



CLOUD-ISOLATED

PRIVATE CLOUD

FULLY ON-PREMISE

PRIVATE & ON-PREMISE ENTERPRISE AI

Private / On-Premise Enterprise AI

Secure AI inside your organization — your data, your infrastructure, your control.

01

Private LLMs & AI Assistants

Generative AI deployed within approved infrastructure — no dependency on public AI services for sensitive workloads.

02

RAG & Document Intelligence

Policies, contracts, manuals and reports become searchable knowledge with source citations and traceability.

03

Enterprise AI Agents

HR · Legal · Procurement · Finance · IT
Service Desk · Executive support, under
role-based access control.

INSIDE THIS BROCHURE

02 Overview & Why Private / On-Premise AI**03** Enterprise ChatGPT · Document AI · RAG · Arabic**04** Enterprise AI Agents & Government Use Cases**05** Private AI Architecture · Security & Governance**06** AI Models & Public vs Private Comparison**07** Integrations · Deployment · AI Server Infrastructure**08** Implementation Journey & Example Case Study**08** Next Steps: Workshop, Demo & Assessment

Designed for government, regulated and enterprise environments across the UAE — with Arabic and English capability, enterprise integration and predictable infrastructure architecture.

ON-PREMISE AI UAE

PRIVATE AI UAE

ENTERPRISE AI UAE

PRIVATE CHATGPT UAE

ON-PREMISE LLM UAE

GOVERNMENT AI SOLUTIONS UAE

ARABIC ENTERPRISE AI

AI DATA SOVEREIGNTY UAE

SECTION 01 — SOLUTION OVERVIEW

Generative AI That Never Leaves Your Organization

Swedish Technology designs, deploys and supports **private enterprise AI** — engineered for government, regulated and enterprise environments in Abu Dhabi and Dubai.

PRIVATE ENTERPRISE AI OVERVIEW

Your Data, Your Infrastructure, Your Control

Enterprise AI UAE deployments keep prompts, documents and model inference inside approved infrastructure — so internal knowledge stays governed by your own policies rather than an external service.

DEPLOYMENT FLEXIBILITY

- ◆ Cloud-Isolated Tenancy
- ◆ Dedicated Private Cloud
- ◆ Fully On-Premise

Selected against your security classification, network model and hardware footprint.

01

Deploy Inside the Organization

Generative AI runs within customer-controlled infrastructure.

02

Sensitive Data Stays Approved

Documents remain within sanctioned storage and network zones.

03

Full AI Capability Stack

Private LLMs, assistants, RAG, automation and AI agents.

04

Three Deployment Paths

Cloud-isolated, private cloud or fully on-premise.

05

Built for Regulated Sectors

Government, defence-adjacent, energy, finance and utilities.

SECTION 02 — RATIONALE

Why Private / On-Premise AI

Nine operational drivers behind on-premise generative AI adoption in UAE government and enterprise organizations.

Greater Data Control

Data residency, retention and processing boundaries are defined and enforced by the organization.

Document Confidentiality

Internal policies, contracts and technical files are indexed inside your own environment.

Reduced External Dependency

Core AI workloads continue to operate independently of public AI service availability.

Custom Security Policies

Model access, logging, retention and content policies follow your internal standards.

Roles & Permissions

User, group and department-level entitlements mapped to existing directory structures.

Local Processing

Inference, embedding and OCR execute on infrastructure you operate and monitor.

Enterprise Integration

Native connection to ERP, CRM, DMS, GIS, SharePoint and internal APIs.

Predictable Architecture

Known hardware, capacity and cost profile — sized to measured workload.

Restricted Networks

Where required, the platform can operate within restricted or isolated network segments.

OPERATE & OBSERVE

Usage Monitoring, Audit Logging and Capacity Visibility

Administrators retain visibility of model usage, retrieval activity and infrastructure utilisation — supporting internal AI governance and cost planning for Private AI Infrastructure UAE deployments.

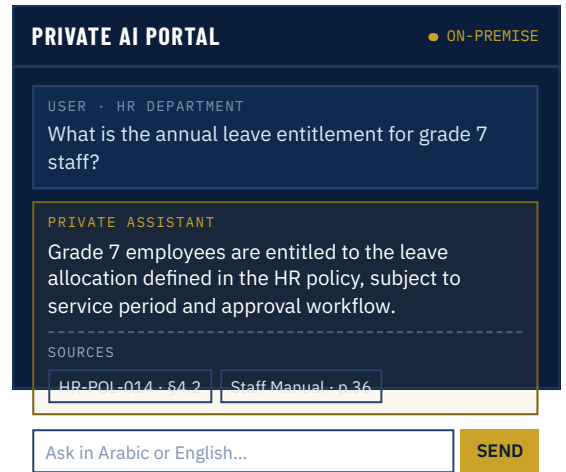
SECTION 03 – PRIVATE AI ASSISTANT

Build Your Own Enterprise ChatGPT

A private organizational AI assistant that answers from your approved internal knowledge — with source citations, document traceability and user-level access control.

ASSISTANT CAPABILITIES

- Chat With Internal Documents
- Manuals & Contract Lookup
- Generate Summaries
- Retrieve Approved Knowledge
- User-Level Access Control
- Search Policies & Procedures
- Arabic or English Questions
- Draft Reports & Correspondence
- Department-Specific Assistants
- Citations & Traceability



SECTION 04 – DOCUMENT AI & RAG

Turn Organizational Content Into Searchable AI Knowledge

An Enterprise RAG Solution converts scattered repositories into a governed, citable knowledge layer — without moving content outside your environment.

Supported Content

PDF	Word	Excel	PowerPoint
Scanned Docs	Policies	Contracts	Manuals
Reports	Drawing Metadata	Internal DBs	Email Archives

Scanned and image-based documents are processed with OCR, including Arabic script.



RAG PROCESSING WORKFLOW

01 Collect > 02 Clean > 03 Classify > 04 Index > 05 Embed > 06 Secure > 07 Retrieve > 08 Generate > 09 Cite Source

Document-level permissions are applied at retrieval time — users receive answers only from content they are entitled to access.

SECTION 05 – BILINGUAL CAPABILITY

Arabic Enterprise AI الذكاء الاصطناعي

Bilingual Government Workflows

- Arabic and English interaction in a single assistant
- Arabic document understanding and bilingual search
- Arabic OCR for scanned records and correspondence
- Arabic speech-to-text for meetings and dictation
- Arabic report and correspondence generation
- Mixed-language repositories queried together

BILINGUAL INTERFACE EXAMPLE

ما هي إجراءات الموافقة على المناقصات الداخلية؟

Tender approval follows the internal procurement procedure, with committee review thresholds defined by contract value.

PROC-SOP-07 AR / EN

Questions in Arabic can retrieve English source documents and vice versa.


One Private Assistant, Every Department

Each department receives its own scoped assistant, drawing only on the knowledge it is permitted to see.

SECTION 06 – ENTERPRISE AI AGENTS

Departmental AI Agents, Scoped to Their Own Knowledge

Each agent operates on a defined document set, with its own permissions, prompts and approval behaviour — all within your private deployment.



HR Agent

- ◆ Policies and employee FAQs
- ◆ Onboarding assistance
- ◆ Training recommendations
- ◆ Leave and HR process guidance



Legal Agent

- ◆ Contract review assistance
- ◆ Clause extraction and comparison
- ◆ Legal knowledge search
- ◆ Drafting support



Procurement Agent

- ◆ Tender analysis
- ◆ Proposal comparison
- ◆ Requirement extraction
- ◆ Supplier document analysis



Finance Agent

- ◆ Financial document search
- ◆ Invoice and report analysis
- ◆ Internal finance Q&A
- ◆ Management summaries



IT Service Desk Agent

- ◆ Internal troubleshooting
- ◆ Knowledge-base support
- ◆ Ticket assistance
- ◆ Technical document search



Executive Assistant

- ◆ Reports and meeting summaries
- ◆ KPI interpretation
- ◆ Decision-support information
- ◆ Technical knowledge lookup

HUMAN APPROVAL WORKFLOWS

Agents can be configured to draft and recommend, with defined actions routed for human review before execution.

SECTION 07 – GOVERNMENT & REGULATED ORGANIZATIONS

Government AI Solutions UAE – Use Cases

Ten deployment patterns commonly requested by authorities, regulators and public-sector entities in Abu Dhabi and Dubai.

01

Internal Government Knowledge Assistant

02

Policies & Regulations Search

03

Secure Employee Chatbot

04

Tender & Procurement AI

05

Inspection Report Assistant

06

Technical Standards Assistant

07

Legal & Compliance Assistant

08

Training Academy AI

09

Arabic Document Intelligence

10

Internal Correspondence & Report Generation



SECURE EMPLOYEE ACCESS

Every Answer Traceable to an Approved Internal Source

Employees query approved knowledge in natural language — in Arabic or English — while document-level permissions keep retrieval within each user's entitlement.

PRIVATE LLM FOR GOVERNMENT

SECURE ENTERPRISE AI

PRIVATE AI ABU DHABI

AGENT SCOPING MODEL

HR

Policies · Staff manuals

Legal

Contracts · Regulations

Procurement

Tenders · Supplier files

Finance

Reports · Invoices

IT

KB · Runbooks

Executive

KPIs · Board packs

AI Agents for Enterprise · Government AI Solutions UAE




www.swedishtechnology.com · +971 56 404 6555 · Abu Dhabi & Dubai

04

SECTION 08 – REFERENCE ARCHITECTURE

Private AI Architecture

A layered design in which every request passes authentication and access control before reaching models, retrieval or source systems.




SECTION 09 – SECURITY & GOVERNANCE

Controls Designed Into the Platform

Controls are configured to your internal standards and reviewed with your security team during design.

Role-Based Access Control Entitlements by role and department	User Authentication Directory-backed identity verification	SSO Integration Single sign-on via existing provider	Data Segregation Separated indexes per department
Encryption In transit and at rest	Audit Logging Queries, retrievals and admin actions	Model Access Policies Which groups may use which models	Document-Level Permissions Retrieval filtered per user rights
Network Segmentation Zoned architecture and firewall rules	Backup & Recovery Index, config and model snapshots	Human Approval Workflows Review before defined agent actions	AI Governance & Usage Monitoring Policy ownership and adoption reporting



SECTION 10 – MODEL ECOSYSTEM

AI Models & Technology Options

Open-weight and specialised models can be deployed within your environment, selected per use case rather than standardising on a single model.

MODEL STRATEGY

Multiple models run side by side — a general assistant model, a compact model for high-volume tasks, plus embedding, OCR and speech models.

EXAMPLE LANGUAGE MODEL FAMILIES (OPEN-WEIGHT)

Llama Broad general-purpose family with wide tooling support	Qwen Strong multilingual coverage including Arabic tasks	Gemma Compact sizes suited to constrained hardware	Mistral Efficient models for high-throughput workloads	DeepSeek Reasoning-oriented options for analytical tasks
--	--	--	--	--

Supporting Model Types <table><tr><td>Embedding Models</td><td>Vision Models</td></tr><tr><td>Speech-to-Text</td><td>Text-to-Speech</td></tr><tr><td>OCR (Arabic / English)</td><td>Computer Vision</td></tr><tr><td>Re-Ranking Models</td><td>Document Layout Parsing</td></tr></table>	Embedding Models	Vision Models	Speech-to-Text	Text-to-Speech	OCR (Arabic / English)	Computer Vision	Re-Ranking Models	Document Layout Parsing	Model Selection Depends On - Security posture - Available hardware - Use case profile - Licensing requirements - Arabic capability - Performance targets - Accuracy requirements - Support & lifecycle Model availability, versions and licence terms are confirmed with you at design stage.
Embedding Models	Vision Models								
Speech-to-Text	Text-to-Speech								
OCR (Arabic / English)	Computer Vision								
Re-Ranking Models	Document Layout Parsing								

MODEL-HARDWARE FIT

Model Size Is Sized to Your GPU Footprint

Quantisation, context length, concurrency and expected user volume determine GPU memory and node count — established during infrastructure sizing.

SECTION 11 – POSITIONING

Public AI Services vs Private Enterprise AI

Both models have valid roles. The comparison below describes deployment characteristics, not vendor quality.

DIMENSION	PUBLIC AI SERVICE	PRIVATE ENTERPRISE AI DEPLOYMENT		
Data Location	Processed in the provider's managed environment	Within customer-approved infrastructure or private cloud		
Data Control	Governed by the provider's service terms and settings	Governed directly by the organization's own policies		
Internal Knowledge	Available where content is supplied to the service	Indexed in place across internal repositories		
Custom Workflows	Configurable within the platform's feature set	Built to the organization's specific processes		
Access Control	Managed through the service's account model	Mapped to internal directory, roles and document rights		
Custom Models	Selection from the provider's model catalogue	Open-weight models selected, tuned and versioned by you		
Enterprise APIs	Integration via the provider's public APIs	Direct integration with internal systems and APIs		
Auditability	Reporting available per the provider's tooling	Full request, retrieval and admin logs retained internally		
Deployment Flexibility	Provider-hosted delivery model	Cloud-isolated, private cloud, on-premise or hybrid		
Organization Agents	Assistants built within the platform environment	Department-specific agents scoped to internal knowledge		
Many organizations operate both: public AI services for general productivity, and a private deployment for internal knowledge, regulated data and department-specific workflows — combined in a hybrid model where appropriate.				
ON-PREMISE GENERATIVE AI	ENTERPRISE GENERATIVE AI DUBAI	SECURE ENTERPRISE AI	PRIVATE LLM FOR GOVERNMENT	ENTERPRISE AI INTEGRATION UAE


SECTION 12 – ENTERPRISE INTEGRATIONS

A Secure Intelligence Layer Above Existing Systems

AI is added above your current landscape — systems of record remain authoritative, while the AI layer reads through governed connectors.

INTEGRATION PRINCIPLE


Connectors run read-only by default. Controlled write-back — such as drafting a ticket or a report — is enabled per use case with approval routing.



Connector scope, authentication method and refresh frequency are defined per system during integration design.

SECTION 13 – DEPLOYMENT OPTIONS

Three Ways to Deploy Private AI



OPTION 01
Fully On-Premise

AI servers and models installed inside customer infrastructure, under the organization's own operations.

- Maximum infrastructure control
- Suits restricted or isolated networks
- Customer-owned GPU hardware
- Internal operations and monitoring



OPTION 02
Private Cloud

A dedicated enterprise cloud environment with isolated tenancy and controlled network access.

- Faster to provision and scale
- Dedicated, isolated environment
- No on-site hardware footprint
- Managed operations optional



OPTION 03
Hybrid Enterprise AI

Private internal AI combined with approved external services where required by the use case.

- Sensitive workloads stay internal
- General tasks may use approved services
- Routing policy defined centrally
- Single portal for end users

SECTION 14 – AI SERVER INFRASTRUCTURE

AI Server UAE – Infrastructure Building Blocks

Indicative components only — the final configuration is defined by your sizing exercise, not fixed to one hardware specification.

Enterprise GPU Servers Rack-mounted nodes for model inference and embedding workloads	Redundant Power Dual PSU and UPS-backed distribution for continuity
NVIDIA GPU Infrastructure Data-centre accelerators selected to model size and concurrency	High-Speed Networking Low-latency interconnect between compute and storage tiers
High-Memory AI Systems System RAM sized for indexing, OCR and document pipelines	Scalable Multi-GPU Architecture Add nodes as user volume and use cases expand
NVMe Storage Low-latency storage for vector indexes and model weights	Cooling & Rack Readiness Thermal and power envelope assessed against your facility

SIZING APPROACH


Configuration Follows Measured Workload

Concurrent users, document volume, context length and response-time targets are measured first — the hardware configuration is then proposed as an example specification for your approval.

SECTION 15 – IMPLEMENTATION JOURNEY

From Readiness Assessment to Continuous Optimization

A staged delivery method — each phase produces a documented output reviewed with your project and security teams before the next begins.



SECTION 16 – EXAMPLE ENTERPRISE CASE STUDY

Government Organization – Private Knowledge AI

Illustrative scenario describing a typical deployment pattern. It does not represent a named customer and contains no measured performance figures.

CHALLENGE

Knowledge Spread Across Repositories

Thousands of internal policies, manuals, technical documents and procedures distributed across multiple repositories, with staff relying on manual searching and informal knowledge to locate the correct version.

SOLUTION

Swedish Technology deploys a private AI assistant using an on-premise LLM, RAG retrieval, secure document indexing and role-based access — integrated with the organization's directory and document repositories.

EXAMPLE OUTCOME

- Employees search approved knowledge using natural language
- Arabic and English support in one assistant
- Answers reference internal source documents
- Department-specific access enforcement
- Reduced manual document searching
- Faster internal knowledge retrieval

ADOPTION & TRAINING

Knowledge Reaches the Field, Not Just the Office

Technical staff query standards, manuals and procedures from the assistant on desktop or mobile, with training delivered per department during rollout.

SECTION 17 – NEXT STEPS

Build Your Own Private Enterprise AI

Keep Your Knowledge Inside Your Organization

Start with a secure AI proof of concept — a bounded pilot on a defined document set, delivered inside your own environment.

Request AI Workshop

Half-day use case discovery session

Request Private AI Demo

Live private assistant walkthrough

AI Infrastructure Assessment

GPU, network and facility sizing review

Talk to an AI Specialist

Direct line to our UAE AI team