



# Enterprise Scale AI Factory

AI Ready infrastructure automation  
+ Templates for DataOps, MLOps, GenAIOps

Enterprise Scale AI Factory is open source (MIT license) on [GITHUB](#)  
+5 years old



jostrm 2019  
azure-enterprise-scale-ml



Azure 2024  
enterprise-scale-aifactory



# AI

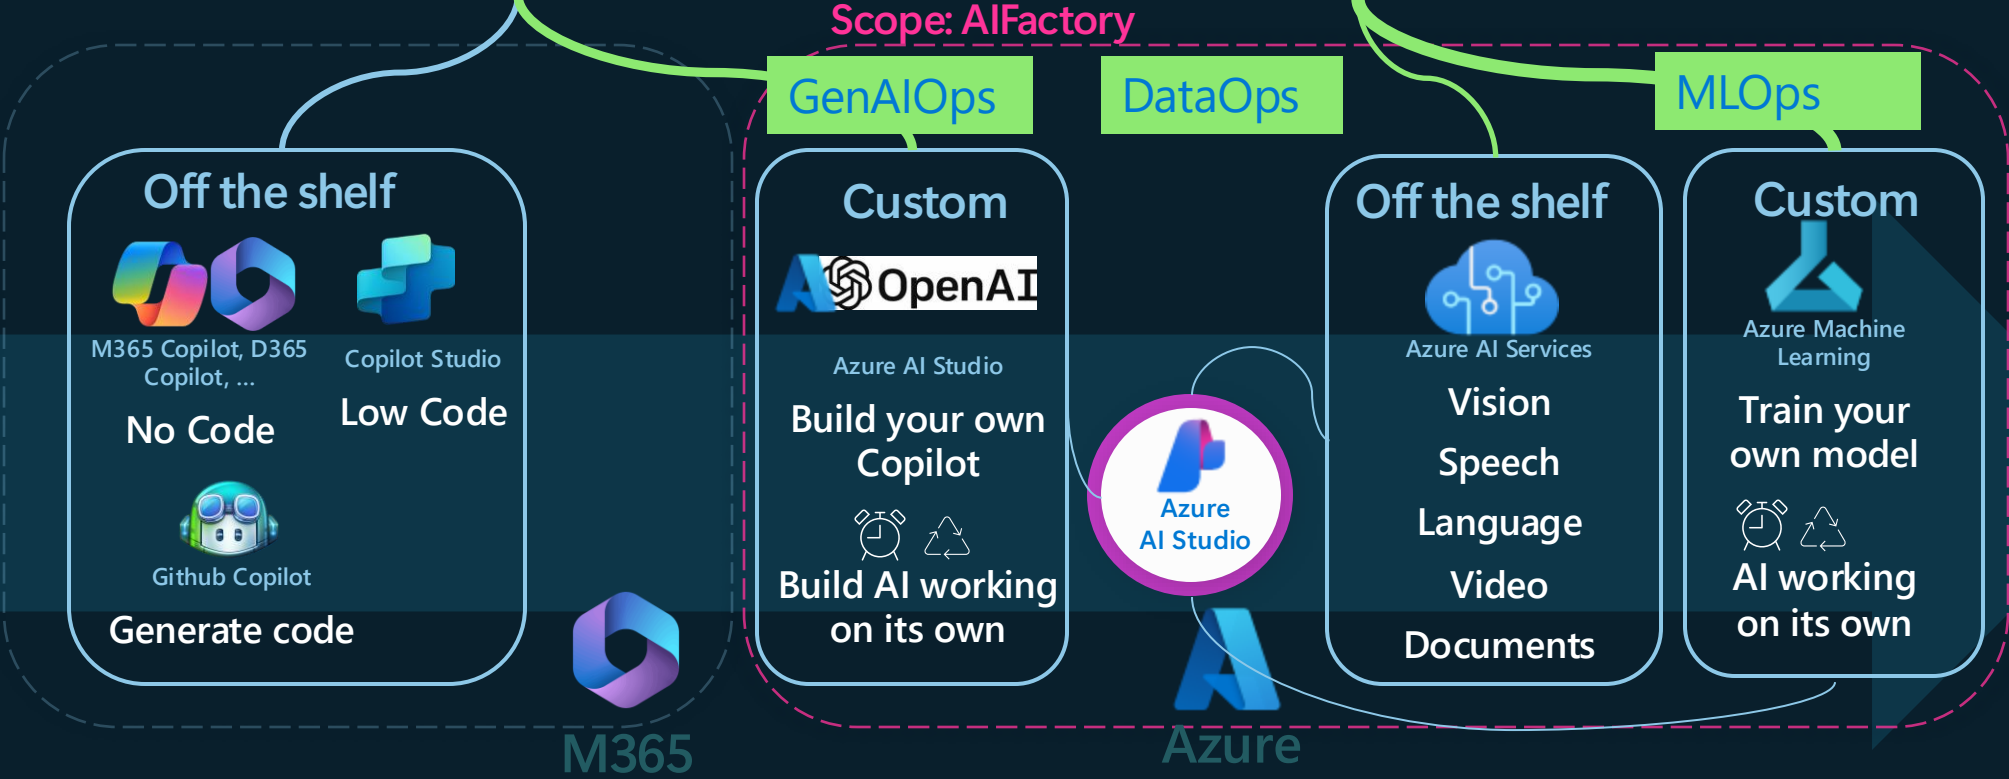
Categories of AI

**Your use case**  
Summarize surveys?  
Analyze IT logs?  
Internal chat?

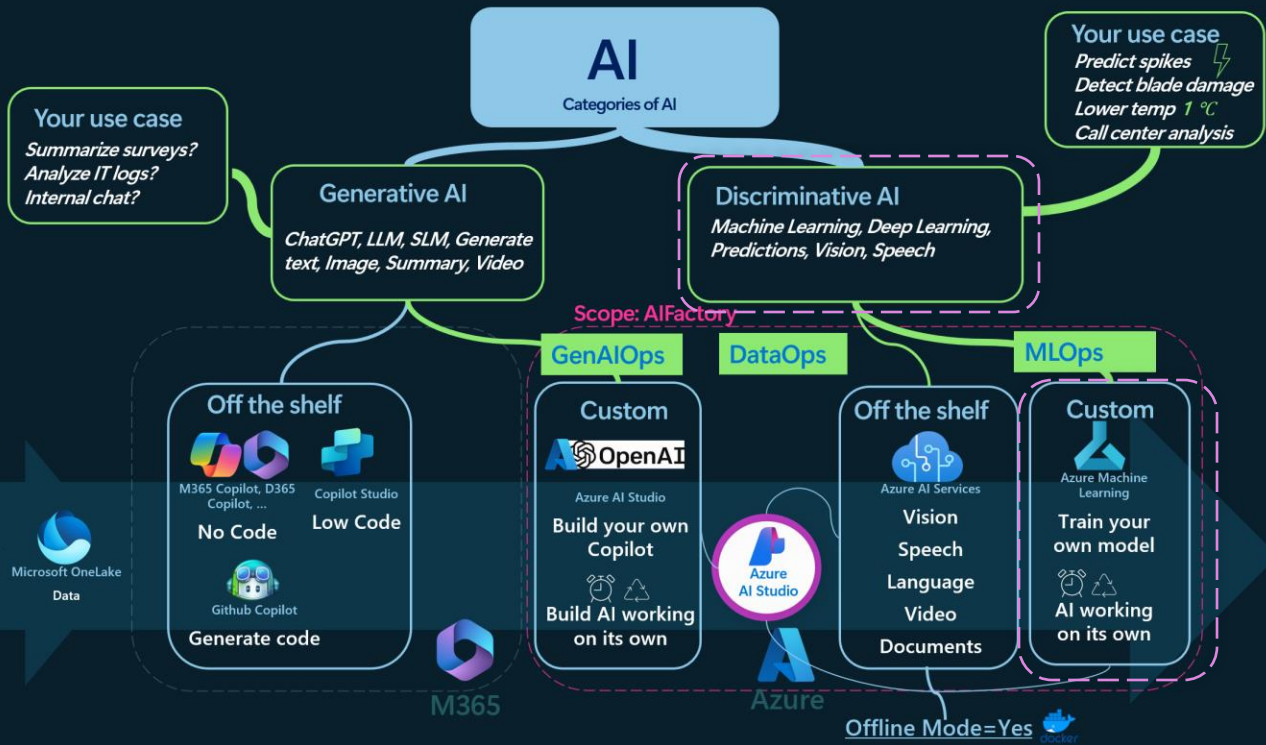
**Generative AI**  
ChatGPT, LLM, SLM, Generate text, Image, Summary, Video

**Discriminative AI**  
Machine Learning, Deep Learning, Predictions, Vision, Speech

**Your use case**  
Predict spikes ⚡  
Detect blade damage  
Lower temp 1°C  
Call center analysis



# AI types and use cases

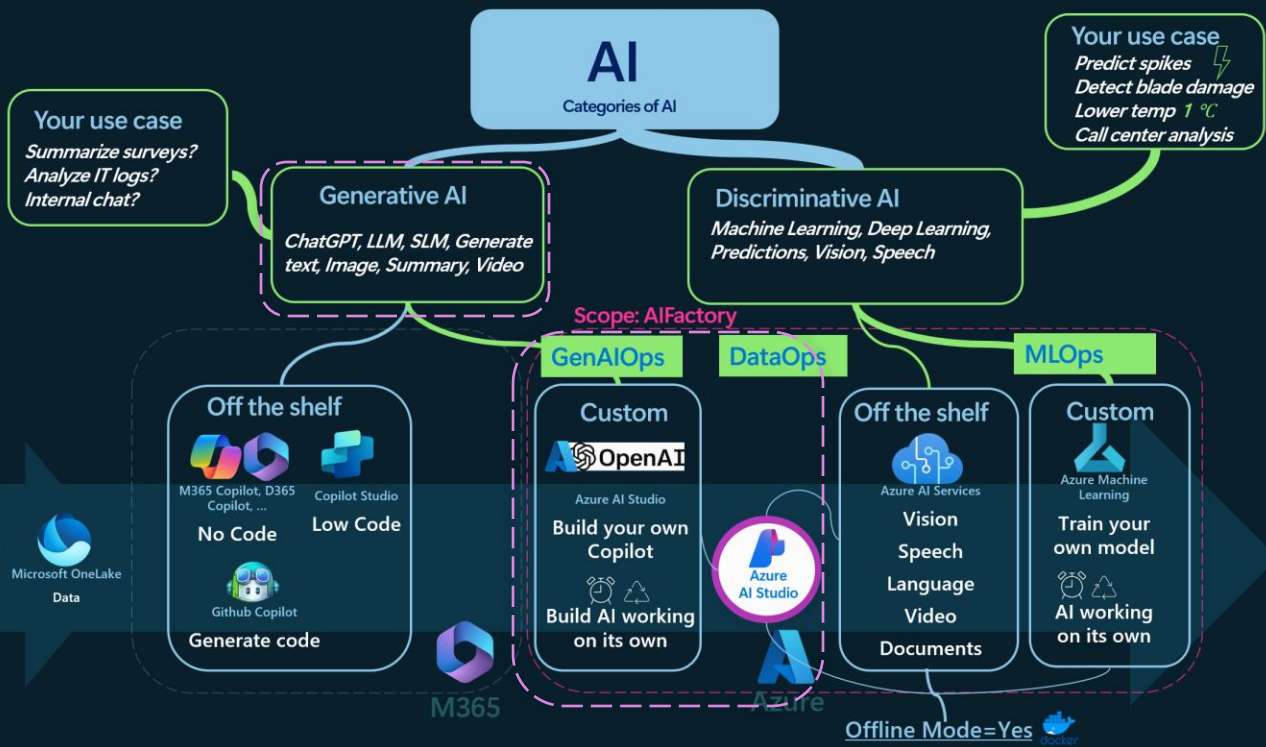


Epiroc – “Predict steel quality”

SAS – “Reduce 45% food waste”

Sweco – “Contract analysis”

# AI types and use cases



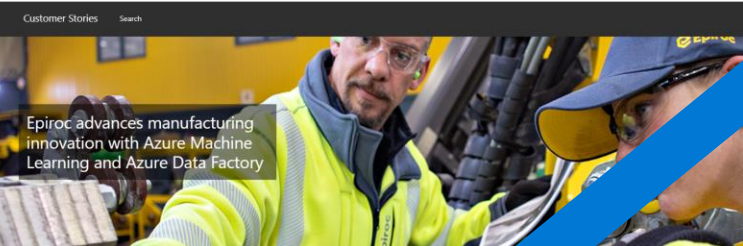
Epiroc – “Predict steel quality”

SAS – “Reduce 45% food waste”

Sweco – “Contract analysis”

**Epiroc**  
Year: 2021

https://customers.microsoft.com/en-us/story/1653030140221000726-epiroc-manufacturing-azure-machine-learning



Customer Stories Search

Epiroc advances manufacturing innovation with Azure Machine Learning and Azure Data Factory

September 1, 2023

With its headquarters in Sweden, Epiroc is a leading manufacturer of rock excavators. The company sources premium steel and processes it at locations across the globe to produce tools and equipment for the mining and construction industries. Epiroc also provides maintenance, repair, and replacement service and technical expertise. To improve efficiency, Epiroc wanted to improve and automate its processes at its manufacturing locations worldwide. To ensure consistent quality of the steel used in its products, the help of Microsoft, Epiroc created an ESML AI Factory on Microsoft Azure. The company was able to get it up and running within 60 hours. The scalable solution gave Epiroc's manufacturing teams across the globe the ability to share data, best practices, and create consistent, repeatable processes to ensure the highest quality steel—and equipment—while increasing efficiencies and reducing waste.

Learn More

- Microsoft Azure
- Azure Machine Learning
- Azure Data Factory
- Azure Databricks

Customer: Epiroc

Products and Services: Azure Data Factory, Azure Databricks (d), Azure Machine Learning, Microsoft Intelligent Data Platform, Power BI

Industry: Manufacturing

Organization Size: Enterprise (50,000+ employees)

Country: Sweden

Downloadable: Epiroc story summary slide


"We felt confident that with Microsoft we could create a..."

**Acceleration**  
4-60h  
VS  
**Challenge**  
~1000-2000h (~years)

Public customer references  
<https://customers.microsoft.com/>

**SAS**  
Year: 2019

https://customers.microsoft.com/en-us/story/781802-sas-travel-transportation-azure-machine-learning



Customer Stories Search

Scandinavian Airlines reduces loyalty program fraud with Microsoft Azure Machine Learning

May 13, 2020

To derive meaningful business value from the wealth of data it generates every day, Scandinavian Airlines (SAS) uses Microsoft Azure Machine Learning to make predictive models to improve everyday operations and identify patterns in the company's data. From reducing fresh food waste to cutting down on fraud in its EuroBonus loyalty program, SAS is harnessing the power of its data using Azure Machine Learning and its responsible machine learning capabilities, including model interpretability and explainability. These capabilities help data scientists understand models better, create transparency about their processes, build trust in machine learning solutions, and create better customer experiences.

Learn More

- Microsoft Azure AI
- Azure Machine Learning
- Machine learning operations

Customer: Scandinavian Airlines

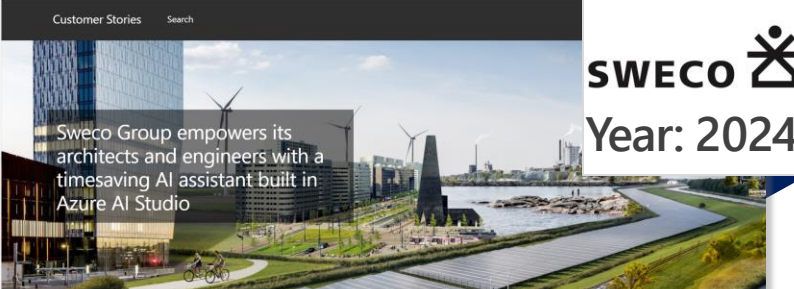
Products and Services: Azure, Azure Data Factory, Azure Databricks (d), Azure DevOps, Azure Subscriptions Service, Azure Machine Learning, Microsoft Visual Studio

Industry: Travel and Transportation

Organization Size: Enterprise

"We use Azure Machine Learning to solve real business problems without worrying about building and managing infrastructure or creating new tools—we can focus directly on gaining value from the..."

customers.microsoft.com/en-us/story/1167395222121336377-swecogroup-azure-ai-services-professional-services-en-sweden



Customer Stories Search

Sweco Group empowers its architects and engineers with a timesaving AI assistant built in Azure AI Studio

May 11, 2024

As one of Europe's leading architecture and engineering firms, Sweco makes it a priority to empower its more than 22,000 consultants with innovative technology. To free up time for more creative solutions in client projects, Sweco used Microsoft Azure AI Studio to rapidly develop a digital assistant to automate document creation and analysis and enhance search capabilities across the enterprise's global network. With the assistant, called SwecoGPT, consultants can now find critical project information quickly, create and analyze documents easily, and use the time they save to deliver more personalized service to customers around the world. Consultants are reporting increased productivity and added client value. With the success of SwecoGPT for internal use, the company plans to expand its use of Azure AI Studio to support digital innovation and to create solutions to enhance the...

Learn More

- Azure AI Studio
- Azure OpenAI Service
- Azure AI Services
- Azure prompt flow

Customer: Sweco


Products and Services: Azure AI Services, Azure AI Studio, Azure Machine Learning

Year: 2024

ESML AI Factory as a service = "SARAH"

**sogeti**  
Part of Capgemini  
Year: 2020  
Läs mer om vårt Azure Intelligence Center här!

https://www.sogeti.se/sarah-eng



sogeti  
Part of Capgemini

Om oss Varför Sogeti? Lösningar Utforska Jobba hos oss Kontakt

Sogeti's Mark Oost about SARAH

Watch on YouTube

The scaling technology

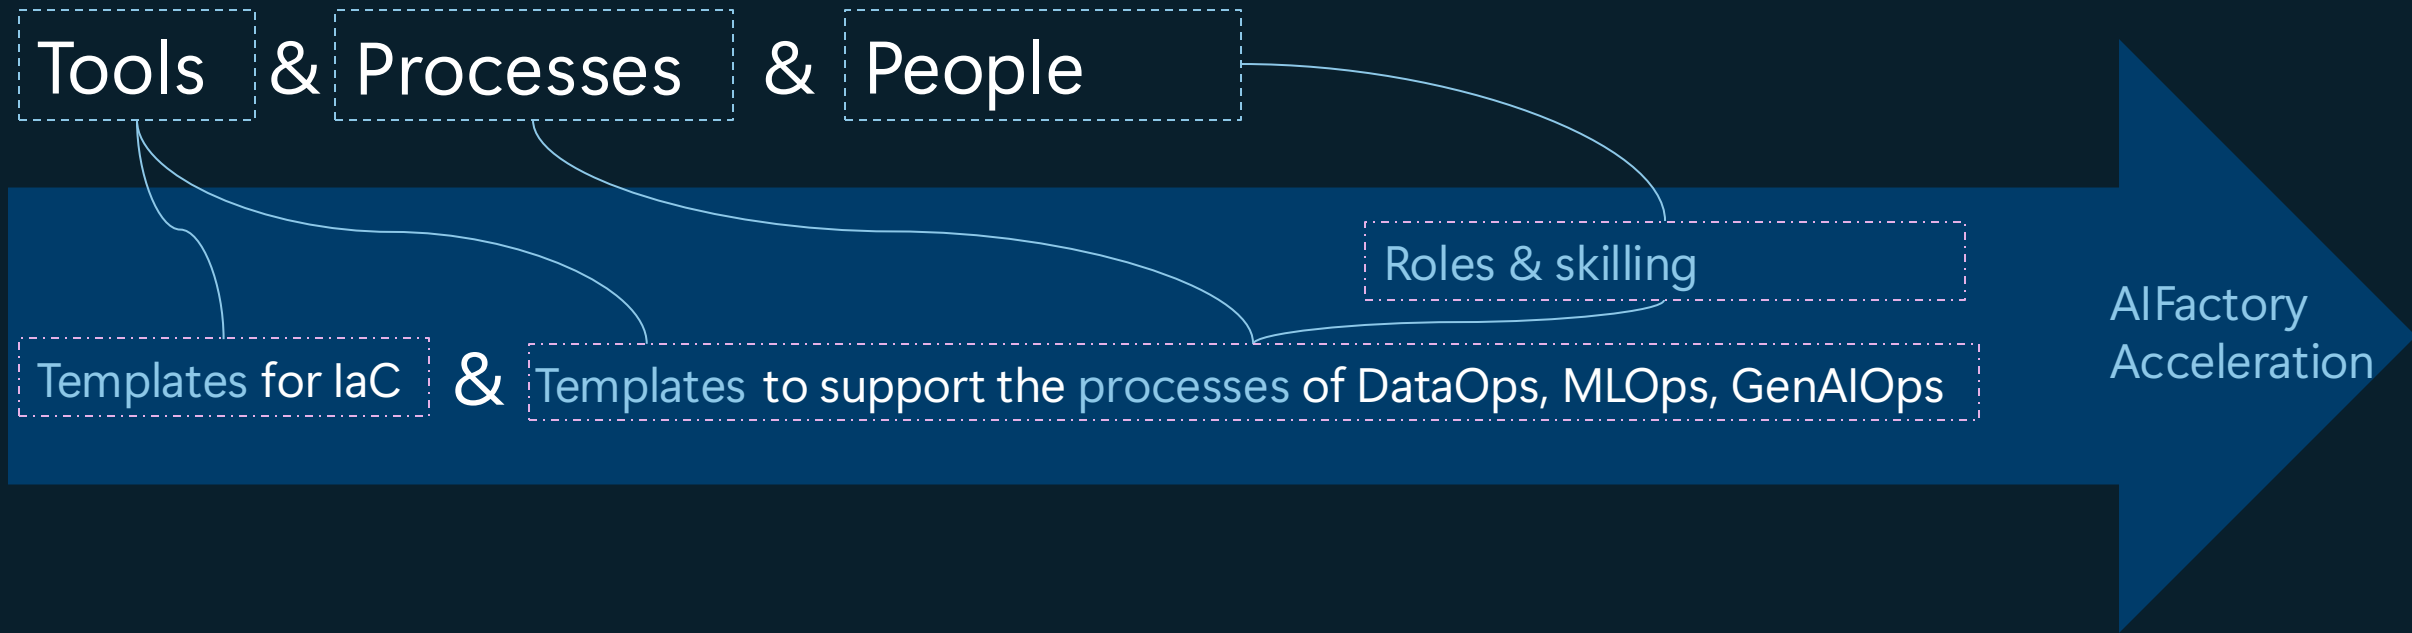
With code, you can set-up development, testing, and production infrastructure for various initiatives, by using ready-to-use scripts. No manual work is required, your infrastructure is set up automatically in Azure. You only pay for the Azure consumption you're using, and when a project isn't needed anymore, you can safely bring it down and know that all the code is left to re-script the entire setup again.

Here are further details on how SARAH gives you an accomplished Data & AI platform. For example, best practice for a modern data platform with lambda architecture, is implemented, including Auto-Lake. The HILops Framework allows you to create a pipeline to train a Machine Learning model with only two lines of code, something that typically takes close to 2000 lines of code. Of course, well-known frameworks such as CAF (Cloud Adoption Framework) and WAF (Well-Architected Framework) are also implemented which provides you with a safe, stable, and secure environment. The next generation of data platforms with Data Mesh architecture are also built-in, through support via a 'Share Back' function that enables various initiatives/projects to subscribe to data sources created by other parts of your organization.

Combined expertise with big muscles

SARAH is the result of a close collaboration between Sogeti, Microsoft and the team behind the Open Source Framework ESML (Enterprise Scale Machine Learning). With ESML as a foundation, we have added additional functionality based on our knowledge and experience of large-scale Data & AI

# WHAT is the AI Factory?



# AI Factory 4 components, assets (97 109 25 1 5 14)

AI ready landing zones:  
Enterprise scale: brown field & green field



Datalake template:  
Datamesh, Medallion, IoT



DataOps, MLOps,  
GenAIOps: End-2End templates



Scale & Delivery: common way of working, project way of working.


+Accelerator *ESML SDK*

01 Resiliency, Security from WAF & CAF, +Automation

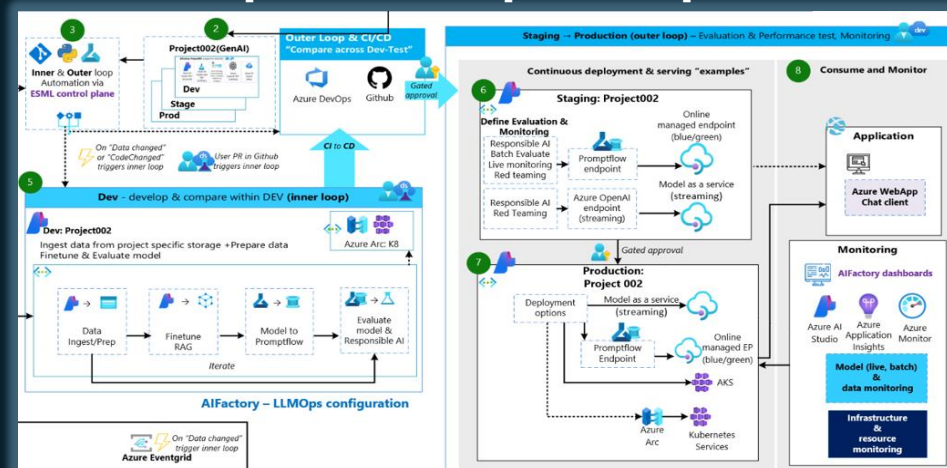
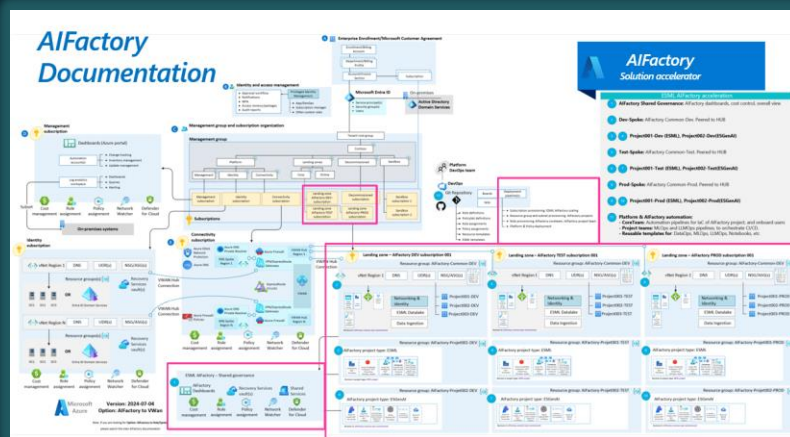
02 Battle-tested lake design

03 Configure VS Build

04 MS EDE/ISD, Partners



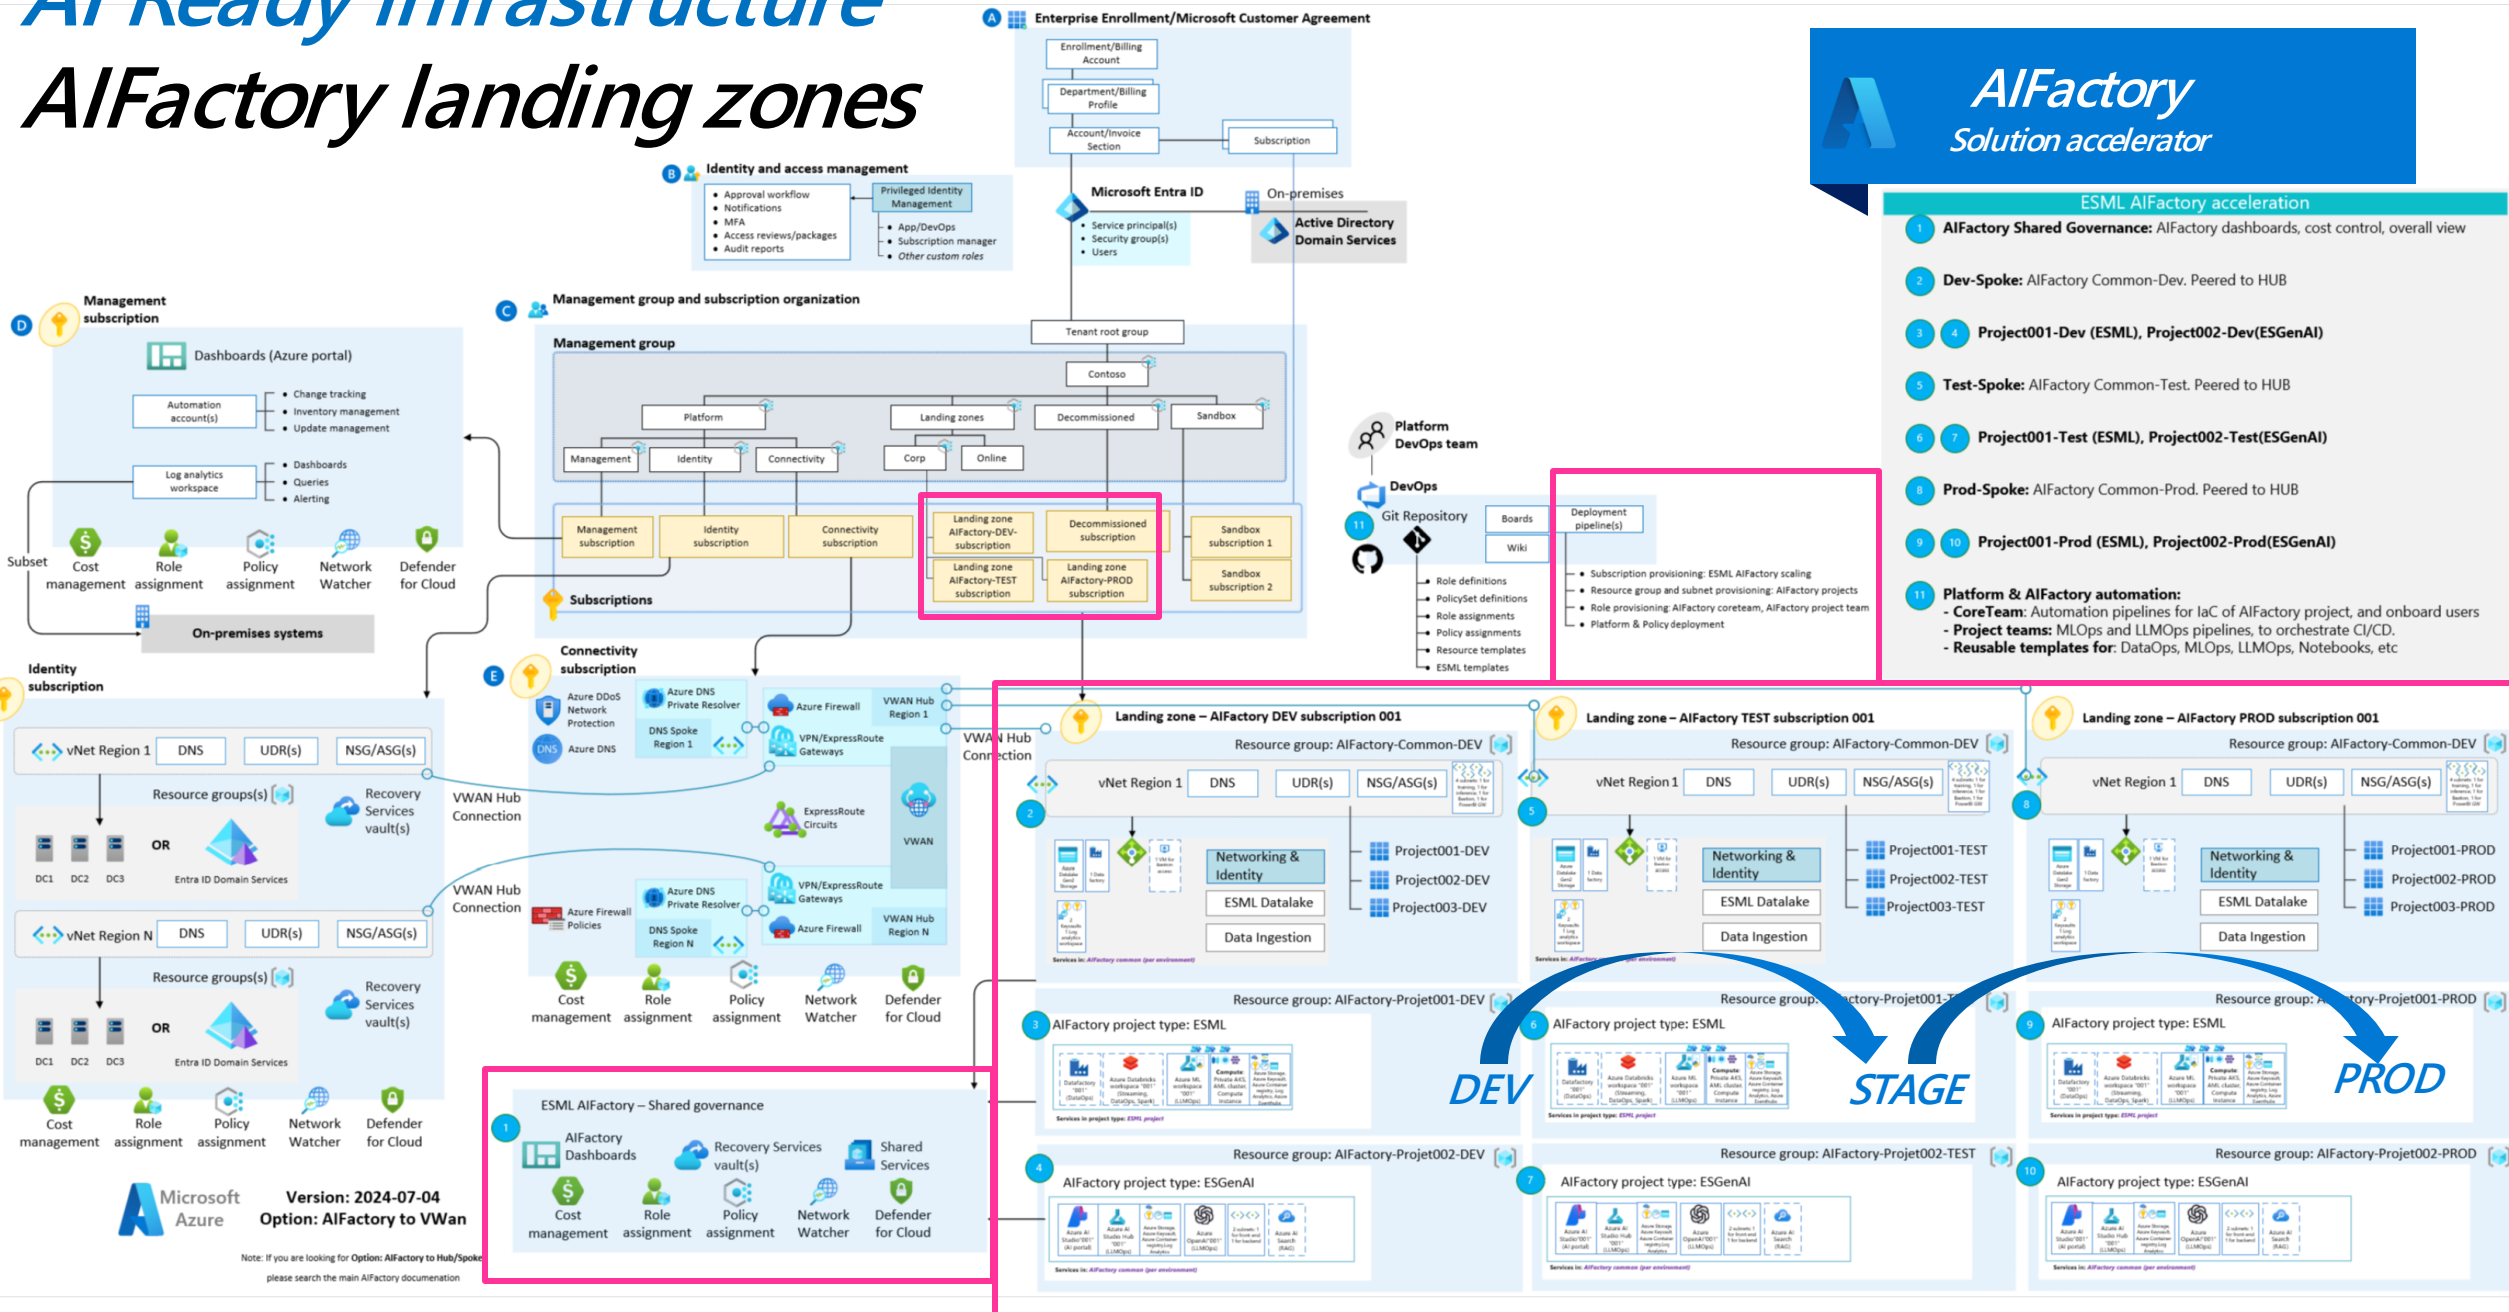
## DataOps+GenAIOps/MLOps



Your AI Factory  
Your use cases

# AI Ready infrastructure

## AI Factory landing zones



- ### ESML AI Factory acceleration
- AI Factory Shared Governance:** AI Factory dashboards, cost control, overall view
  - Dev-Spoke:** AI Factory Common-Dev. Peered to HUB
  - Project001-Dev (ESML), Project002-Dev(ESGenAI)**
  - Test-Spoke:** AI Factory Common-Test. Peered to HUB
  - Project001-Test (ESML), Project002-Test(ESGenAI)**
  - Prod-Spoke:** AI Factory Common-Prod. Peered to HUB
  - Project001-Prod (ESML), Project002-Prod(ESGenAI)**
  - Platform & AI Factory automation:**
    - Core Team:** Automation pipelines for IaC of AI Factory project, and onboard users
    - Project teams:** MLOps and LLMops pipelines, to orchestrate CI/CD.
    - Reusable templates for:** DataOps, MLOps, LLMops, Notebooks, etc
  - Platform DevOps team**
  - DevOps**
  - Git Repository**
  - Boards**
  - Wiki**
  - Deployment pipeline(s)**
    - Subscription provisioning: ESML AI Factory scaling
    - Resource group and subnet provisioning: AI Factory projects
    - Role provisioning: AI Factory coreteam, AI Factory project team
    - Resource templates
    - ESML templates



Version: 2024-07-04  
Option: AI Factory to VWAN

Note: If you are looking for Option: AI Factory to Hub/Spoke please search the main AI Factory documentation










# 2 packages of services

Discriminative AI "Classical AI" MLOps + DataOps  
*Machine Learning, Deep Learning, Predictions, Vision, Speech*

ESML (est 2019)








Services in project type: **ESML project**

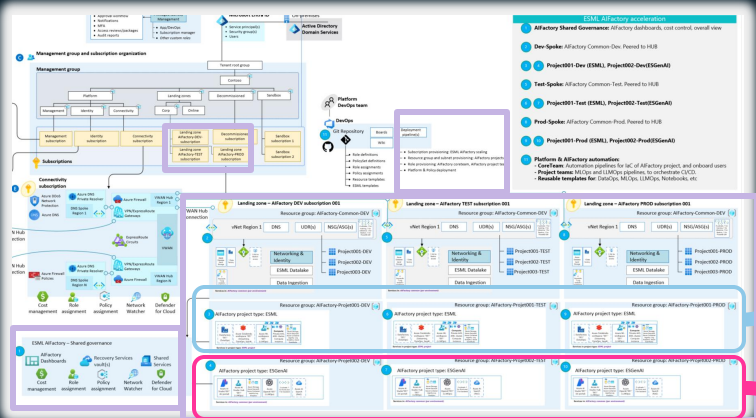
 Datafactory "001" (DataOps)	 Azure Databricks workspace "001" (Streaming, DataOps, Spark)	 Azure ML workspace "001" (LLMOps)	 <b>Compute:</b> Private AKS, AML cluster, Compute Instance	 Azure Storage, Azure Keyvault, Azure Container registry, Log Analytics, Azure Eventhubs, VM	 3 subnets: 2 for databricks, 1 for AKS	 1 VM for Bastion access
--	---	--	---	--	---	--

Generative AI "ChatGPT" – GenAIOps  
*ChatGPT, RAG, Agentic, Generate text, Image, Summary, Video*

GenAI (est 2024)


Services in project type: **ESGenAI project**

 Azure AI Studio "001" (AI portal)	 Azure AI Studio Hub "001" (LLMOps)	 Azure Storage, Azure Keyvault, Azure Container registry, Log Analytics	 Azure OpenAI "001" (LLMOps)	 2 subnets: 1 for front-end 1 for backend	 Azure AI Search (RAG)	 Azure Webapp (UI)
--	---	---	--	---	--	--



# AI Factory 4 components, assets: 97 109 25 1 5 14


AI ready landing zones:  
Enterprise scale: brown field & green field




Datalake template:  
Datamesh, Medallion, IoT



DataOps, MLOps,  
GenAIOps: End-2End templates




Scale & Delivery: common way of working, project way of working.  
+Accelerator *ESML SDK* 

01 Resiliency, Security from WAF & CAF, +Automation

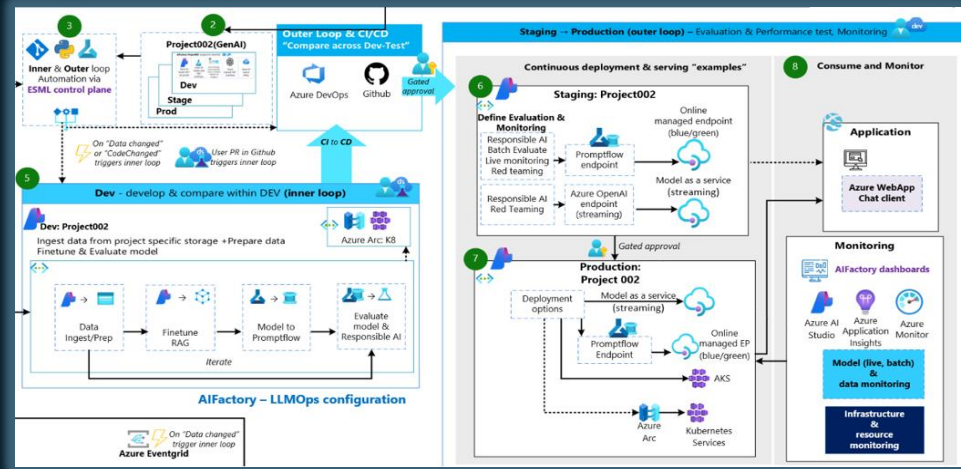
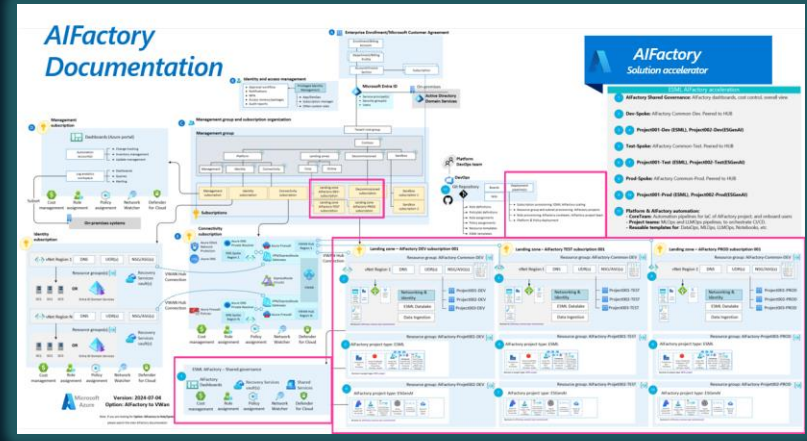
02 Battle-tested lake design

03 Configure VS Build

04 MS EDE/ISD, Partners



## DataOps+GenAIOps/MLOps



Your AI Factory  
Your use cases



Tools

&

Processes

&

People

# Governance – cost, automation, onboarding

Process to create more projects?

Process of GenAIOps, MLOps?

How to onboard a team to the AI Factory?

# Ticket to get a AIFactory project

## ESMLProject onboarding

How the business can order a ESML project to be setup in the ESML AI factory

Result 1: Coreteam review, and auto-provision an ESMLProject in Azure (including AD-group)

Result 2: Email sent out with ESML Project number. Example: "Project003"

Project owner \*

Business unit \*

Cost center \*

AD-Users/group \*

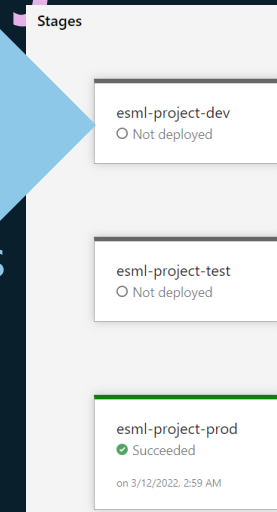
Model/Project folder

Model prefix:3 chars   
Maximum of 3 characters. *Currently Used: 0 characters.*

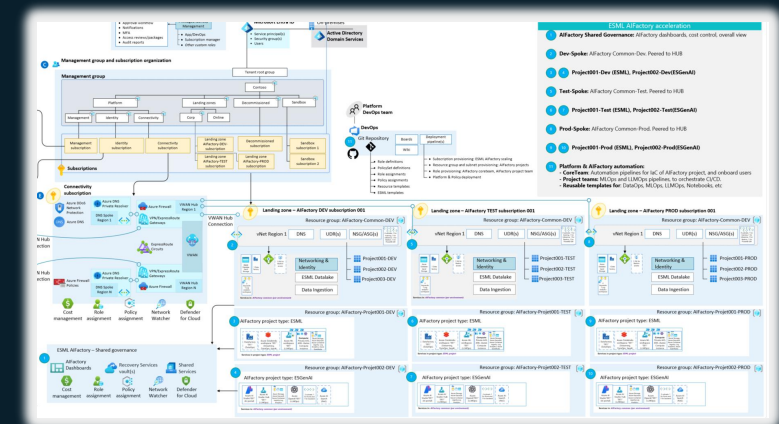
Project type (ml,generic):

Project environments  Dev  Test&Prod  Dev,Test,Prod

1-250 projects



Wait 30 min



✓ Each ESMLProject own area



Tools & Processes & People

Processes: DataOps, MLOps, GenAIOps

Templates: GenAIOps example

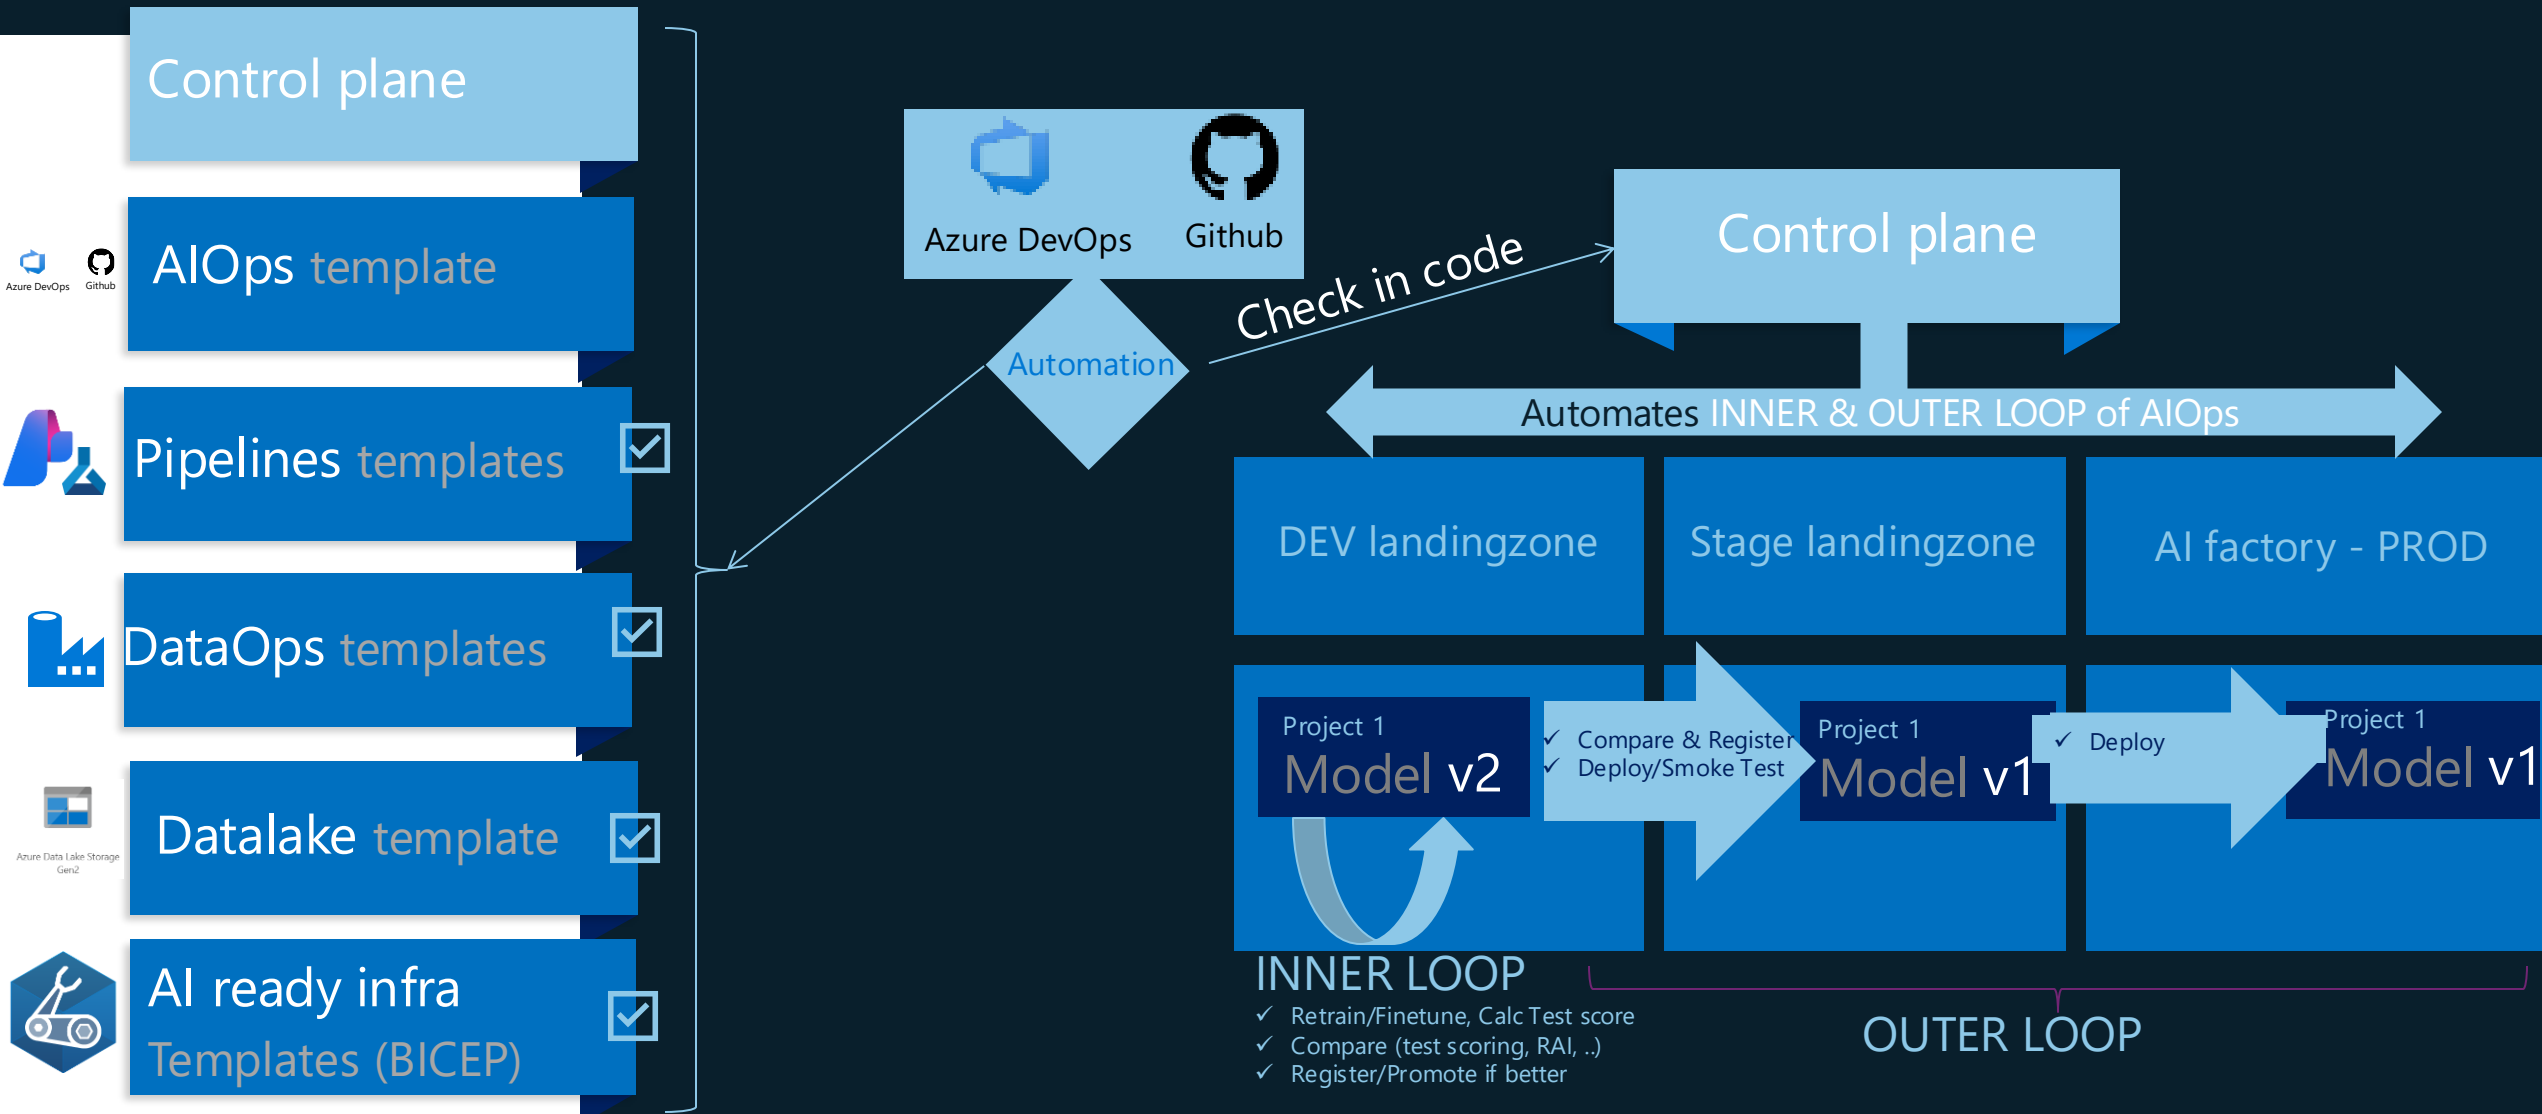


# Maturity stages: 5 levels of AIOps

	Maturity Level	Finetuning/Evaluation Process	Release Process	Process automated
Inner loop	Level 1 – No Ops	UNTRACKED	MANUAL	-
	Level 2 – Pipeline for Finetuning, Grounding, Evaluation	Automated	MANUAL	AIOps (inner)
	Level 3 – Release Pipeline	Automated(Code)	“CodeChanged” with DevOps	AIOps (inner)
	Level 4 –CI/CD + DataChanged	Automated (Code+Data)	“DataChanged” trigger with EventGrid	DataOps+AIOps (inner)
Outer loop	Level 5 – Outer loop AIOps: Enterprise security & monitoring. App integration	Dev,Test,Prod environments No Public IP	Dev,Test,Prod environments No Public IP	DataOps+AIOps (inner / outer) + Feedback loop: data back for training

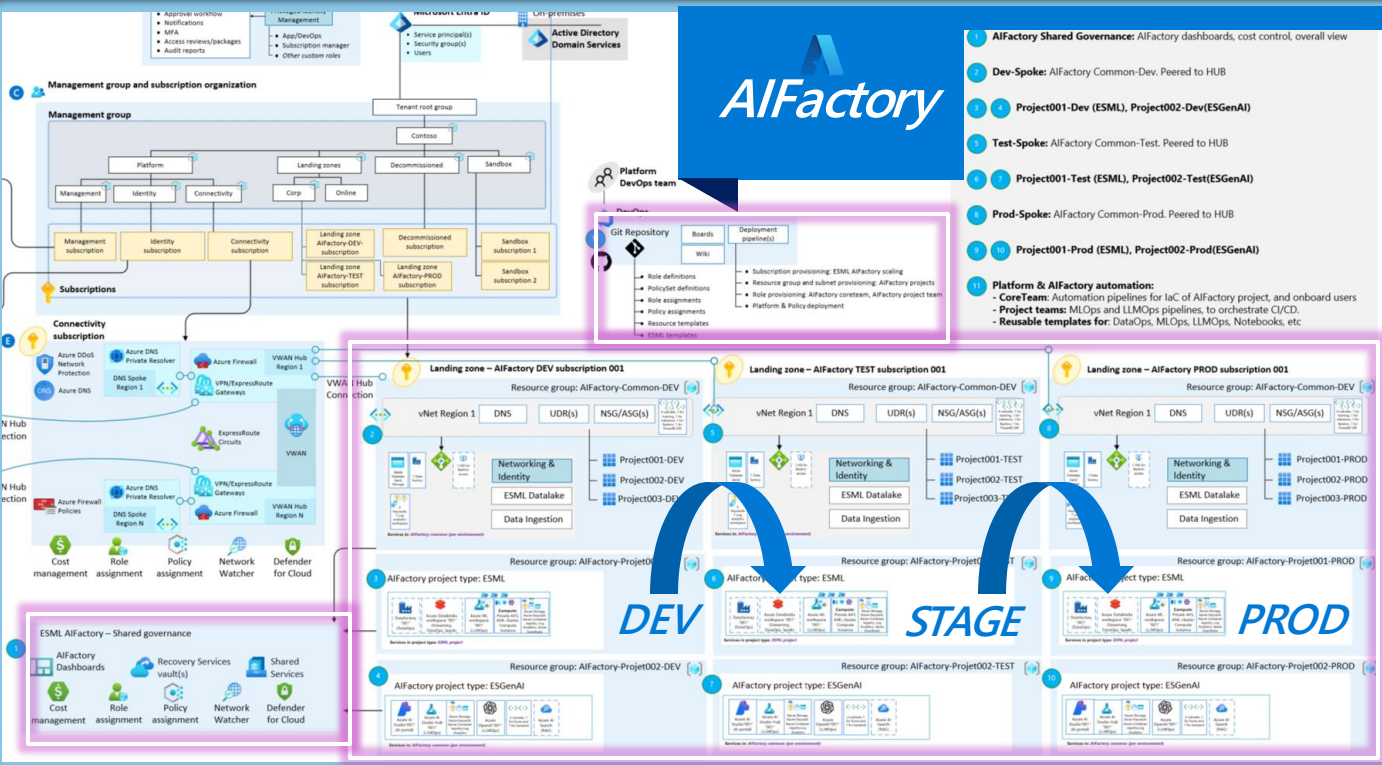


# AIOps – INNER & OUTER loop





# AIReady infra - Enterprise Scale AI (GenAIOps)

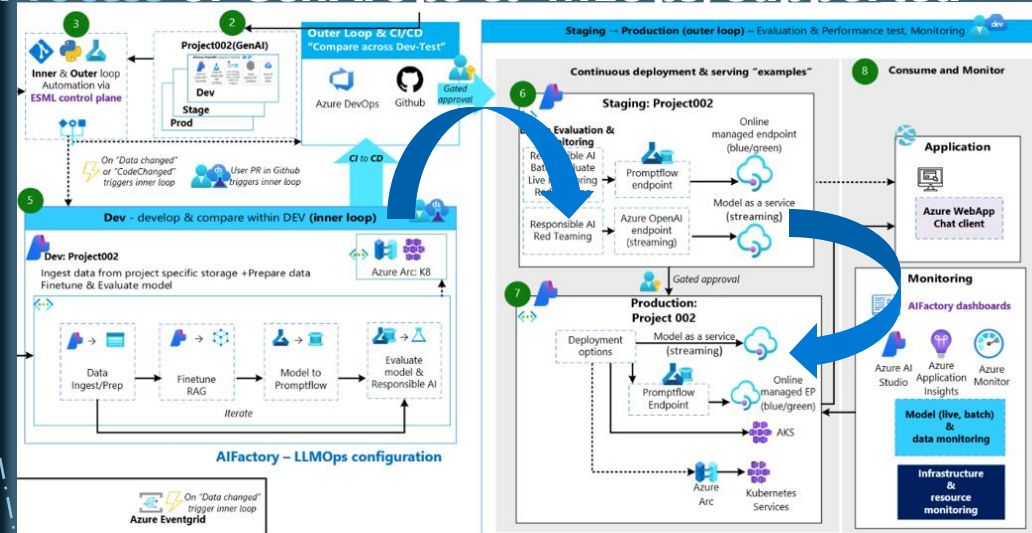


DEV STAGE PROD

## Process of GenAIOps / MLOps

- ✓ AI Ready Infrastructure – Dev, Stage, Production
- ✓ AIOps Templates: move from: Dev->Stage->Production

## Process of GenAIOps & MLOps, supported





# Processes: DataOps, MLOps, GenAIOps

Templates: DataOps, MLOps example

# DataOps



## Ingest

Core team  
Data ingestion team

transform indata to .parquet /  
.delta

# MLOps



Azure Machine Learning

## Azure ML - Train



Project member  
Data scientist, MLOps engineer

# MLOps

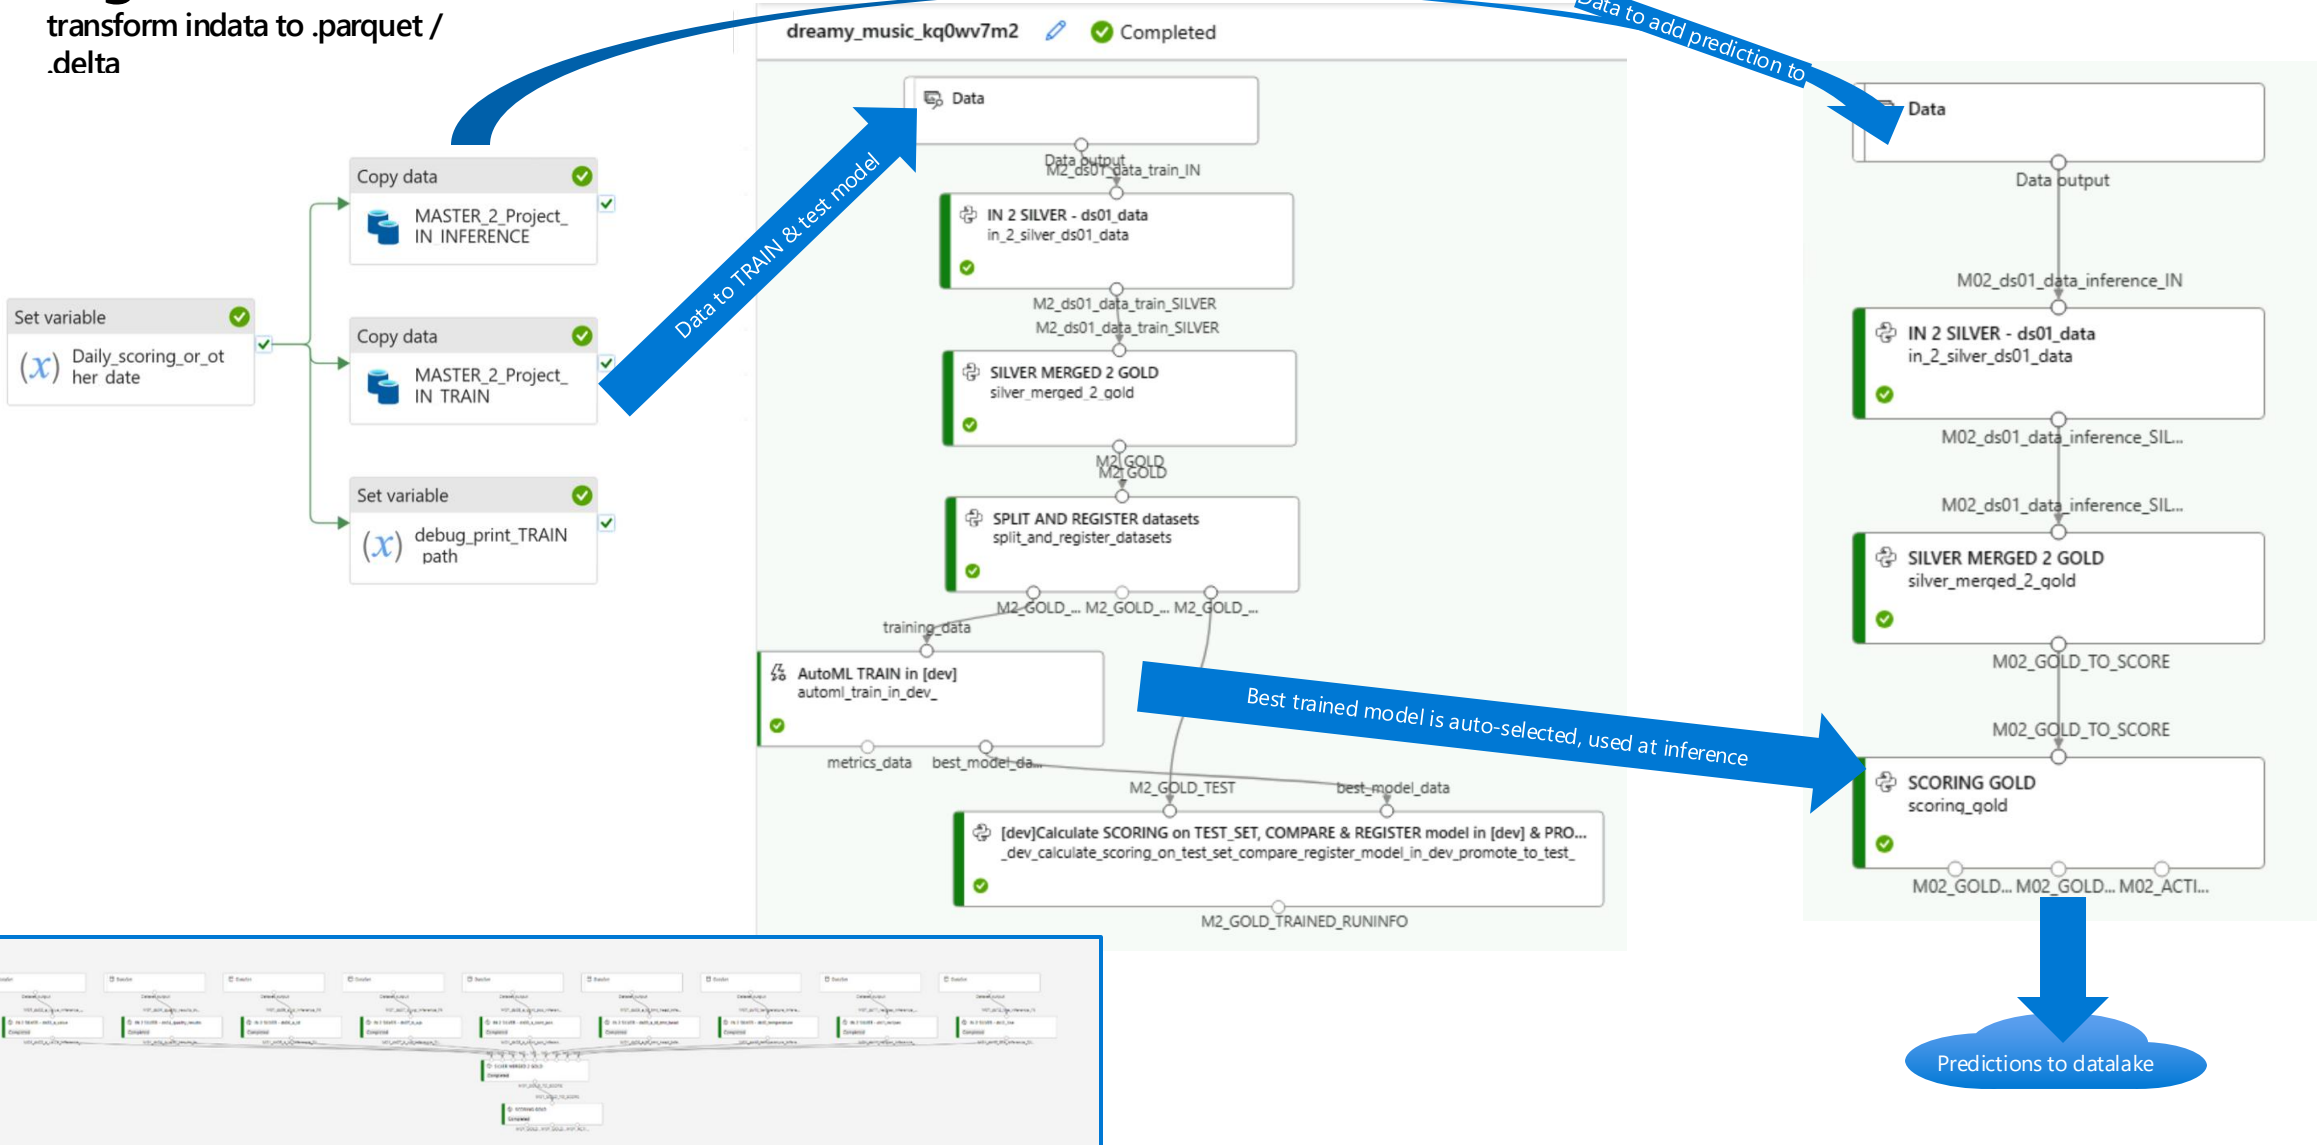


Azure Machine Learning

## Azure ML - Inference



Project member  
Data scientist, MLOps engineer



# AI Factory in context of CAF Enterprise Scale LZ

## Green Field

- 1) Setup ESLZ first (4-16h)
- 2) Setup AI Factory landingzones (4h)

## Brown Field

- 1) Setup AI Factory landingzones (4h)

## CAF ESLZ

### Implement Cloud Adoption Framework enterprise-scale landing zones in Azure

Example deployment	Description	GitHub repo	Deploy to Azure
Enterprise-scale foundation	The suggested foundation for enterprise-scale adoption.	<a href="#">Example in GitHub</a>	<a href="#">Deploy to Azure</a>
Enterprise-scale hub and spoke	Add a <a href="#">hub and spoke</a> network module to the enterprise-scale foundation.	<a href="#">Example in GitHub</a>	<a href="#">Deploy to Azure</a>
Enterprise-scale Virtual WAN	Add a <a href="#">Virtual WAN</a> network module to the enterprise-scale foundation.	<a href="#">Example in GitHub</a>	<a href="#">Deploy to Azure</a>
Enterprise-scale for small enterprises	Add a <a href="#">hub and spoke</a> network architecture for small organizations.	<a href="#">Example in GitHub</a>	<a href="#">Deploy to Azure</a>
Enterprise-scale for Azure Government	Reference implementation that can be deployed to Azure Government and includes all options in a converged portal experience.	<a href="#">Example in GitHub</a>	<a href="#">Deploy to Azure</a>

Configure AI Factory to map with resiliency configuration & policy's.  
(Default: it inherits from WAF/CAF)

# Cloud Adoption Framework - About AI Factory in 2020

## (DataOps + MLOps)

## AI factory for organization machine learning operations

A data science team might decide it can manage multiple machine learning use cases internally. Adopting machine learning operations helps an organization set up project teams for better quality, reliability, and maintainability of solutions. Through balanced teams, supported processes, and technology automation, a team that adopts machine learning operations can scale and focus on developing new use cases.

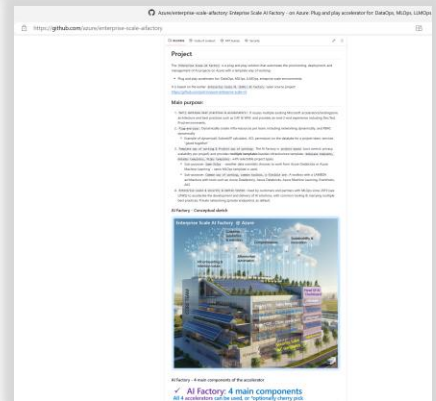
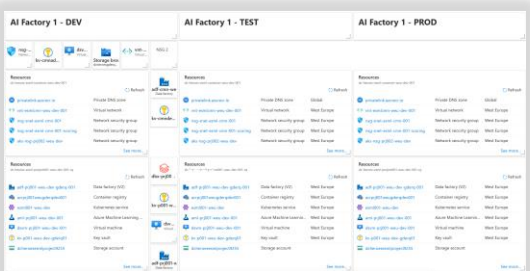
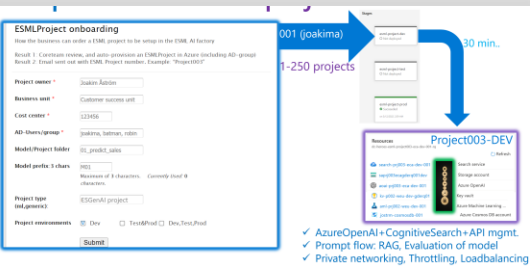
As the number of use cases grows in an organization, the management burden of supporting the use cases grows linearly, or even more. The challenge for the organization becomes how to accelerate time to market, support quicker assessment of use case feasibility, implement repeatability, and best use available resources and skill sets on a range of projects. For many organizations, developing an AI factory is the solution.

An AI factory is a system of repeatable business processes and standardized artifacts that facilitates developing and deploying a large set of machine learning use cases. An AI factory optimizes team setup, recommended practices, machine learning operations strategy, architectural patterns, and reusable templates that are tailored to business requirements.

Source: CAF AI Factory Documentation (2020) [Machine learning operations - Cloud Adoption Framework | Microsoft Learn](#)

GITHUB Repo (est. 2019, MIT): **Enterprise Scale AI Factory**: "Start small (4h setup) but secure, and scale with automation & templates"

<https://github.com/jostrm/azure-enterprise-scale-ml>  
<https://github.com/azure/enterprise-scale-aifactory>



# Documentation

## Documentation - Executive summary

### Quick Setup

1. [Prerequisites - to setup AIFactory](#) - Estimated setup time: 1-2h
2. [End-2-End SETUP tutorial - AIFactory + 1 ESMLProject](#) - Estimated setup time: 4-8h

### Governance related - relevant for central IT, networking team (CoreTeam: 10-29)

- [11\) Infra:AIFactory: Static documentation \(CoreTeam\)](#)
- [12\) Infra:AIFactory: Roles & Permissions for users \(CoreTeam\)](#)
- [13\) Infra:AIFactory: Flow diagrams \(CoreTeam\)](#)
- [14\) Infra:AIFactory: Networking \(CoreTeam\)](#)
- [15\) Infra:AIFactory: Overview of services: Naming convention\(CoreTeam\)](#)
- [21\) Infra:AIFactory: How-to: Onboarding of CoreTeam users and ProjectMembers via Pipelines \(CoreTeam\)](#)
- [22\) Datalake template: How-to: Setup Datalake & Onboard ProjectTeam permissions \(CoreTeam\)](#)
- [23\) DataOps template: How-to: Setup DataOps via PIPELINE templates \(CoreTeam\)](#)

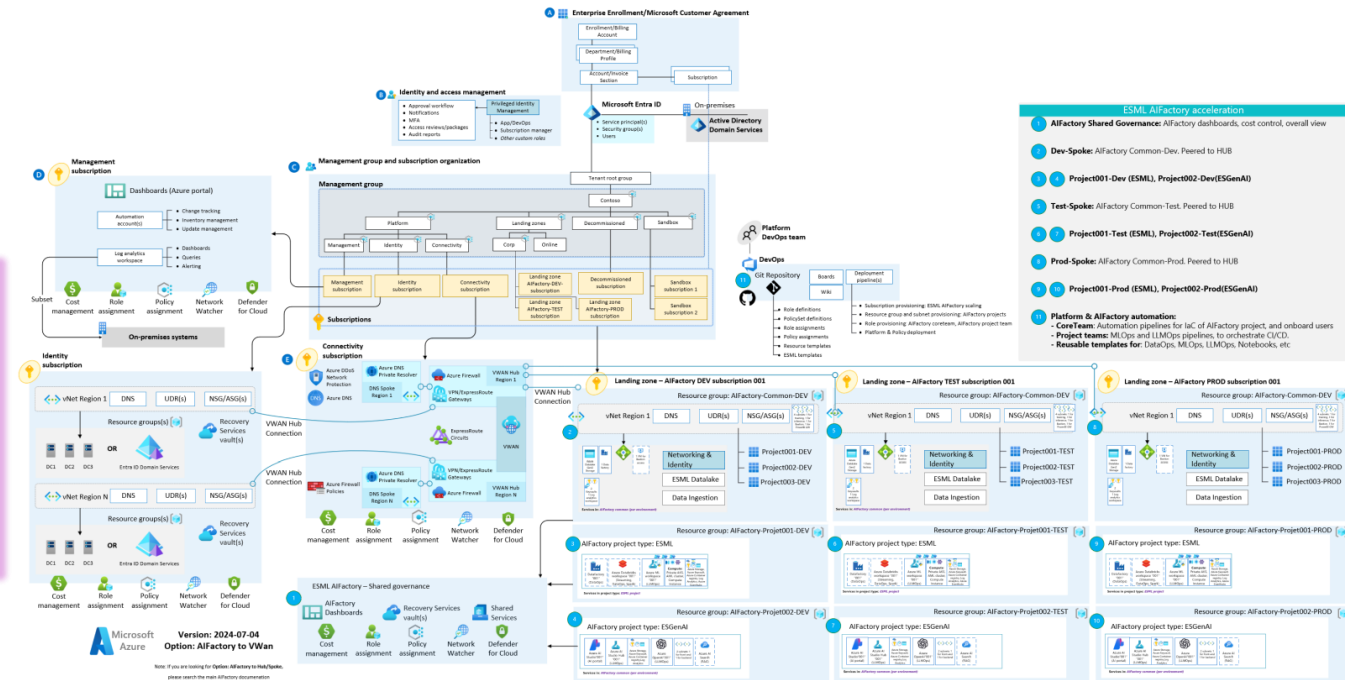
### Consumer related - relevant for developers, data scientists, data engineers (ProjectTeam: 30-39)

- [30\) Usage: Dashboard, Available Tools & Services, DataOps, MLOps, Access options to AIFactory \(ProjectTeam\)](#)
- [31\) How-to guide: Get access to the AIFactory:RBAC & Networking \(ProjectTeam\)](#)
- [32\) Overview: Dashboards, Services & Acceleration in AIFactory \(ProjectTeam, CoreTeam\)](#)
- [33\) Setup: Install AzureML SDK v1+v2 and ESML accelerator library](#)
- [34\) Setup: Datalake: How-to onboard your own data, to project folder R&D purpose](#)
- [35\) Setup: ESML SDK accelerated Notebook templates](#)
- [36\) How-to guide: DataOps](#)
- [37\) How-to guide: MLOps](#)
- [38\) How-to guide: LLMops & RAG Chat Agent](#)
- [39\) End-2-End config tutorial - ESML Project, ESGenAI Project](#) - Estimated config time: 1-2h

### FAQ - Trouble shooting - relevant for all

- [FAQ - Core team & AFactory infra](#)
- [FAQ - Data scientist & Azure ML pipelines](#)
- [FAQ - Data engineering & Azure ML pipelines](#)

## AIFactory's Enterprise Scale Landing Zones: Virtual Wan



# Documentation

## Documentation - Executive summary

### Quick Setup

1. [Prerequisites - to setup AIFactory](#) - Estimated setup time: 1-2h
2. [End-2-End SETUP tutorial - AIFactory + 1 ESMLProject](#) - Estimated setup time: 4-8h

### Governance related - relevant for central IT, networking team (CoreTeam: 10-29)

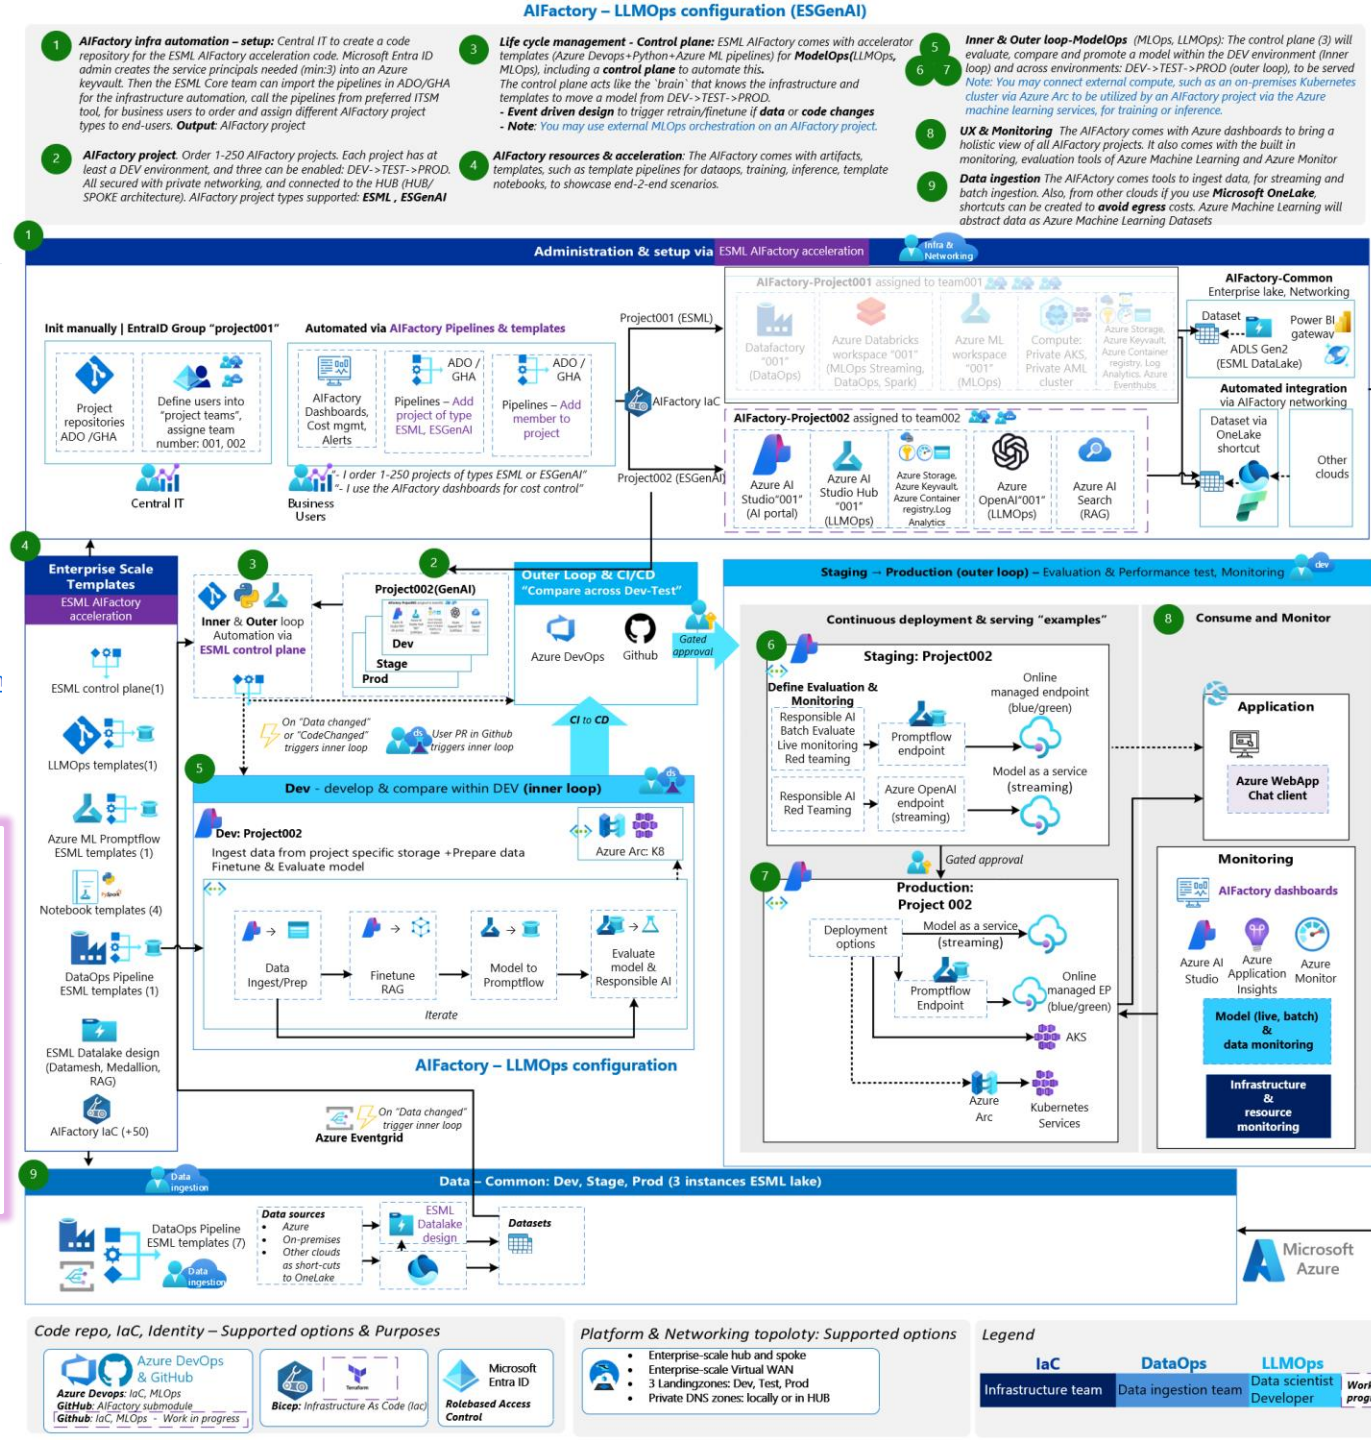
- [11\) Infra:AIFactory: Static documentation \(CoreTeam\)](#)
- [12\) Infra:AIFactory: Roles & Permissions for users \(CoreTeam\)](#)
- [13\) Infra:AIFactory: Flow diagrams \(CoreTeam\)](#)
- [14\) Infra:AIFactory: Networking \(CoreTeam\)](#)
- [15\) Infra:AIFactory: Overview of services: Naming convention\(CoreTeam\)](#)
- [21\) Infra:AIFactory: How-to: Onboarding of CoreTeam users and ProjectMembers via Pipelines \(CoreTeam\)](#)
- [22\) Datalake template: How-to: Setup Datalake & Onboard ProjectTeam permissions \(CoreTeam\)](#)
- [23\) DataOps template: How-to: Setup DataOps via PIPELINE templates \(CoreTeam\)](#)

### Consumer related - relevant for developers, data scientists, data engineers (ProjectTeam: 30-39)

- [30\) Usage: Dashboard, Available Tools & Services, DataOps, MLOps, Access options to AIFactory \(ProjectTeam\)](#)
- [31\) How-to guide: Get access to the AIFactory:RBAC & Networking \(ProjectTeam\)](#)
- [32\) Overview: Dashboards, Services & Acceleration in AIFactory \(ProjectTeam, CoreTeam\)](#)
- [33\) Setup: Install AzureML SDK v1+v2 and ESML accelerator library](#)
- [34\) Setup: Datalake: How-to onboard your own data, to project folder R&D purpose](#)
- [35\) Setup: ESML SDK accelerated Notebook templates](#)
- [36\) How-to guide: DataOps](#)
- [37\) How-to guide: MLOps](#)
- [38\) How-to guide: LLMOps & RAG Chat Agent](#)
- [39\) End-2-End config tutorial - ESML Project, ESGenAI Project](#) - Estimated config time: 1-2h

### FAQ - Trouble shooting - relevant for all

- [FAQ - Core team & AFactory infra](#)
- [FAQ - Data scientist & Azure ML pipelines](#)
- [FAQ - Data engineering & Azure ML pipelines](#)



# Documentation

## Documentation - Executive summary

### Quick Setup

1. [Prerequisites - to setup AIFactory](#) - Estimated setup time: 1-2h
2. [End-2-End SETUP tutorial - AIFactory + 1 ESMLProject](#) - Estimated setup time: 4-8h

### Governance related - relevant for central IT, networking team (CoreTeam: 10-29)

- [11\) Infra:AIFactory: Static documentation \(CoreTeam\)](#)
- [12\) Infra:AIFactory: Roles & Permissions for users \(CoreTeam\)](#)
- [13\) Infra:AIFactory: Flow diagrams \(CoreTeam\)](#)
- [14\) Infra:AIFactory: Networking \(CoreTeam\)](#)
- [15\) Infra:AIFactory: Overview of services: Naming convention\(CoreTeam\)](#)
- [21\) Infra:AIFactory: How-to: Onboarding of CoreTeam users and ProjectMembers via Pipelines \(CoreTeam\)](#)
- [22\) Datalake template: How-to: Setup Datalake & Onboard ProjectTeam permissions \(CoreTeam\)](#)
- [23\) DataOps template: How-to: Setup DataOps via PIPELINE templates \(CoreTeam\)](#)

### Consumer related - relevant for developers, data scientists, data engineers (ProjectTeam: 30-39)

- [30\) Usage: Dashboard, Available Tools & Services, DataOps, MLOps, Access options to AIFactory \(ProjectTeam\)](#)
- [31\) How-to guide: Get access to the AIFactory:RBAC & Networking \(ProjectTeam\)](#)
- [32\) Overview: Dashboards, Services & Acceleration in AIFactory \(ProjectTeam, CoreTeam\)](#)
- [33\) Setup: Install AzureML SDK v1+v2 and ESML accelerator library](#)
- [34\) Setup: Datalake: How-to onboard your own data, to project folder R&D purpose](#)
- [35\) Setup: ESML SDK accelerated Notebook templates](#)
- [36\) How-to guide: DataOps](#)
- [37\) How-to guide: MLOps](#)
- [38\) How-to guide: LLMOps & RAG Chat Agent](#)
- [39\) End-2-End config tutorial - ESML Project, ESGenAI Project](#) - Estimated config time: 1-2h

### FAQ - Troubleshooting - relevant for all

- [FAQ - Core team & AFactory infra](#)
- [FAQ - Data scientist & Azure ML pipelines](#)
- [FAQ - Data engineering & Azure ML pipelines](#)

## Infra:AIFactory : Flow diagrams - Add AIFactory project, Add users (CoreTeam)

### How to get Object ID of user?

- Option A) If read access in Microsoft EntraID you can see your own user. Administrators can see all users.
- Option B) Contact the administrator in your IT department, and ask for ObjectID's for users in Microsoft EntraID. Provide the admin a list of email addresses of users, to use for lookup of the ObjectID.
- Option C) Powershell. Get it by using `Get-AzADUser` or `Get-AzADServicePrincipal` cmdlet

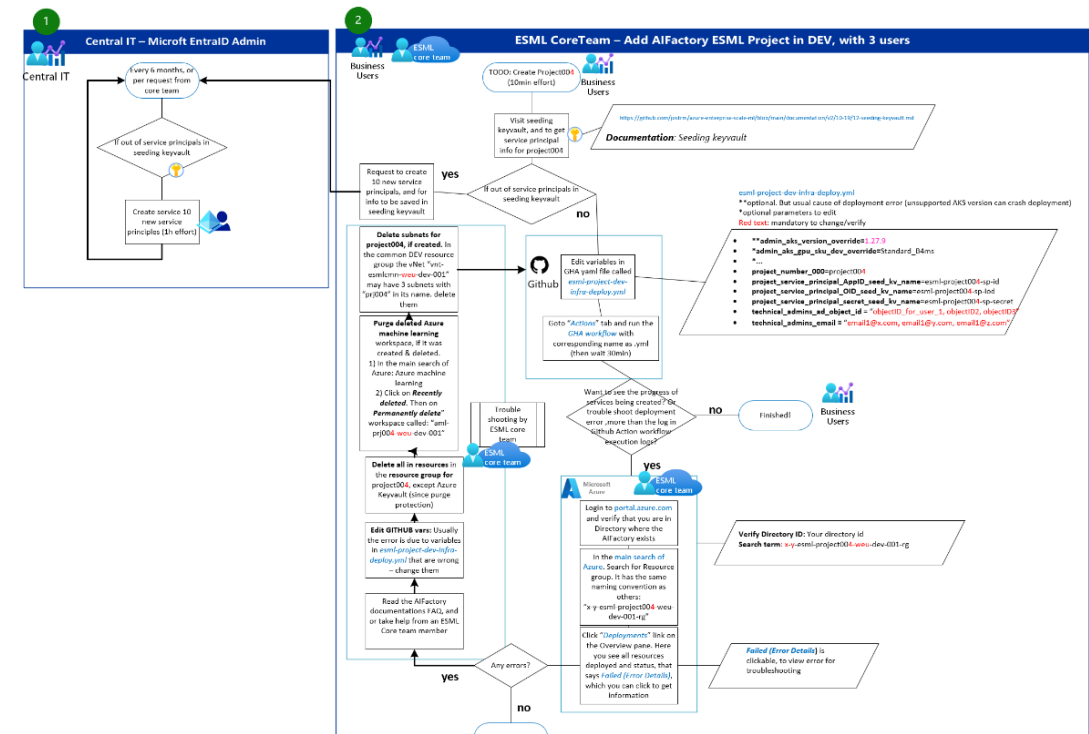
### Option C) Get user by User Principal Name (UPN)

```
$user = Get-AzADUser -UserPrincipalName "user@example.com"
$objectId = $user.Id
Write-Output $objectId
```

### Option C) Get service principal by Display Name

```
$servicePrincipal = Get-AzADServicePrincipal -DisplayName "MyServicePrincipal"
$objectId = $servicePrincipal.Id
Write-Output $objectId
```

## Flow: Add AIFactory ESML Project in DEV, with 3 users



Ops,

es. Also how

Index
<a href="#">31,32</a>
<a href="#">34</a>
36,37
<a href="#">33,35</a>

am)

# Documentation

## Documentation - Executive summary

### Quick Setup

1. [Prerequisites - to setup AIFactory](#) - Estimated setup time: 1-2h
2. [End-2-End SETUP tutorial - AIFactory + 1 ESMLProject](#) - Estimated setup time: 4-8h

### Governance related - relevant for central IT, networking team (CoreTeam: 10-29)

- [11\) Infra:AIFactory: Static documentation \(CoreTeam\)](#)
- [12\) Infra:AIFactory: Roles & Permissions for users \(CoreTeam\)](#)
- [13\) Infra:AIFactory: Flow diagrams \(CoreTeam\)](#)
- [14\) Infra:AIFactory: Networking \(CoreTeam\)](#)
- [15\) Infra:AIFactory: Overview of services: Naming convention\(CoreTeam\)](#)
- [21\) Infra:AIFactory: How-to: Onboarding of CoreTeam users and ProjectMembers via Pipelines \(CoreTeam\)](#)
- [22\) Datalake template: How-to: Setup Datalake & Onboard ProjectTeam permissions \(CoreTeam\)](#)
- [23\) DataOps template: How-to: Setup DataOps via PIPELINE templates \(CoreTeam\)](#)

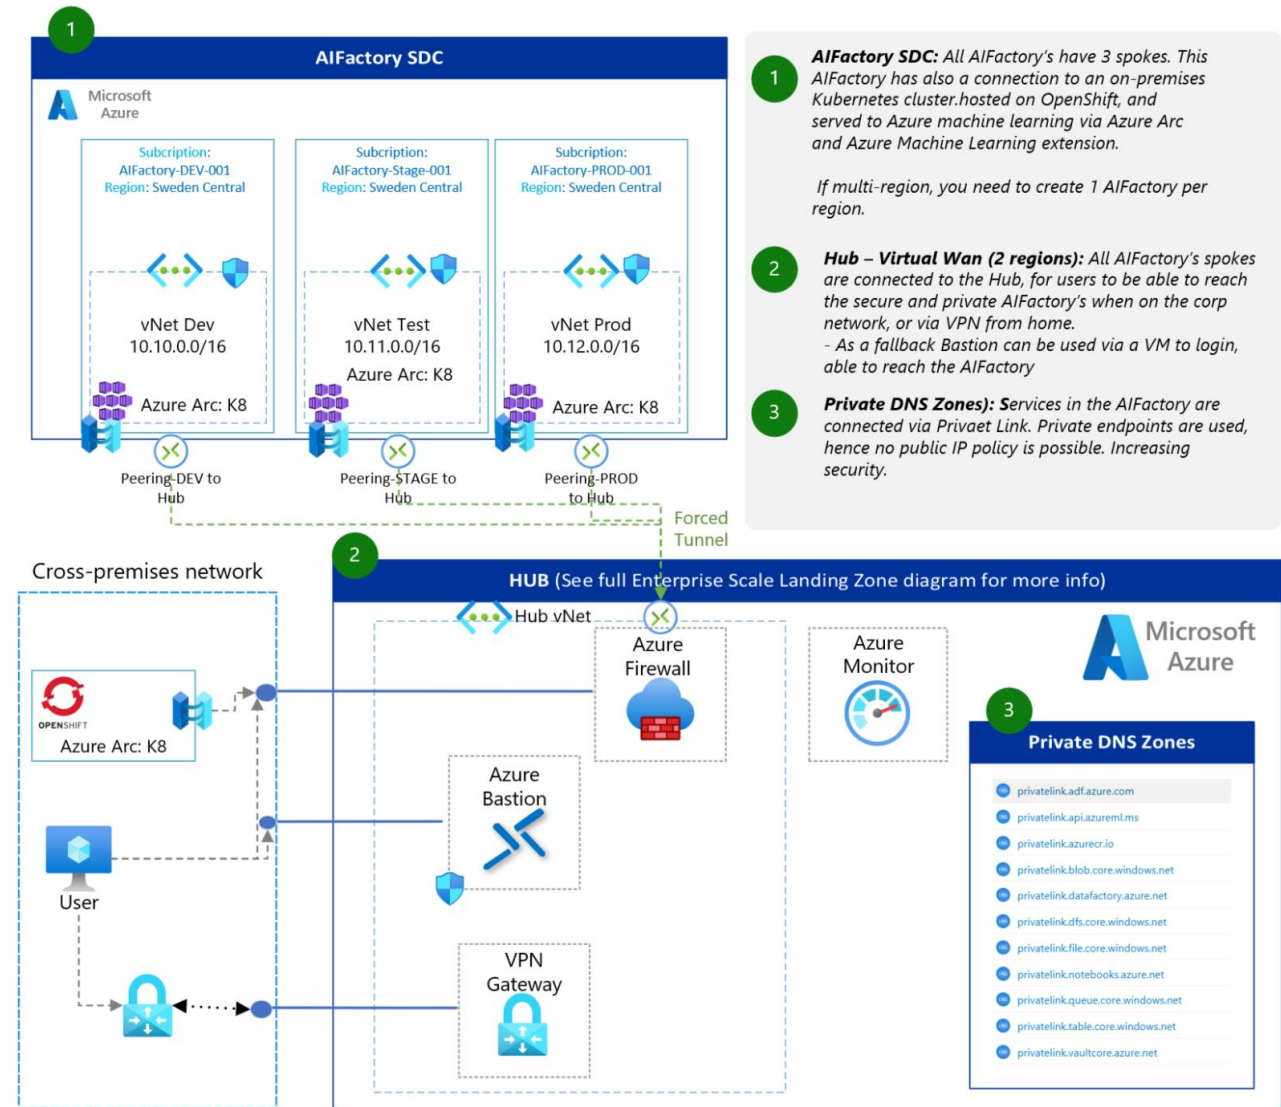
### Consumer related - relevant for developers, data scientists, data engineers (ProjectTeam: 30-39)

- [30\) Usage: Dashboard, Available Tools & Services, DataOps, MLOps, Access options to AIFactory \(ProjectTeam\)](#)
- [31\) How-to guide: Get access to the AIFactory:RBAC & Networking \(ProjectTeam\)](#)
- [32\) Overview: Dashboards, Services & Acceleration in AIFactory \(ProjectTeam, CoreTeam\)](#)
- [33\) Setup: Install AzureML SDK v1+v2 and ESML accelerator library](#)
- [34\) Setup: Datalake: How-to onboard your own data, to project folder R&D purpose](#)
- [35\) Setup: ESML SDK accelerated Notebook templates](#)
- [36\) How-to guide: DataOps](#)
- [37\) How-to guide: MLOps](#)
- [38\) How-to guide: LLMOps & RAG Chat Agent](#)
- [39\) End-2-End config tutorial - ESML Project, ESGenAI Project](#) - Estimated config time: 1-2h

### FAQ - Trouble shooting - relevant for all

- [FAQ - Core team & AFactory infra](#)
- [FAQ - Data scientist & Azure ML pipelines](#)
- [FAQ - Data engineering & Azure ML pipelines](#)

## AIFactory Hub-Spoke - Light & Simplified view



- Simplified view
- Not all Private DNS zones are listed. [To see all - go here](#)
- Not all networking components are shown. [To see more - go here \(this doc\)](#) and [go here \(this doc\)](#)