).
Process creates a system that is comprised of all data stack stages:

- Design and implement data infrastructure
- Update and optimize existing data infrastructure & storage
- Discover data sources and information (*e.g. Search, Enrichment*)
- Create multi-modal data flow (*e.g. Ingestion, Integration*)
- Monitor and remediate data flows *across the data life-cycle*
- Ensure governance requirements are enforced
- Uncover actionable insights (*and potential for decision making*)

Impact: Democratize value creation from data at speed and scale. Bring time to value from weeks down to hours by decreasing task times and reducing human touchpoints / bottlenecks.

Example: Fund Performance Forecasting

Create financial product that forecasts performance across Asian and European mutual funds.

Additional requirements include low latency, GDPR compliance, price-performance, high data quality, restricted user access.

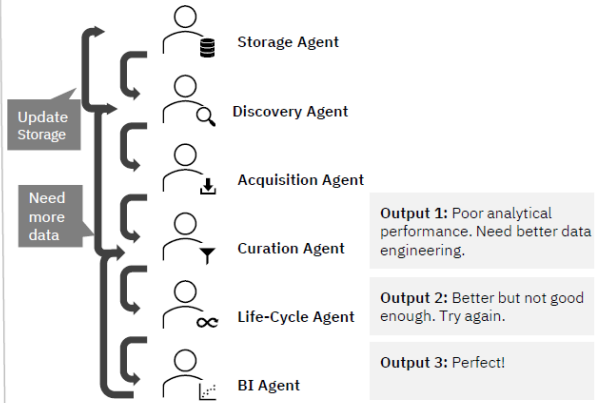

IBM Research

Fontys

[Keynote Lisa Amini-What's Next in AI for Data and Data Management--Pydata Global 2025]
[2512.07926] Can AI autonomously build, operate, and use the entire data stack?

Building AI-Enabled Metadata-Driven Data Platform Teams

Jaco van der Laan he/him

905 followers

Exploring Business & Logical Data Modeling.
Writing on Clarity, Structure & Creative Approaches to Data Architecture.

Follow

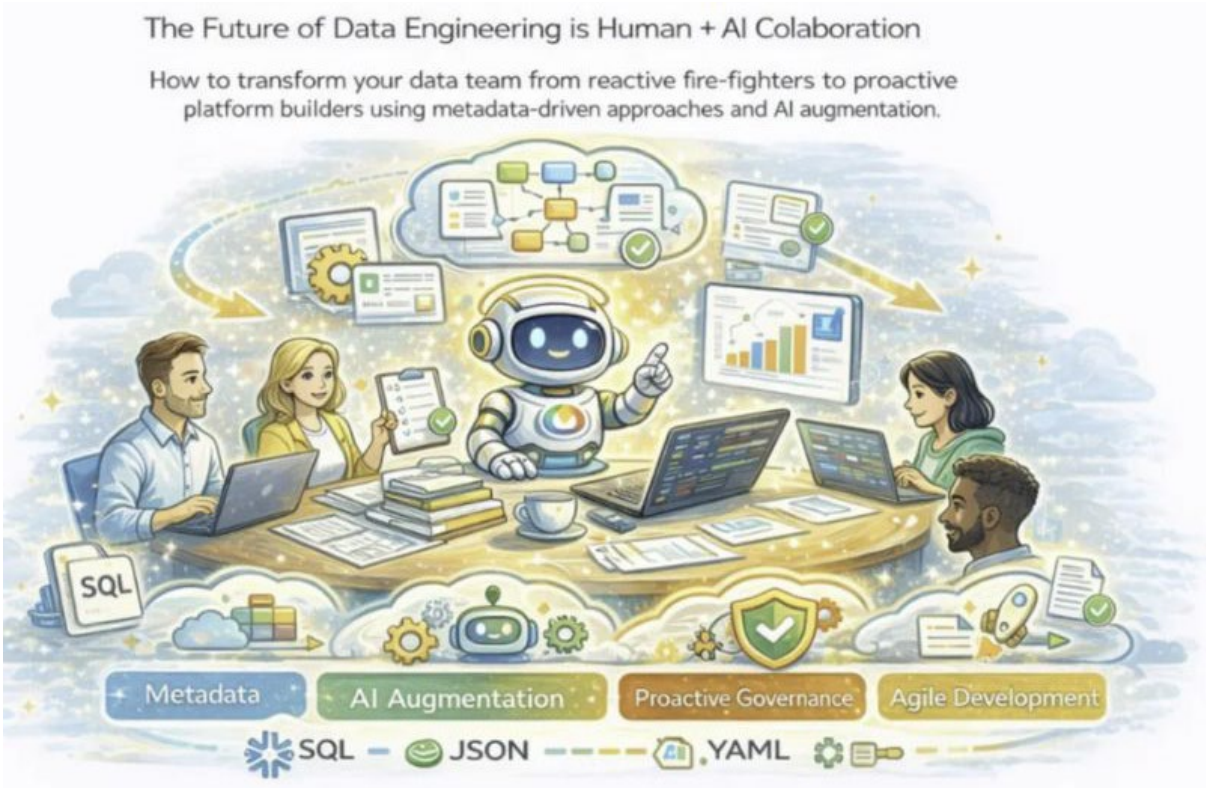

(13) Post | LinkedIn

Where Should Data Documentation Live?

Building AI-Enabled, Metadata-Driven Data Platform Teams

The Knowledge Repository Pattern: One Folder Per Artifact

Fontys

Other Examples

- AI Agents With Dataiku | Dataiku
- Introducing Genie Code | Databricks Blog
- Actian introduceert Data Observability Agents | BI-Platform
- IBM Db2 Genius Hub voegt autonome mogelijkheden toe aan Db2 | BI-Platform
- LLMs4OM: Matching Ontologies with Large Language Models

Fontys

Fontys > FOR SOCIETY

GENAI ENGINEERING @ FONTYS ICT

Fontys ICT – Applied GenAI

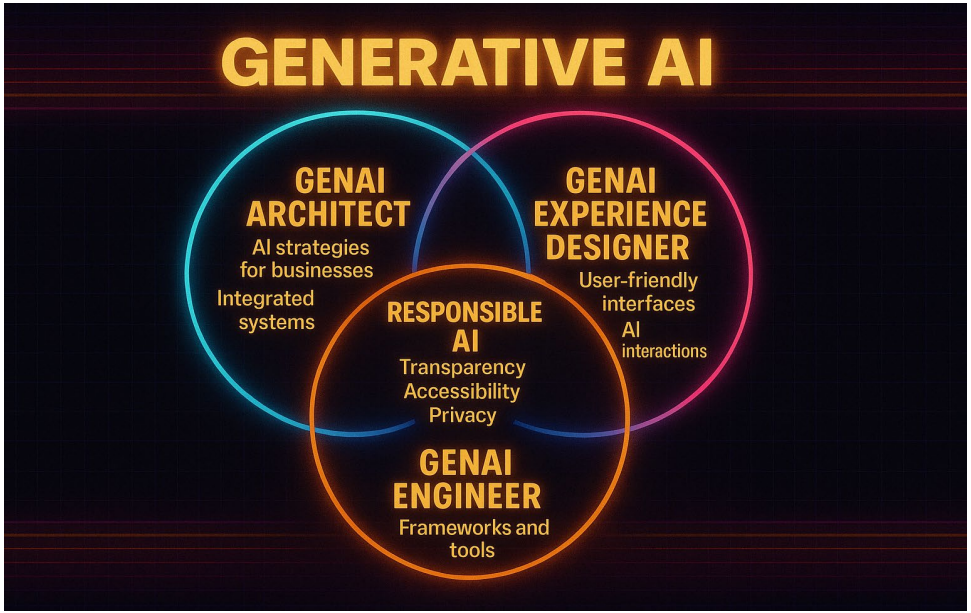

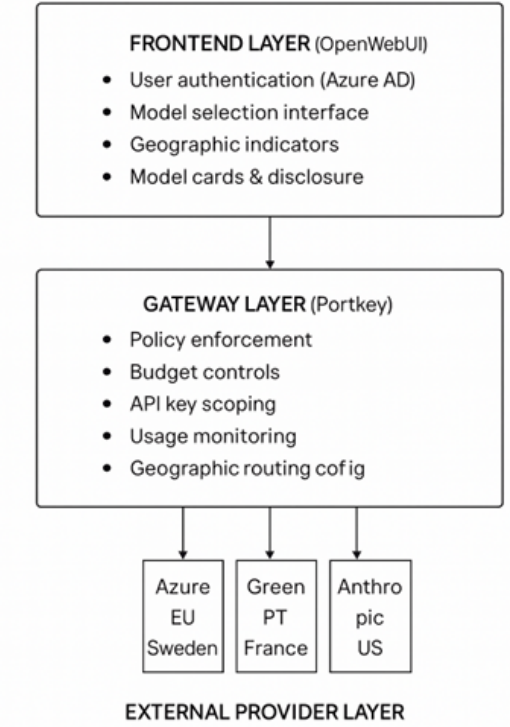

Fontys

Study Guide Fontys ICT - Applied Generative AI

[2512.08978] Institutional AI Sovereignty Through Gateway Architecture: Implementation Report from Fontys ICT

Fontys ICT – Study Programmes

Study programme Software Engineering

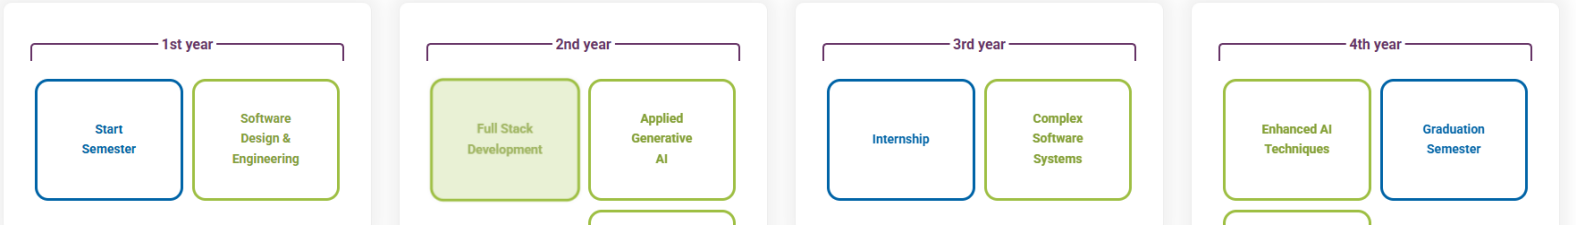

Function Prompt Engineer

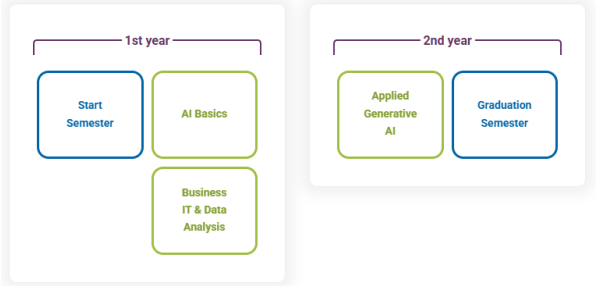

Function Software Prompt (AI Integration) Engineer

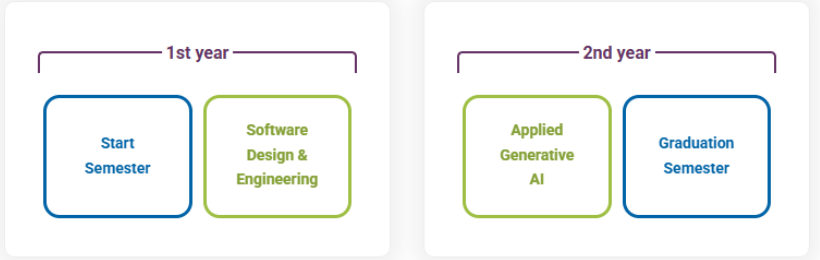

[Study Guide Fontys ICT - Programmes]

Applied Research Projects

- “A Validation Framework for LLM applications” – Professional Doctorate Leon Schrijvers
- “Architecture patterns for modifiable AI systems” – Phd Merel Veracx
- “Active Model Cards” – Koen Suilen & Ruud Huijts – under submission
- “What is a GenAI engineer?” – Petra Heck & Leon van Bokhorst
- DEMAND (Sprong) – Data engineering and data management in data chains – with HAN and Saxion
- ….
- **Several projects where we apply LLMs to innovate professional practice of other domains**

Fontys

Fontys > FOR SOCIETY

CONCLUSION

Lessons We Are Learning

- GenAI engineering is software engineering
  - New **architectures** for systems, new **tech stack**
  - From Devops to MLOps to **LLMOps** to **AgentOps**
- Don’t forget about the **data** (and data engineering)
- From automation to augmentation, new processes?
- Today’s impossible is tomorrow’s possible
  - So start yesterday, learning by doing

Let’s learn together! p.heck@fontys.nl

DEMAND

Fontys

DEMAND – Handling Data Safely, Responsibly, and Efficiently

Questions?

Petra Heck, auteur op Fontys, p.heck@fontys.nl, https://www.linkedin.com/in/petraheck/

Fontys › FOR SOCIETY

How do software engineers use AI-coding assistants?

The benefits of AI are not equally distributed!

How do others use it? — What practices and patterns are developing? — How to avoid damaging the talent-pipeline? — How to use it to improve quality and productivity? — How to keep developers happy and healthy?

We are starting a **qualitative research** effort about how software engineers use AI in their work (using rigorous research methods and frameworks such as DORA & SPACE)

🥇 We would like to find **partners, contributors and participants** for going on this journey with us.

❓ You can help us shape this research, contribute expertise, and get answers

💬 Interested? Send a mail to marc.vangrootel@fontys.nl or scan this code