For financial services and enterprises

Generative AI Implementation on Amazon Bedrock



Secure in the Seoul Region, from assessment to operations

Generative AI on Amazon Bedrock Implementation by MetanetX

November 2026

MetanetX

AT A GLANCE

One partner for generative AI, from assessment to operations

WHAT

Implementation service on Amazon Bedrock in the Seoul Region

We take business-specific generative AI from assessment and PoC through build and operations

WHO

Financial services, regulated industries, and domain-specific AI

Organizations with network separation or in-country processing needs, and teams automating content and workflows

HOW

Three paths and four packages

Choose secure enterprise AI, business-specific AI applications, or landing-zone-based secure AI, and start at the stage you need

VALUE

Security review readiness, operable systems, transparent costs

We keep invocation logs and cost records, and provide MSP operations as well

WHY NOW

The models are ready. Operations is what remains

SEOUL REGION



2026.9

In-Region inference for Claude now available in the Seoul Region

Processed in-Region

Requests and data sent to Seoul Region endpoints are processed within the Region

Three barriers that keep PoCs from reaching production, and how we solve them

1 **Security review**

Hard to know which design will pass

→ **Private calls, multi-layer masking, and audit logs by design**

Page 8

2 **Cost predictability**

Costs vary widely by model and call volume

→ **Cost records and up-front cost estimates**

Page 13

3 **Operations staffing**

No team to run the system and update models

→ **MSP operations, model updates, monthly reports**

Page 13

OFFERING

Three paths and four packages

Path A



Secure enterprise AI environment

For	Finance, public sector, network-separated firms
Use cases	Document Q&A, summarization, review support
Core	Private network, AI App, LLM Gateway

Path B



Business-specific AI applications

For	Media, content, and workflow automation
Use cases	Content generation, planning agents, analysis reports
Core	AI App, agent loop, model roles by task

Path C



Landing-zone-based secure AI

For	Financial firms with the strictest security needs
Use cases	Enterprise-wide AI, customer-facing services
Core	Separate Network, Security, and AI accounts

All three paths follow the same four packages, and you can start at the stage you need

01
GenAI Assessment
2 weeks

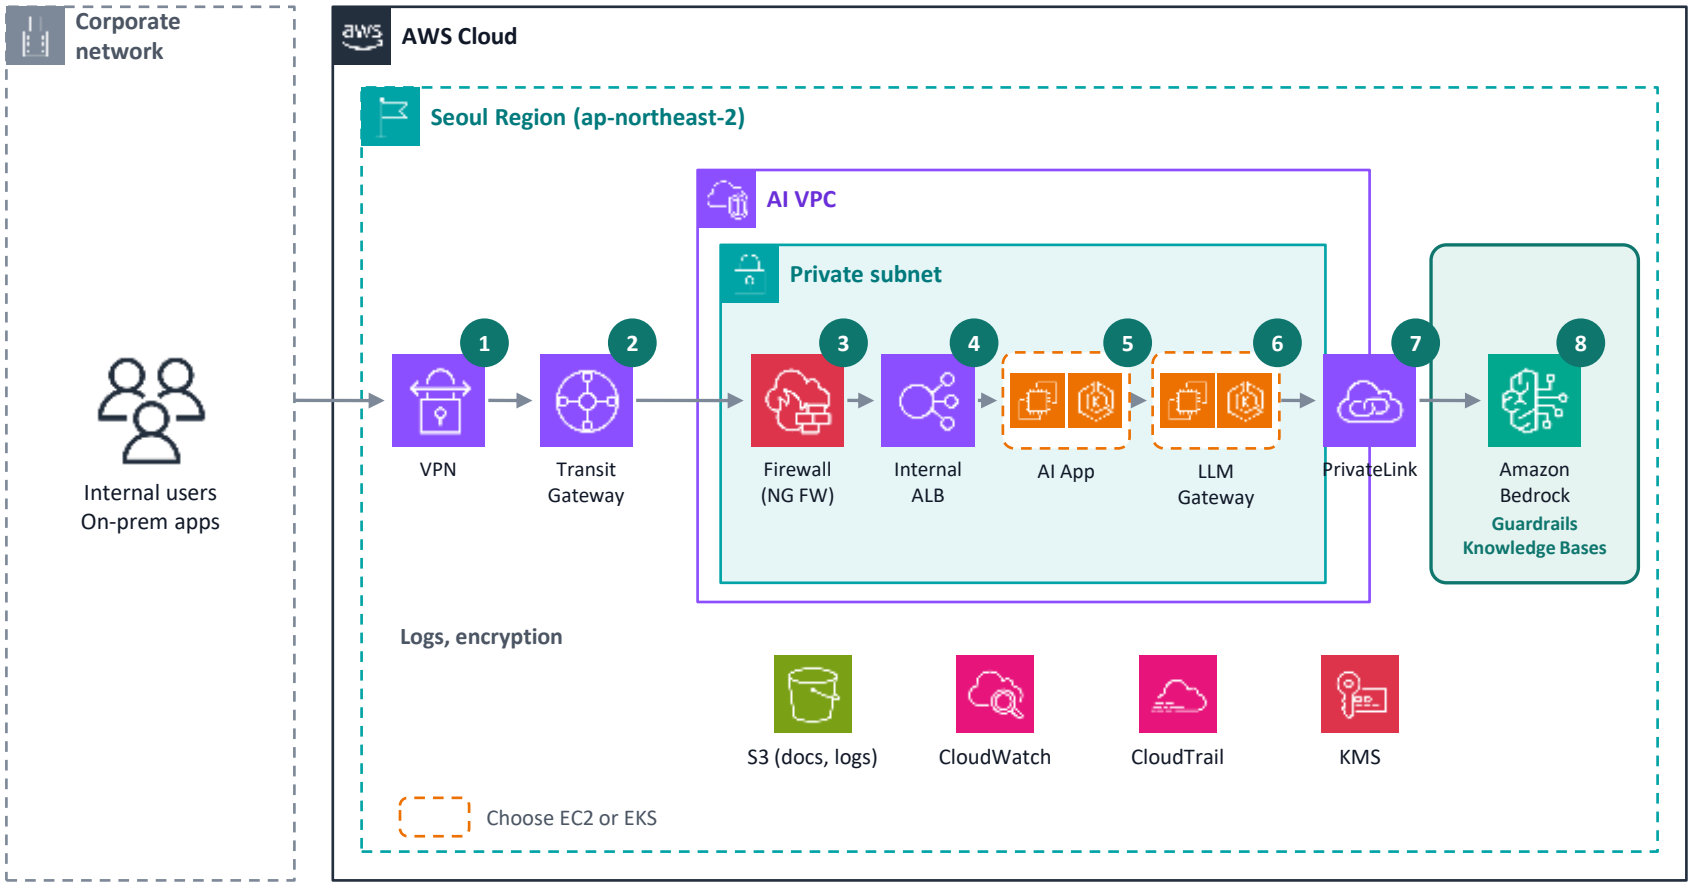
02
GenAI PoC
6–8 weeks

03
GenAI Production Build
8–12 weeks

04
GenAI Managed Operations
Monthly

ARCHITECTURE A

Path A, a secure AI environment reachable only from your network



Private calls only

Bedrock is called only over IPSec VPN or Direct Connect and PrivateLink

Separate control layer

The AI App handles code, prompts, and documents, and the LLM Gateway handles masking

Multi-layer masking

The AI App, LLM Gateway, and Guardrails filter sensitive data step by step

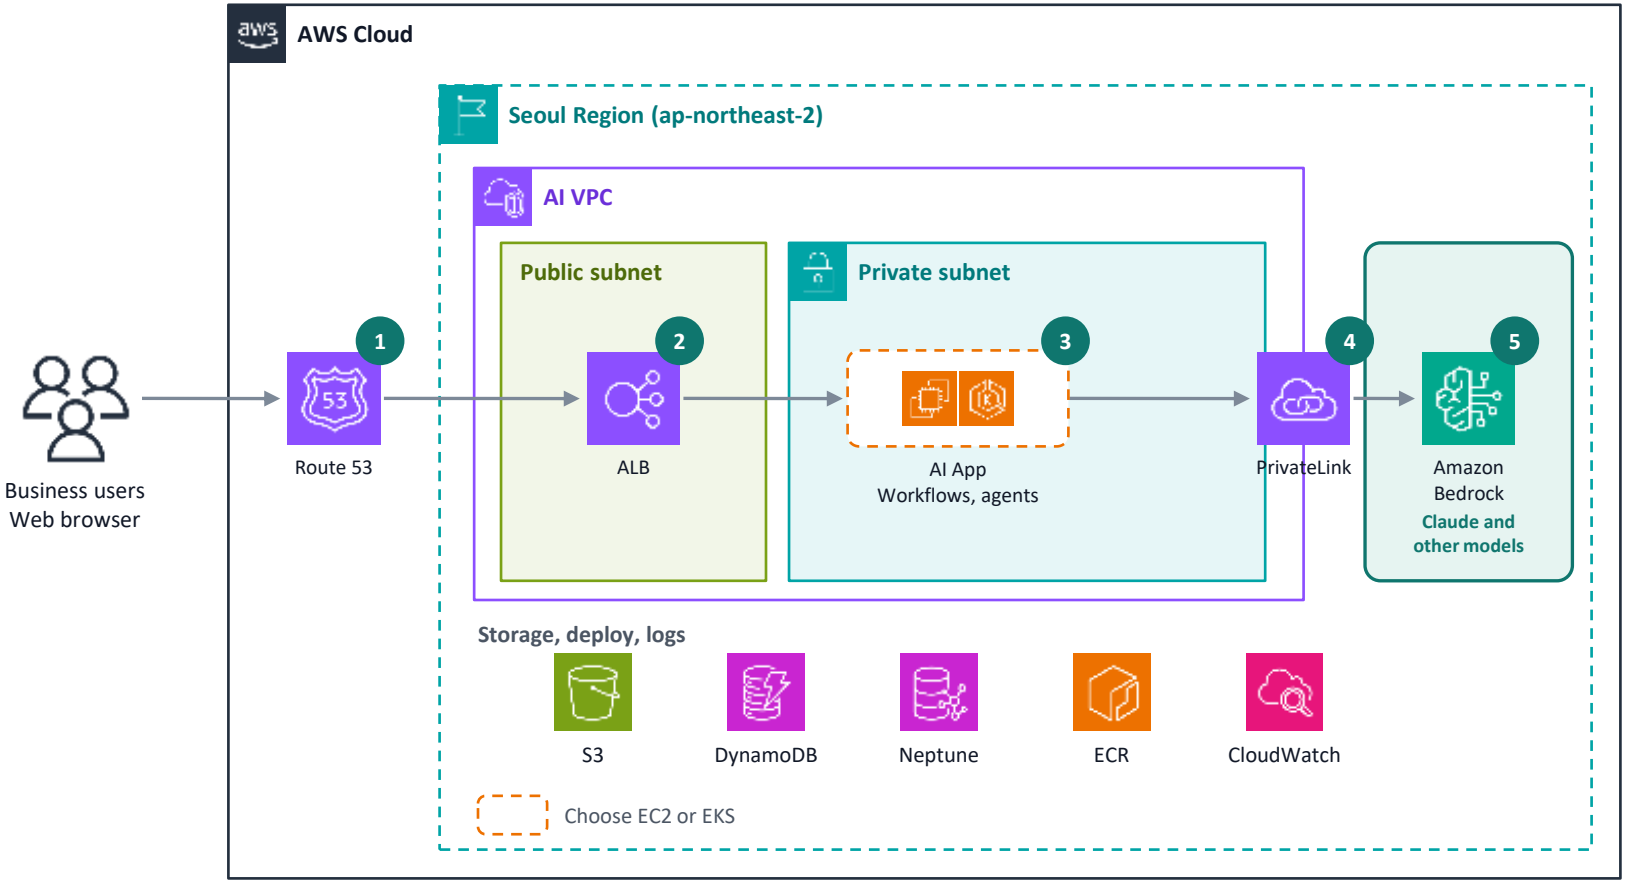
Audit logs

Every call is recorded in CloudWatch, CloudTrail, and KMS-encrypted S3

Call flow: ① Private connection → ② Transit Gateway → ③ Firewall → ④ Internal ALB → ⑤ AI App → ⑥ LLM Gateway → ⑦ PrivateLink → ⑧ Bedrock

ARCHITECTURE B

Path B, AI applications that mirror your business workflow



Human-in-the-loop

Review results at every step and regenerate only what needs to change

Model roles by task

The right model for each task, such as reading, summarizing, and reviewing

Background jobs

Long-running jobs finish on the server and results are saved

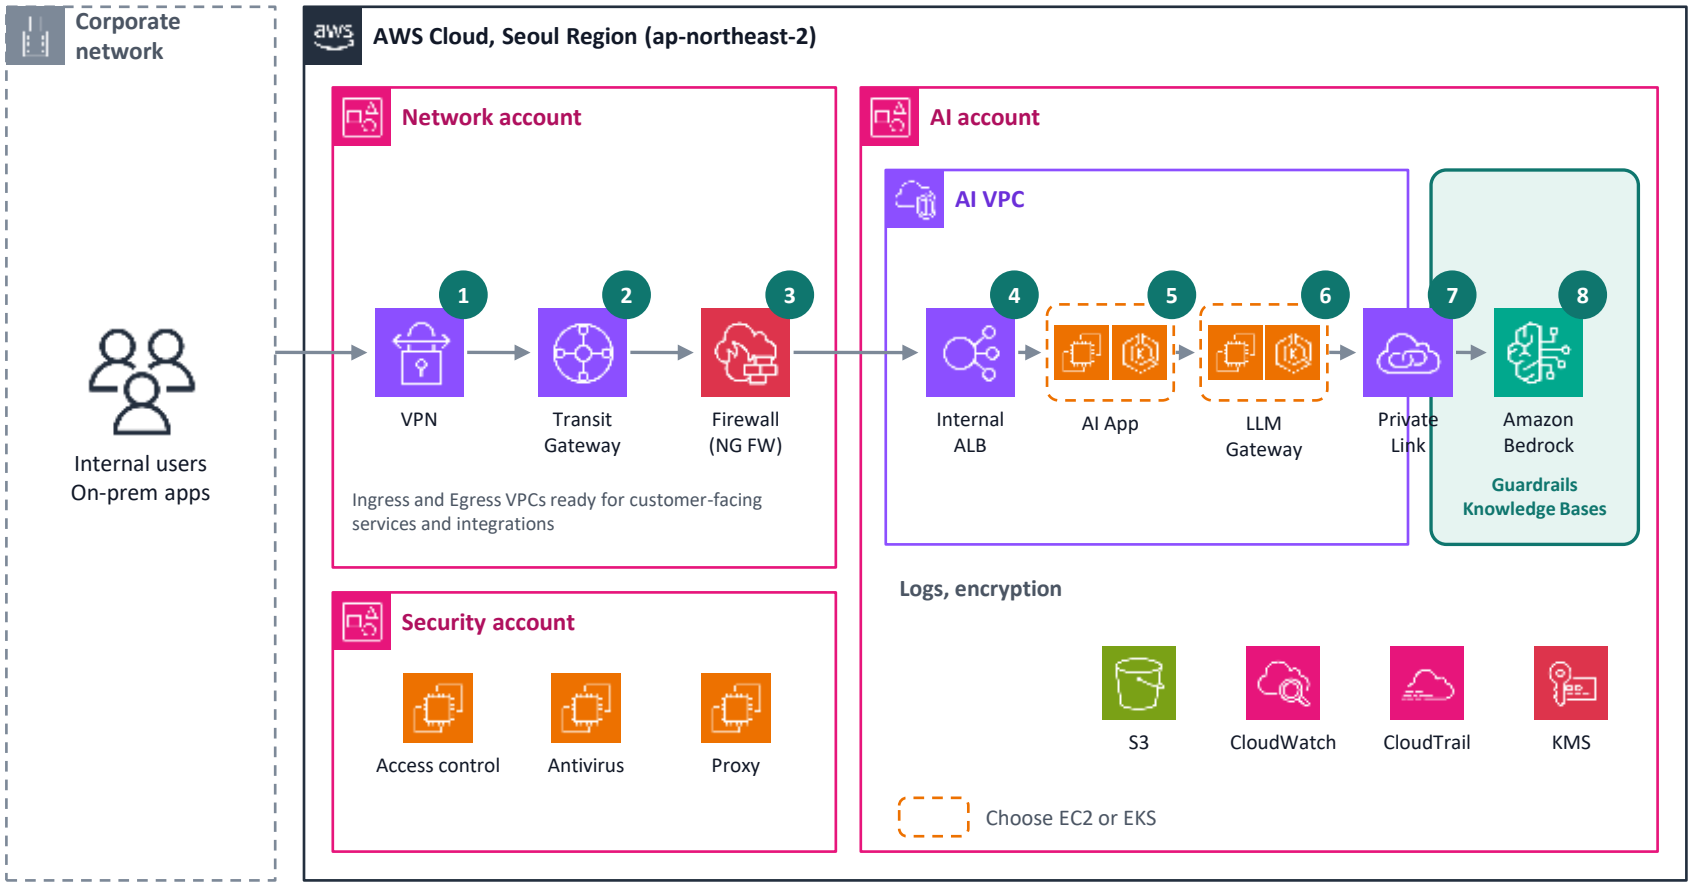
Compute choice

Run the AI App on EKS or EC2 to fit your scale and operating model

Call flow: ① Route 53 → ② ALB → ③ AI App → ④ PrivateLink → ⑤ Bedrock Deployment: GitLab CI, ECR, ArgoCD

ARCHITECTURE C

Path C, financial-grade secure AI wrapped in a landing zone



Account separation

Separate Network, Security, and AI accounts isolate permissions and failures

Forced firewall path

Transit Gateway routing sends every request through the firewall

Separate control layer

The AI App runs business logic, and the LLM Gateway handles masking and usage control

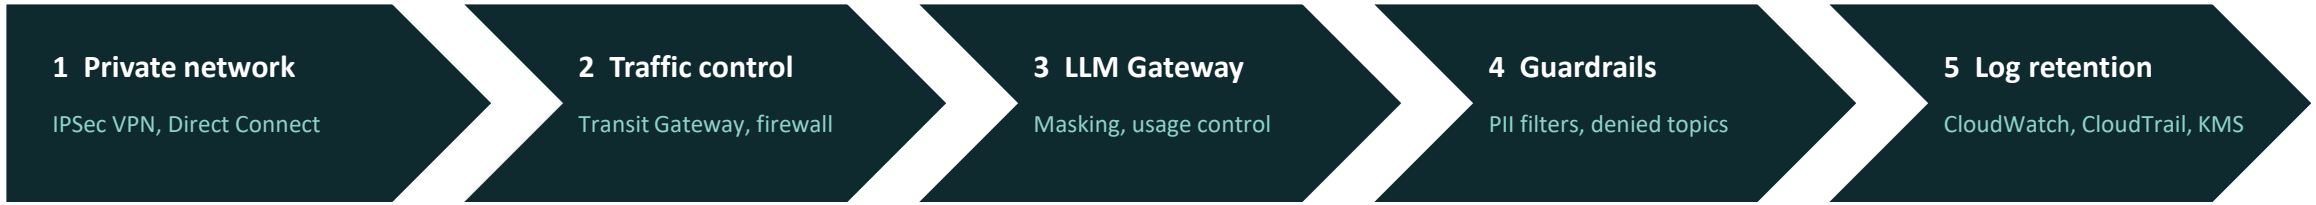
Integrated security tools

Access control, antivirus, and proxy run centrally in the Security account

Call flow: ① Private connection → ② Transit Gateway → ③ Firewall → ④ Internal ALB → ⑤ AI App → ⑥ LLM Gateway → ⑦ PrivateLink → ⑧ Bedrock

SECURITY

Security requirements built in from the design stage



Requirement	How we address it
In-country processing	Only Seoul Region endpoints are called, and cross-Region inference is blocked with IAM
Network separation	Models are called only over private connections separated from production (IPSec VPN or Direct Connect, PrivateLink)
Personal and credit data	Pseudonymized or sample data is used, and LLM Gateway masking and Guardrails block sensitive patterns
Access control	Developer, operator, and auditor roles are separated in IAM with MFA enforced
Audit and logging	Model invocation logs and CloudTrail are encrypted with KMS and retained
AI accountability	Answers show their sources, no source means no answer, and staff make the final decision
Account separation	Path C separates Network, Security, and AI accounts and routes every request through the firewall

Security review package included: configuration checklist, architecture, data flow diagrams, log samples

WHY METANETX

Fast delivery with methods proven in real projects



Standard design

- Single Seoul Region setup
- Single Bedrock call path
- Proven VPC and EKS layout



Human review

- Review at every step
- Regenerate only what changed
- Versioning and rollback



Model operations

- Model roles by task
- Fallback models for uninterrupted calls
- Cross-checks with another model
- Prompt caching to cut cost



CI/CD automation

- GitLab CI, ECR, ArgoCD
- Linked to MSP operations



Patterns repeated across delivered projects are our standard starting point

Single Seoul Region setup, single Bedrock call path, human review, call logs and cost records, automated deployment

CASES

Building generative AI in content and finance



Media

Webtoon to video

7 steps

Reads webtoon images and produces a 30-second video via script, storyboard, and previz



Media

Story development

8 steps

Writes from logline to script with AI, with every result under version control



Media

Planning agent

9 tools

The agent chooses web search, 15 writing skills, and document tools on its own



Media

Webtoon analysis

4 models

Splits episode reading, summaries, season analysis, and cross-review across models



Securities firm

Cloud build

Multi-account

EKS in a multi-account setup calls Bedrock in a separate account over a private path

SHOWCASE

Webtoon analysis service: how AI works on the real screen





1 Episode reading and summary tables

Upload webtoon images in episode order, and AI reads scenes and dialogue to build a table for each episode. Edits carry over to the next episode

2 Season analysis with independent review

One model writes the season structure and genre essentials, and another model re-checks the evidence and adds supporting or opposing views

PACKAGES

Start at the stage you need

01 GenAI Assessment

2 weeks

Key activities

- Use-case discovery workshop
- Use-case prioritization
- Security and regulatory needs
- Model and Region check
- Cost estimate

Deliverables

Use-case selection, target architecture, draft KPIs

02 GenAI PoC

6–8 weeks

Key activities

- Dev environment (VPC, PrivateLink, IAM, logging)
- 1–2 use cases built
- Model comparison with 100–200 test questions

Deliverables

PoC app, evaluation report, production recommendation

03 GenAI Production Build

8–12 weeks (scoped per project)

Key activities

- Production architecture (EKS or EC2)
- CI/CD, Guardrails, cost records
- Security review support
- User training

Deliverables

Production system, architecture guide, security review package

04 GenAI Managed Operations

Monthly

Key activities

- Monitoring and incident response
- Model version updates
- Cost reports
- Prompt improvement

Deliverables

Monthly operations report

PoC targets (draft)

Agreed per project

85%+	Answer accuracy
95%+	Source citation rate
30%+	Time saved

COST AND OPERATIONS

Transparent costs and end-to-end operations

Cost

Pay per use	Bedrock charges only for what you use
Cost records	Every paid call is logged and can be reviewed anytime
No double charges	Interrupted jobs are not charged again for the same call
Model by task	Pick the right model for each task to control cost

See the estimated cost on screen before you run

Work starts only after the user confirms the amount

Operations

MSP monitoring	Detects anomalies and responds to incidents
Model updates	New models are validated before switching, with fallback models on failure
Monthly report	Cost, usage, and prompt improvements shared every month
CI/CD automation	GitLab CI and ArgoCD apply changes safely

Delivered monthly through the GenAI Managed Operations package

ABOUT METANETX

An MSP partner running cloud and AI together

2007	KRW 538.8B	800+	600+
Founded	Revenue (FY2024, consolidated)	Employees	IT professionals



24x7 Managed Service

Dedicated engineers handle operations, monitoring, and incident response around the clock, every day of the year



Cloud SOC

Cloud security monitoring as a security monitoring provider designated by the Korean Ministry of Science and ICT



ISO/IEC 27001:2022

Certified information security management system (valid until May 2028)



Data and AI capabilities

We build data and AI services together with our subsidiary specializing in LLM-based conversational AI

HOW TO BUY

Buy through a Private Offer on AWS Marketplace

STEP 1

Review listing

Check this product's packages and scope on AWS Marketplace

STEP 2

Request offer

Request a Private Offer from the product page

STEP 3

Agree on terms

Agree on use cases, schedule, and price, and set the SOW and contract terms

STEP 4

Accept offer

Accept the Private Offer from MetanetX in your AWS account

STEP 5

Kick off

Once the contract is final, start with Assessment or PoC

Pay on your AWS bill

Charges appear on your existing AWS invoice

Flexible payment

Choose upfront, installment, or milestone payments in the Private Offer

Custom terms

Attach contract documents such as the SOW and price sheet to the offer

NEXT STEP

Start with a GenAI Assessment

1

Discovery workshop

Collect candidate use cases and set priorities

2

PoC scope and KPIs

Agree on what to validate and how to measure success

3

Production plan

Build an execution plan covering security, cost, and operations

Let's design your generative AI adoption around security, speed, cost, and operations.

mx.awssalest@metanetx.com

MetanetX