R. Rockefeller S.C.

AI & CYBER DEPARTMENT

SOVEREIGN AI	SOFTWARE ENGINEERING	AI RESEARCH	CYBERSECURITY
Sovereign Multi-Agent Systems	**Custom AI Software Engineering**	**Advanced AI Model & Architecture Research**	**Cybersecurity & Resilience Engineering**
Pages 34 to 43	Pages 6, 7 and 30 to 32	Page 50 - Research programme	Page 33

THE REPORT

Agents that never see the data

Sovereign multi-agent systems, custom AI software engineering and cybersecurity.
Three tiers, one security gateway, two deployment tracks and four classes of agent.
Alongside delivery: advanced AI model and architecture research.

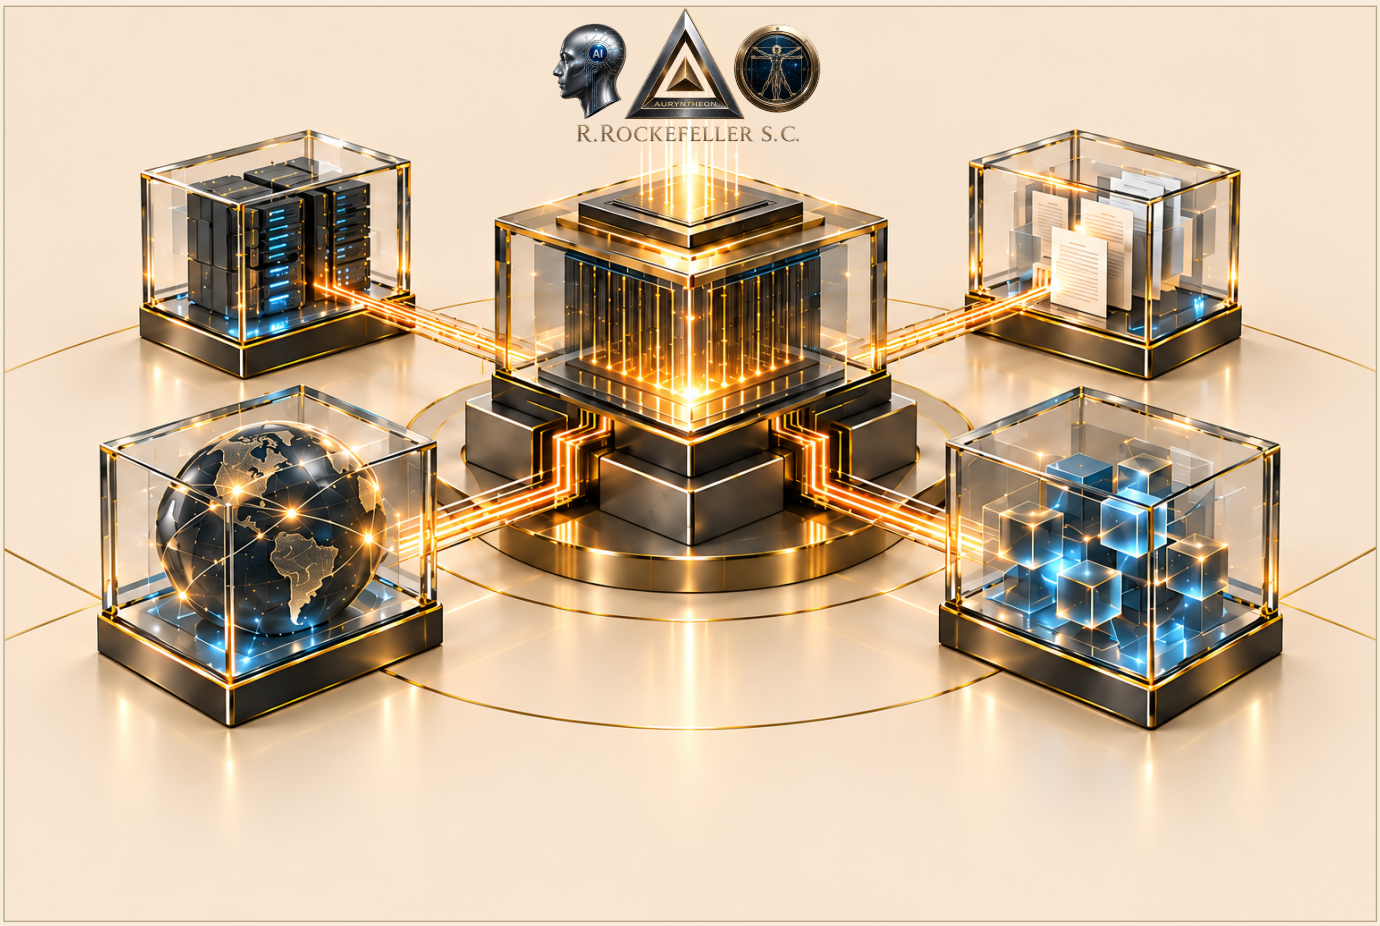

Sovereign AI engineering: secure orchestration, institutional workflows and model infrastructure.

The AI & Cyber Department of R. Rockefeller S.C. delivers custom AI software engineering and cybersecurity for institutional clients. Its end-to-end capabilities span model development, secure architecture, data engineering and deployment.

The sovereign multi-agent architecture defines records, documents, compliance and analysis agents, with human oversight and one security gateway.
Advanced AI model and architecture research remains a separate programme, reported at TRL 3.

IN THIS EDITION

INSIDE THIS EDITION

Contents, and how the report is ordered

What the Department does and the sectors it serves, then the team, then relations and centres; after that the method, what is new, the market and the evidence.

HOW TO READ THIS REPORT

For a decision maker. Part I, the sectors of interest in Part II, chapters 30 to 32 and chapter 42. **For a technical adviser.** Part V, chapters 31 to 38, and chapters 44, 45 and 47 for the evidence.

THE REPORT ON ONE PAGE

At a glance

Numbers, structure and status of the Department. The pages that follow explain each line; this page lets a reader hold the whole before reading the parts.

THE DEPARTMENT

2	50+	15	13	200+	14+
AI & CYBER CENTRES	CERTIFIED ENGINEERS IN LALITPUR	SENIOR MANAGERS AND EXPERTS	SENIOR KEY RESOURCES	CUMULATIVE PROFESSIONAL YEARS	YEARS OF INSTITUTIONAL DELIVERY

WHAT IS NEW, IN NUMBERS

3	1	2	4	5	0
TIERS IN THE ARCHITECTURE	SECURITY GATEWAY FOR EVERY AGENT AND MODEL	DEPLOYMENT TRACKS	CLASSES OF AGENT	CONTROLS IN THE GATEWAY	ITEMS OF CLIENT DATA SEEN BY THE MODEL

TWO PILLARS, AND THE LINES OF WORK INSIDE THEM

PILLAR ONE

Artificial intelligence

Sovereign build from zero	Architecture, corpus, tokeniser, training, platform and deployment: eight stages inside one perimeter.
Sovereign multi agent architecture NEW	Three tiers, one AI security gateway, two deployment tracks, four classes of agent.
Private model operation	The Group's private models operated for institutional clients through its centres, at its sole discretion.

PILLAR TWO

Cyber

Complete cyber practice	Offensive, defensive, architectural, cryptographic and regulatory: twelve domains, five forms of engagement.
Security of artificial intelligence	Model artefacts, training pipeline, inference boundary, agentic autonomy, model behaviour, deployment estate.
Strategic partnership	The Strategic Alliance signed in 2025 with SecurityScorecard: continuous rating and supply chain risk.

STATUS REGISTER: WHAT IS MATURE, WHAT IS IN FORMATION

WHAT	STATUS	ON WHAT EVIDENCE
AI software from zero, platforms and e-governance systems	Mature, in daily production	ISO 9001:2015 quality system; fourteen years of institutional delivery; executed references
Cyber practice	In delivery	Twelve domains; VAPT validated government deployments; strategic partnership with SecurityScorecard since 2025
Sovereign multi agent architecture	Specified, September 2026	Technical specification by the Scientific and Technical Lead; both deployment tracks defined
AI Center of Lalitpur, Nepal	Operational since 5 May 2026	More than 5,000 sq. ft., 50+ certified engineers, 15 senior managers and experts
AI & Cyber centre in the Emirates	In constitution	Institutional relations team resident; technical centre with a local partner, on the bench of Lalitpur
European operational base, Sicily	In formation	AI software development and research inside the European Union; work in progress
ISO/IEC 42001, AI management system	In formation	Implemented on the same lifecycle artefacts as the quality system; not reported as held
Proprietary architecture research programme	Research, TRL 3	Analytical proof complete; training phase pending and not claimed

THE DISCIPLINE OF STATEMENT What is mature is evidenced, what is in formation is declared as such, and what belongs to a later phase is not claimed in advance.

WHAT IS NEW

What is new in this edition

Since the Presentation of August 2026 the Department has added one line of work: multi agent systems for institutions that must keep their data under their own control. This page records what was added, in the order in which the report explains it.

	WHAT IS NEW	WHAT IT IS, IN ONE LINE	EXPLAINED IN
1	The sovereign multi agent architecture	A system of AI agents built in three tiers: access and identity, orchestration and AI security gateway, inference.	Ch. 31
2	The AI security gateway	One component through which every agent and every model passes: tokenization, injection defence, data loss prevention, human approval, audit.	Ch. 31, 32
3	The tokenization boundary	The model receives tokens, never data. The vault and the keys stay with the client; re-identification happens only inside the perimeter.	Ch. 32
4	Two deployment tracks	A proprietary track on the client's own servers, including disconnected environments, and a Microsoft aligned track for estates already in place.	Ch. 33
5	Four classes of agent	Records, documents, compliance, analysis: each with a defined scope, defined tools and defined permissions.	Ch. 34
6	The person in the loop	No decision issued, document sent or record changed without a person's approval where the institution's rules require one.	Ch. 35
7	Public and private institutions	The same design for ministries and for banks, corporate groups and family holdings: the rules change, the mechanism does not.	Ch. 34
8	A verified market position	Who delivers systems of this kind today, checked against public sources in September 2026.	Part VII

The Presentation of August 2026 documented a Department of two pillars: artificial intelligence software produced from zero, and a complete cyber practice. Both stand unchanged, and Parts I and V of this report record them. What is new is a line of work that institutions now ask for by name: systems of AI agents that work on their files, under their control.

Why it was added

Agents are the form in which artificial intelligence enters the daily work of an institution: not a chat window, but software that finds the file, drafts the document, runs the check and answers the question. The August Presentation already treated the security of agentic systems as a discipline of the cyber pillar. The new architecture turns that discipline into something an institution can commission.

What makes it different

Most agent offers place an orchestration layer on a foreign model and send the institution's data to it. The Department's architecture puts a security gateway between the agents and the model, replaces every identifying element with a token before anything reaches the model, runs on infrastructure the client controls, and treats the model as a part that can be replaced.

How to read the register

Each row is one addition, stated in one line, with the chapter that explains it. Part VI takes them in order and assumes no technical background. Part VII then sets the architecture against the market, on sources checked in September 2026.

IN ONE SENTENCE

The architecture puts the client's data behind a boundary that the AI model never crosses, runs the whole system on infrastructure the client controls, and makes the model a replaceable part.

CHAPTER 01 · THE DEPARTMENT

Two pillars, one institution

Within R. Rockefeller S.C. the AI & Cyber Department is the flagship. It stands on two pillars of equal weight, exercised under one technical authority, one quality system and one accountable management.

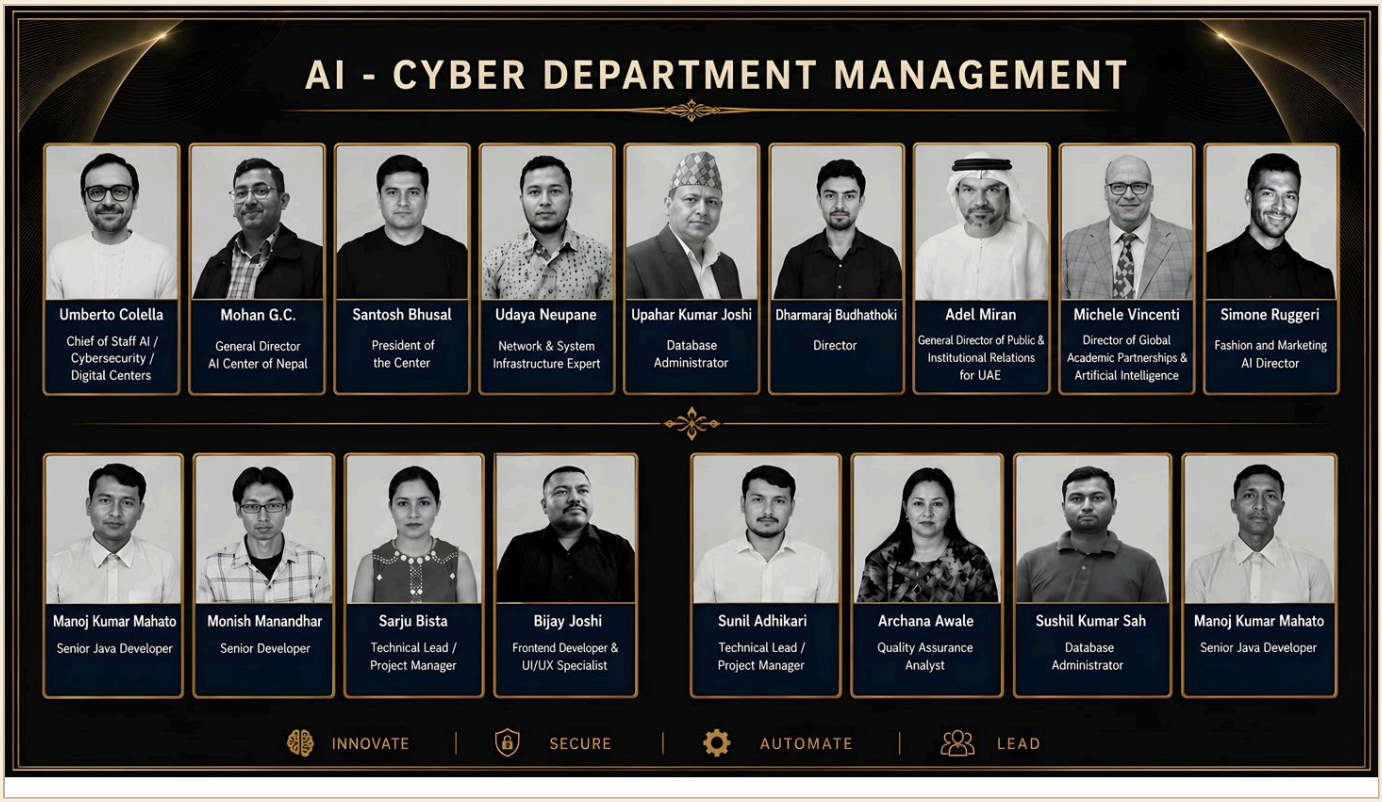

The management of the AI & Cyber Department: scientific authority and direction in the upper row, the engineering bench in daily production in the lower row.

The AI & Cyber Department is the flagship of R. Rockefeller S.C., and it stands on two pillars of equal weight. The first is a capability held by very few institutions in the world: to produce artificial intelligence software from zero, for any sector, end to end and under sovereign control. The second is a complete cyber practice, from offensive testing and security operations to cryptographic engineering, the protection of operational technology and regulatory assurance.

One technical authority

Both pillars are exercised under the technical authority of Mr Umberto Colella, Chief of Staff and Lead AI System Architect, through the R. Rockefeller S.C. AI Center of Lalitpur, Nepal, the Global Innovation Hub of the Group. One architectural authority, one quality system and one accountable management answer for everything the Department delivers.

Two centres, kept distinct

The Department works through two AI & Cyber centres: Lalitpur, operational since 5 May 2026, and the Emirates, in constitution. Its European operational base is being established in Sicily. They are a different structure from the Group's engineering centres for energy and technology, which are not the subject of this report.

The first horizon

The Department looks first to the Middle East, where it is present through its institutional relations team in the Emirates. Part IV is given to the region.

What this edition adds

To the two pillars the Department has added, in 2026, a new line of work: multi agent systems for institutions that must keep their data under their own control. Part VI explains it from the beginning.

WHAT THE PANEL RECORDS

The management of the Department. Upper row: the scientific authority and the direction of the centres, of institutional relations, of academic partnerships and of applied AI for fashion and marketing. Lower row: the engineering bench in daily production.

CHAPTER 02 · AI SOFTWARE FROM ZERO

AI software, engineered from zero

The core business of the Department, and the reason its counterparties come to it. Very few institutions in the world do this. The Department does it in daily production, under certification.

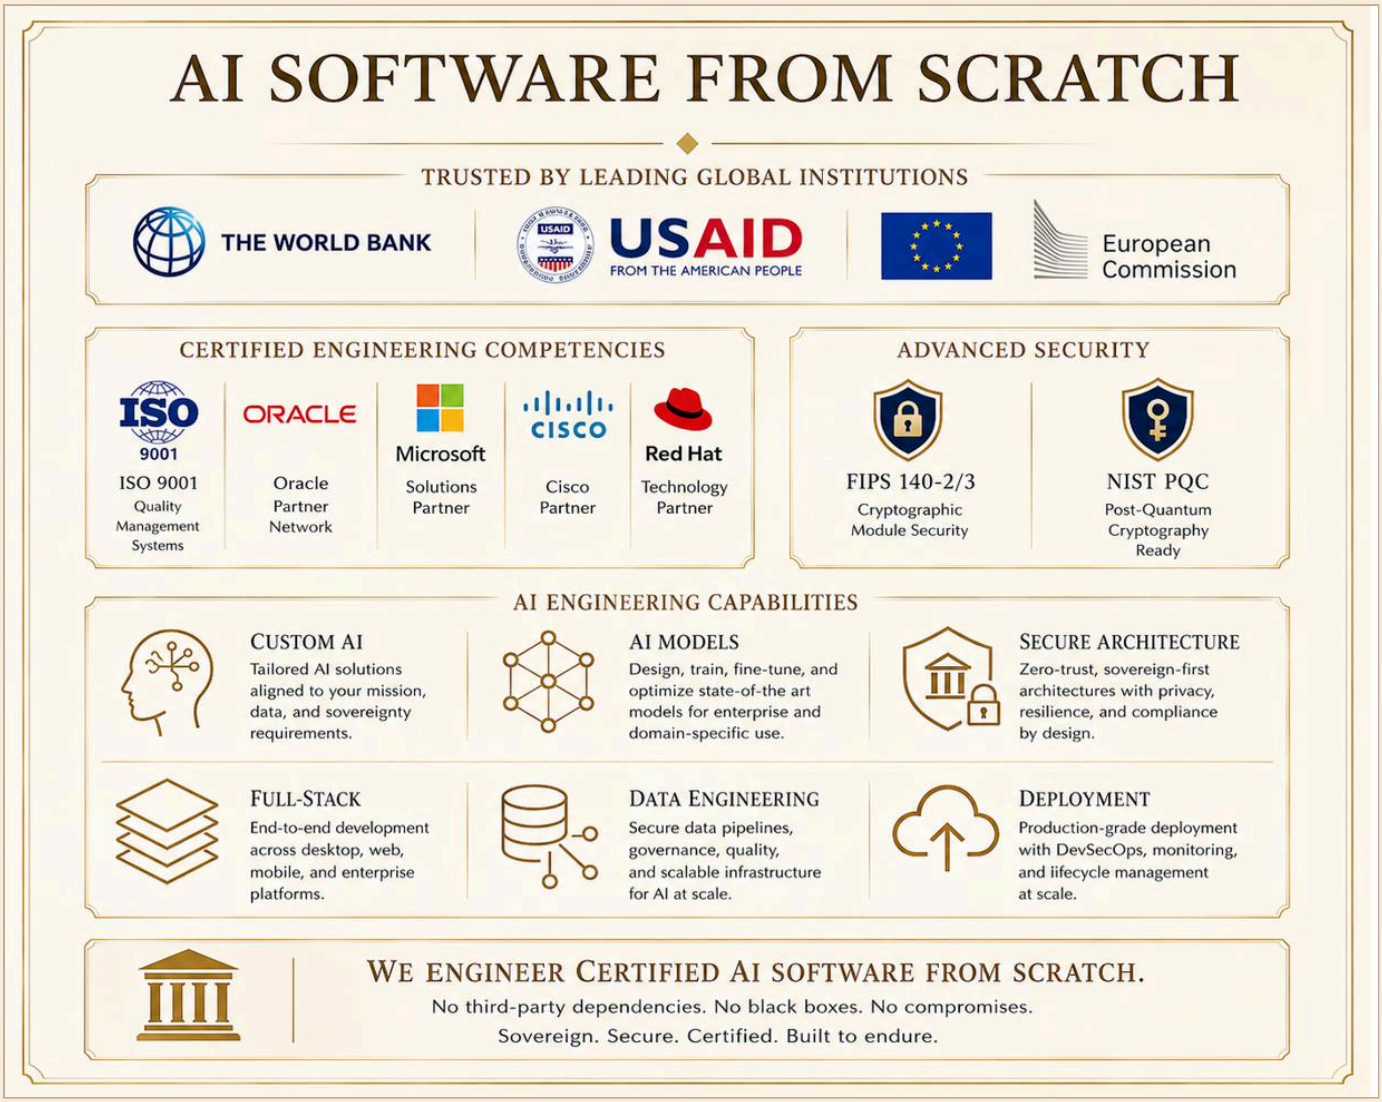

The engineering proposition on one panel: the institutions served, the certified competencies, the advanced security credentials and the six engineering capabilities.

Producing AI software from zero is not the assembly of third party components. It means authoring the architecture, training the models on the client's own data and controlled infrastructure, engineering the cryptography that protects them, building the platform that carries them, and operating the whole inside the client's sovereign perimeter.

Three things at once

It requires an architectural authority at the frontier of the field, a certified production floor with real engineers in real numbers, and a delivery record institutions can verify. The Department holds all three.

In figures

More than fifty certified engineers on the production floor, more than two hundred professional years across the senior team, fourteen years of institutional delivery.

THE LINE THAT DOES NOT CHANGE

In a sovereign build every line of code is written in house and the models are trained from scratch. Part VI adds a second line of work; it does not replace this one.

CHAPTER 03 · THE SPECIALISATION

What the Department does, capability by capability

The specialisation of the Department is narrow and deep: it writes artificial intelligence software from zero, and it secures it. Six engineering capabilities make that possible, held together under one quality system.

CAPABILITY ONE
Custom AI
Built for one institution's mission, data and sovereignty requirements.

CAPABILITY TWO
AI models
Designed, trained, adapted and optimised for institutional and domain specific use.

CAPABILITY THREE
Secure architecture
Zero trust, sovereign first, with privacy, resilience and compliance by design.

CAPABILITY FOUR
Full stack
Desktop, web, mobile and enterprise platforms, end to end.

CAPABILITY FIVE
Data engineering
Secure pipelines, governance, quality and scalable infrastructure.

CAPABILITY SIX
Deployment
Production grade, with DevSecOps, monitoring and lifecycle management.

Each capability is stated here with the way it is actually carried out, because the difference between the Department and an integrator lies in the how: the code is original, the models are sized to the problem, and the security is engineered in from the first stage.

Custom AI

Solutions built for one institution's mission, data and sovereignty requirements. The work begins with a sovereignty specification, which fixes the jurisdiction, the classification of the data, the infrastructure and the exit position, and against which every later decision is tested. It ends with a system the institution owns.

AI models

The Department designs, trains, adapts and optimises models for institutional use: conversational and assistant models, multilingual language models, document and process intelligence, analytical and forecasting models, and perception models for machines. It trains from one billion parameters upward and sizes each model to the problem rather than to fashion.

Secure architecture

Zero trust, sovereign first architectures, with privacy, resilience and compliance by design: a proprietary cryptographic stack that includes the 4096-bit algorithmic work of the Scientific and Technical Lead, custody of keys in hardware security modules, and vulnerability assessment and penetration testing before go live.

Full stack

End to end development across desktop, web, mobile and enterprise platforms: Django, React, .NET, Java and the Spring ecosystem, PHP and Laravel, with microservice and API architectures and enterprise integration.

Data engineering

Secure data pipelines, governance and quality: Oracle, PostgreSQL, MySQL and SQL Server at national scale, with high availability, disaster recovery and certified database administration.

Deployment

Production grade deployment with DevSecOps, monitoring and lifecycle management, on premise, on government cloud or on sovereign cloud, followed by the transfer of knowledge to the institution's own staff.

And, since 2026, agents

On these six capabilities the Department has built its new line of work: multi agent systems in which the institution's data never reaches the model. Part VI explains it from the beginning.

IN ONE LINE

We engineer the intelligence, and we engineer its security.

CHAPTER 04 · THE SECTORS

The sectors the Department serves

Nine sectors, one method. What changes from a ministry to a bank to a power utility is the data, the rules and the machines. The way the Department works does not change.

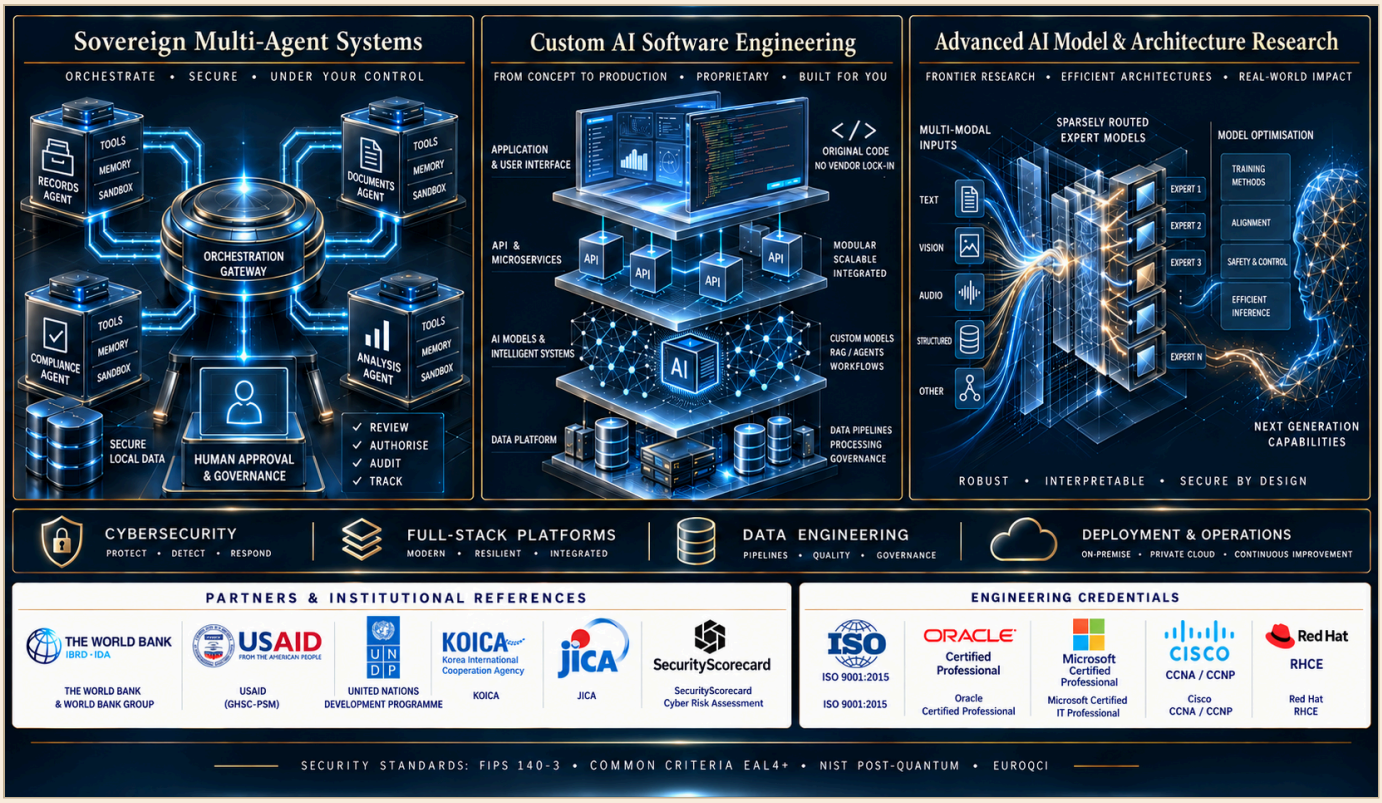

Technical specialisations of the AI & Cyber Department, with institutional references and engineering credentials.

The Department engineers AI software for any sector, and it names nine in which institutions ask for it most: government and public administration, financial services and banking, healthcare and life sciences, energy and utilities, industrial manufacturing, critical infrastructure, education and research, telecommunications, and smart cities with digital transformation.

How the work is done, in every sector

It starts from the institution, not from a product. The first stage fixes the jurisdiction and the regulatory perimeter, the classification of the data, the infrastructure on which the system will run and the exit position. Then the Department builds what the problem requires: the models, sized to the task; the platform that carries them; the security that protects them; and the agents that work inside the institution's processes. Everything is deployed inside the institution's own perimeter and handed over with documentation and trained operators.

What is common to all nine

Original code and no vendor lock in; security and privacy integrated in every layer of the lifecycle; an architecture that scales to mission critical use; clear processes with full documentation. The three chapters that follow take the sectors one by one: what the Department builds, how, and on what evidence.

Where agents enter

In every sector the new multi agent architecture adds the same thing: agents that find the file, draft the document, run the check and answer the question, behind a gateway that keeps the institution's data away from the model.

THE PANEL

The Department's technical specialisations and engineering foundations, with institutional references, professional credentials and security standards.

CHAPTER 05 · GOVERNMENT, BANKING, HEALTH

Government, banking, healthcare

The three sectors in which the Department's record is deepest, and in which the data is most sensitive.

13
MODULES OF THE LABOUR MANAGEMENT PLATFORM ILMIS

58
SITES OF THE ELMIS HEALTH LOGISTICS ROLLOUT

82%
FINANCIAL SERVICES RESPONDENTS WHO PUT DATA LEAKAGE AND COMPLIANCE FIRST

In these three sectors the institution holds files about people, and the law says what may be done with them. The work of the Department therefore rests on two things: systems of record that a state or a bank can run on, and intelligence that works on them without the data leaving.

Government and public administration

What. Treasury and budgetary information systems, revenue and tax platforms, labour and licensing systems, grievance redress with citizen channels, case management, registries and reporting. **How.** The platform and the intelligence are engineered together. Document and process models extract, classify and reconcile across registries, filings and case records. Multilingual assistants answer on the institution's own documents, including in non Latin scripts, with access control preserved and answers attributable to sources. Systems are deployed inside government data centres and government cloud, after vulnerability assessment and penetration testing. **Evidence.** Public finance through LMBIS and the Treasury Single Account; ILMIS, the labour management platform of thirteen modules; the grievance redress system of the Department of Roads, with SMS gateway and nationwide training; land records; electronic building permits; an integrated tax system.

Financial services and banking

What. Document and process intelligence, risk and reconciliation systems, and multilingual assistants that operate inside the bank's own perimeter. **How.** This is the sector in which the multi agent architecture is most immediately useful. The compliance agent runs the checks that anti money laundering, know your customer and sanctions rules require, on tokenized data, and escalates to a person every case that needs judgement. The records agent works under the same access rules the bank applies to its people. The model never receives a name or an account number, and every request leaves a record for the regulator. **Evidence.** A centralised core banking system and an online money booking platform among the delivered systems. In the IDC survey of 2026, 82 per cent of financial services respondents put data leakage and compliance first.

Healthcare and life sciences

What. Logistics and supply chain platforms for health commodities, digital health profiling and its linkage to social protection. **How.** Health data is personal data of the most sensitive kind. Systems are built for the health authority's own infrastructure, and where agents are used, the identifiers of patients are tokenized before anything reaches a model. **Evidence.** The eLMIS health logistics system, rolled out across fifty eight sites at two levels of the supply chain under the USAID Global Health Supply Chain Program; the digital health profile of two municipalities, linked to social health protection, delivered for GIZ.

What the three have in common

All three are governed by written rules applied to individual files, and all three are audited. The Department delivers them with the same discipline: requirement specification, configuration management, test strategy, penetration testing before go live, formal acceptance, and the transfer of knowledge to the institution's own staff.

IN ONE LINE

Systems of record a state can run on, and intelligence that works on them without the data leaving.

CHAPTER 06 · ENERGY, INDUSTRY, INFRASTRUCTURE

Energy, industry, critical infrastructure

Where intelligence meets machines and networks that cannot stop, and where the system must often run with no connection to the outside.

SECTOR
Energy and utilities
Forecasting and optimisation on the operator's own infrastructure; protection of SCADA and industrial control.

SECTOR
Industrial manufacturing
Operational intelligence for industrial estates; perception models that run on the device in the plant.

SECTOR
Critical infrastructure and defence
Systems that run with no external connectivity, under the artefact protection regime.

SECTOR
Robotics and autonomous systems
Embedded perception and conversational intelligence, with custody of the models at fleet level.

In these sectors the counterpart of the AI system is physical: a grid, a plant, a pipeline, a fleet. The Department brings its two pillars to them together, because a model that optimises an installation is also a new way into it.

Energy and utilities

What. Operational intelligence, forecasting and optimisation across energy infrastructure. **How.** Analytical and forecasting models, built on time series and structured data, support budgetary, logistical, energy and operational planning. They are sized to the problem and deployed on the operator's own infrastructure. The cyber pillar protects the operational technology around them: industrial control and SCADA environments, with segmentation that understands the industrial protocols and passive monitoring where active scanning is prohibited.

Industrial manufacturing

What. Operational intelligence for industrial estates, and perception for inspection and for the machines on the floor. **How.** Vision, speech and sensor fusion models are sized for constrained hardware and compressed by quantisation and distillation, so that they run on the device in the plant and not in a distant cloud. Custody of the models is kept at fleet level, under the same protection regime as every other artefact.

Critical infrastructure and defence

What. Systems engineered to operate without external connectivity, which is the ordinary condition of defence, critical infrastructure and classified deployments. **How.** Under the artefact protection regime, and within the dual use discipline of the Group: model weights protected at rest, in transit and at inference, keys in hardware security modules under split custody, integrity verified at each call. For agents, the proprietary track installs the whole system on the operator's own servers, including fully disconnected environments, with no dependency on any single external vendor.

Robotics and autonomous systems

What. Embedded perception and conversational intelligence for machines in the field. **How.** Models compressed for the hardware of the machine, with custody at fleet level.

Evidence inside the Group

The geothermal units of the Group's engineering centres are supervised by AI monitoring and remote control: real time connectivity, predictive diagnostics and autonomous optimisation of the system, with an operator dashboard for energy, temperature and flow.

Why both pillars are needed here

An AI system introduces exposures that conventional security programmes were not designed to see. In sectors where the consequence is physical, the Department builds the intelligence and engineers its security in one perimeter and with the same engineers.

IN ONE LINE

A model that optimises an installation is also a new way into it.

CHAPTER 07 · TELECOMS, CITIES, EDUCATION

Telecommunications, smart cities, education

Three sectors in which the volume of data is large, the public is the user, and the institution must answer for what the system says.

SECTOR
Telecommunications
Multilingual assistants on the operator's own procedures; document and process intelligence; risk and reconciliation.

SECTOR
Smart cities and digital transformation
Dashboards, building permits, road accident information, land records, citizen channels.

SECTOR
Education and research
Academic partnerships, joint centres, role based training and the formation of resident teams.

Here the intelligence speaks to people: subscribers, citizens, students. What matters is the language it speaks, the sources it can show, and the fact that the data of millions of people stays where the law says it must stay.

Telecommunications

What. Document and process intelligence, risk and reconciliation systems, and multilingual assistants operating inside the operator's perimeter. **How.** Language coverage is engineered for the jurisdiction, including non Latin scripts, for translation, summarisation, classification and extraction. Assistants answer on the operator's own procedures, with sources. **Evidence.** In the IDC survey of 2026, 75 per cent of telecommunications respondents put data leakage and compliance first among the reasons for sovereign AI.

Smart cities and digital transformation

What. The digital platforms of a city or a province: dashboards for the executive, electronic building permits, road accident information, land records, citizen channels for grievance redress. **How.** Platform and intelligence are engineered together, with analytical models for planning and, with the new architecture, agents inside public bodies under the controls a regulator expects. **Evidence.** The Chief Minister's dashboard of a province, under a programme of the Swiss cooperation; the road accident information system across sixty five sites; the electronic building permit system. The objective of Abu Dhabi, a government that is AI native by 2027, shows where the demand is going.

Education and research

What. The formation of talent, and the dialogue with universities. **How.** Through the academic directorate of Dr Michele Vincenti; through joint centres in which capability is transferred to resident personnel; and through role based training, board briefings and the formation of the client's own resident team. **Evidence.** The National IT Roadmap and the ICT roadmap of a national commission among the advisory references.

And many more

The list of the Department ends with the words and many more, and the reason is the method: a Department that writes its own code from zero is not limited to the sectors for which a vendor has a product.

What the three have in common

In all three the user is the public. The assistant must be able to show where an answer comes from, in the user's own language, and the operator must be able to show the regulator what was done with personal data. Retrieval against the institution's own documents, with sources, answers the first requirement; the tokenization boundary of the new architecture answers the second.

WHAT COMES NEXT

The next chapter records the verticals in a single register, with the five forms under which an institution can commission the work.

CHAPTER 08 · VERTICALS AND ENGAGEMENT

Verticals served, forms of engagement

Where the production capability is applied, and the contractual forms through which institutions engage it.

VERTICAL	SYSTEMS THE DEPARTMENT BUILDS
Public administration and finance	Treasury and budgetary information systems, revenue and tax platforms, labour and licensing systems, grievance redress with citizen channels, case management, registries and reporting.
Defence and critical infrastructure	Systems engineered to operate without external connectivity, under the artefact protection regime, within the dual use discipline of the Group.
Energy and industry	Operational intelligence, forecasting and optimisation across industrial estates and energy infrastructure.
Health and social protection	Logistics and supply chain platforms, digital health profiling and social protection linkage, of the class delivered under multilateral donor programmes.
Robotics and autonomous systems	Embedded perception and conversational intelligence for machines in the field, with custody of the models at fleet level.
Finance, telecommunications, enterprise	Document and process intelligence, risk and reconciliation systems, multilingual assistants operating inside the institution's own perimeter.

The upper register records the verticals in which the Department has delivered or for which it engineers today. The lower register records the five forms under which an institution can commission the work.

One method across sectors

A treasury system, a logistics platform for health supplies and a perception model for a machine in the field have little in common on the surface. They share the method: the same eight stages, the same quality system, the same security regime for the artefacts.

What the institution receives

In a sovereign build the institution receives a complete system on its own infrastructure, with the exit position defined at stage 1. In a joint centre it receives a permanent engineering capability under joint governance, transferred over time to resident personnel.

What each form is for

A sovereign build answers an institution that must own its system. A joint centre answers a country or a group that wants the capability itself to take root locally. Programme delivery answers an institution that already runs a programme and needs certified engineering capacity inside it. Security and assurance answers an institution that needs an independent view of systems it did not build.

The new form

To these forms the Department now adds the multi agent systems of Part VI, delivered on either of two deployment tracks and built as the client's own code.

THE MEASURE OF SUCCESS

An institution able to continue without its builder.

FORM OF ENGAGEMENT	WHAT THE INSTITUTION RECEIVES
Sovereign build	A complete system engineered from zero across the eight stages, on the institution's own infrastructure, with source, documentation and trained operators.
Joint centre	A permanent engineering centre established with a local partner or host institution, under joint governance, with capability transferred to resident personnel.
Programme delivery	Engineering capacity from the certified bench, integrated into the institution's own programme under the Department's quality system.
Private model operation	The qualified operation of the Group's private models for institutional clients through its centres, at the Group's sole discretion.
Security and assurance	Cryptographic architecture, deployment security, penetration testing, red teaming and continuous cyber risk rating, whether or not the Department built the system.

CHAPTER 09 · THE SCIENTIFIC LEAD

The scientific and technical lead

The authority under which both pillars are exercised, and the author of the specification of the new multi agent architecture.

Mr Umberto Colella, Chief of Staff and Lead AI System Architect: core areas of specialisation, track record and signature work.

The scientific and technical authority of the Department rests with Mr Umberto Colella, Chief of Staff and Lead AI System Architect of R. Rockefeller S.C. and the originating architect of its proprietary architecture. He is specialised in full stack algorithm design, advanced mathematics, cybersecurity, encryption and neural architectures.

Formation

He holds Bachelor's (2015) and Master's (2017) degrees in Software Engineering from the University of Technology Sydney, the Master's on the Machine Learning track with a final thesis on a proprietary 4096-bit encryption algorithm. Through the University's competitive international exchange programme he completed selective engagements at Google, on search algorithm optimisation and penetration testing, at Sony Corporation in Tokyo, on real time graphics software and advanced algorithms, and at Microsoft Corporation in Redmond, on encryption systems and adaptive threat detection, followed by cryptographic engineering of cross border settlement protocols at Euroclear.

A record held under confidentiality

A substantial part of his later record concerns dual use and national security programmes for countries of the North Atlantic alliance perimeter. That work is governed by confidentiality and classification obligations and is described here only by area of competence: defence AI, cryptography including quantum resistant protocols, neural architecture design, custom hardware and sovereign distributed systems.

His role in what is new

He is the author of the technical specification of the sovereign multi agent architecture presented in Part VI, and the guarantor of the architectural rigour of every deliverable of the Department.

SIGNATURE WORK

Proprietary 4096-bit encryption and zero trust models, optimised for security, performance and scalability.

CHAPTER 10 · THE TECHNICAL TEAM

The technical team, by function

Six functions, each with a name behind it: research, direction of the centres, software production, technical support, institutional relations, strategy and governance.

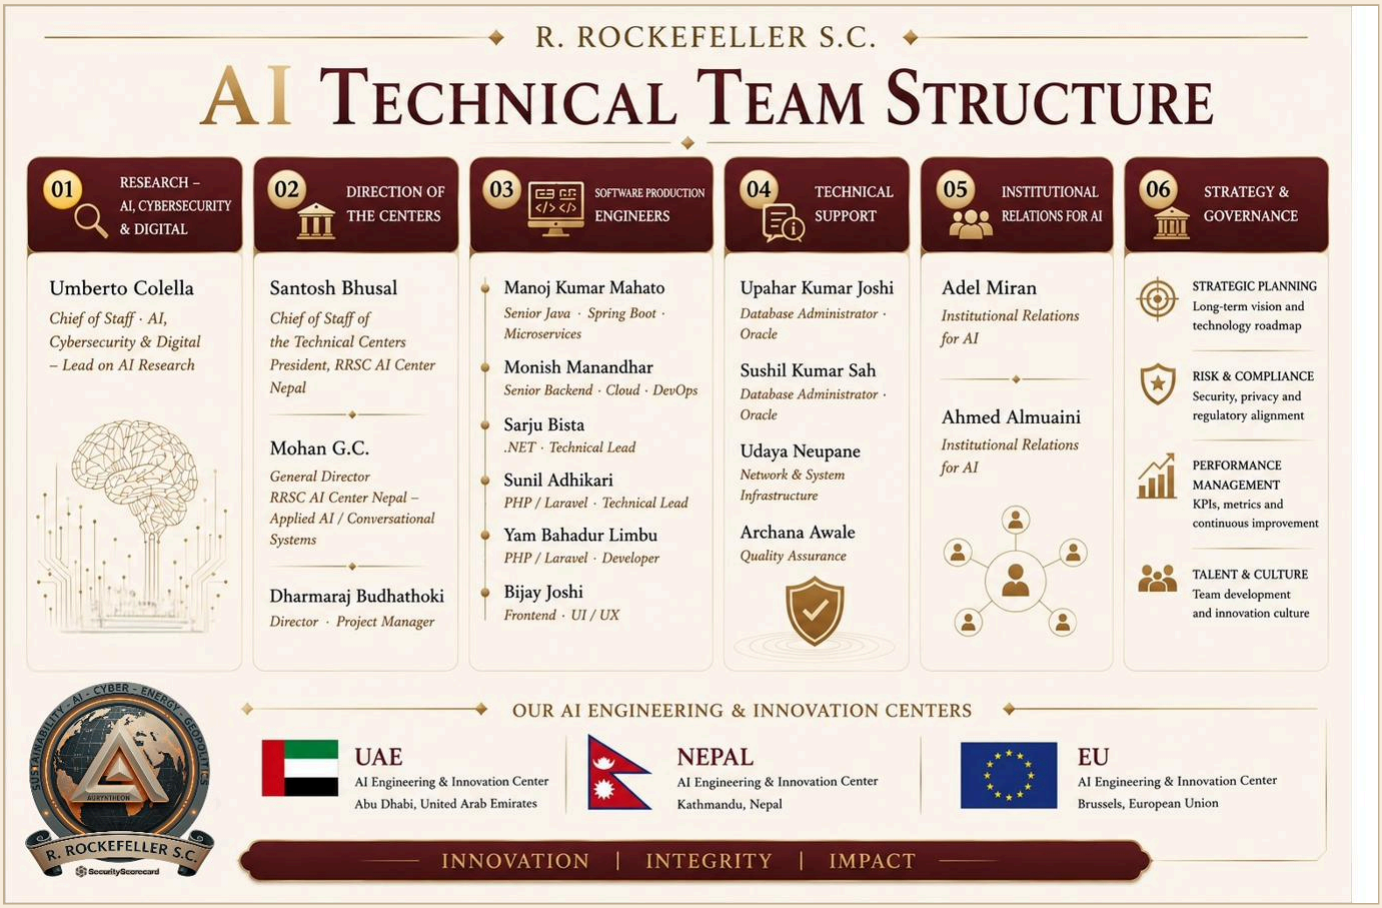

The AI technical team structure: six functions and the three poles of the network.

The Department is organised by function, in six columns, and each column has a name behind it. Research in artificial intelligence, cybersecurity and digital is led by Mr Umberto Colella. The direction of the centres rests with Mr Santosh Bhusal, President of the AI Center of Nepal, Mr Mohan G.C., its General Director, and Mr Dharmaraj Budhathoki, Director.

Production and support

The software production engineers are named individually with their specialisations: Java and Spring, backend and cloud, .NET, PHP and Laravel, frontend and interface design. Technical support covers database administration on Oracle, network and system infrastructure, and quality assurance.

Relations and governance

Institutional relations for AI are held by Dr Adel Miran and Mr Ahmed Almuaini. Strategy and governance covers strategic planning, risk and compliance, performance management and talent.

Relations as a function of its own

Institutional relations for AI is a function of the Department in its own right. It sits in the structure beside research, production and support, and it is held from the Emirates.

How the functions work together

A programme begins in institutional relations, which opens the counterpart. It is framed by research, which fixes the architecture, built by the production engineers, carried by technical support, and governed throughout by strategy and governance, which holds risk, compliance and performance.

THE LOWER BAND

The panel names the three poles of the network: Nepal, the Emirates and the European Union. Their status today: Nepal is operational, the Emirates are in constitution, and the European operational base is being established in Sicily.

CHAPTER 11 · LALITPUR

Lalitpur: the production floor

The R. Rockefeller S.C. AI Center of Nepal, Global Innovation Hub: the principal engineering production floor of the Group, certified, referenced and in daily production.

The R. Rockefeller S.C. AI Center of Nepal, with the three people who hold its direction.

The R. Rockefeller S.C. AI Center of Lalitpur is a centre of the Group: it operates under the Group's name and on the Group's behalf. Its operating entity, Dryice Solutions Pvt. Ltd., accepted full integration into the Group by the corporate instrument executed on 5 May 2026.

What the instrument does

It confers the official designation, places the centre's engineering capacity at the disposal of the Group's programmes, recognises thirteen senior professionals as members of the Group's global professional network, preserves the Group's existing intellectual property and binds the parties to confidentiality that is indefinite for trade secrets.

What the centre produces

Conversational models and vertical assistants with multilingual natural language processing; e-governance and public financial management systems; mission critical database engineering; and full stack development for desktop, web, mobile and enterprise platforms.

IN FIGURES

Lane 2, Kumaripati, Lalitpur. More than 5,000 sq. ft. of dedicated space, more than fifty certified IT engineers and developers, fifteen senior managers and experts, fourteen years of institutional record.

CHAPTER 12 · THE GENERAL DIRECTOR

Mohan G.C., General Director

Eighteen years between national information systems for World Bank programmes and the direction of a production floor of fifty engineers.

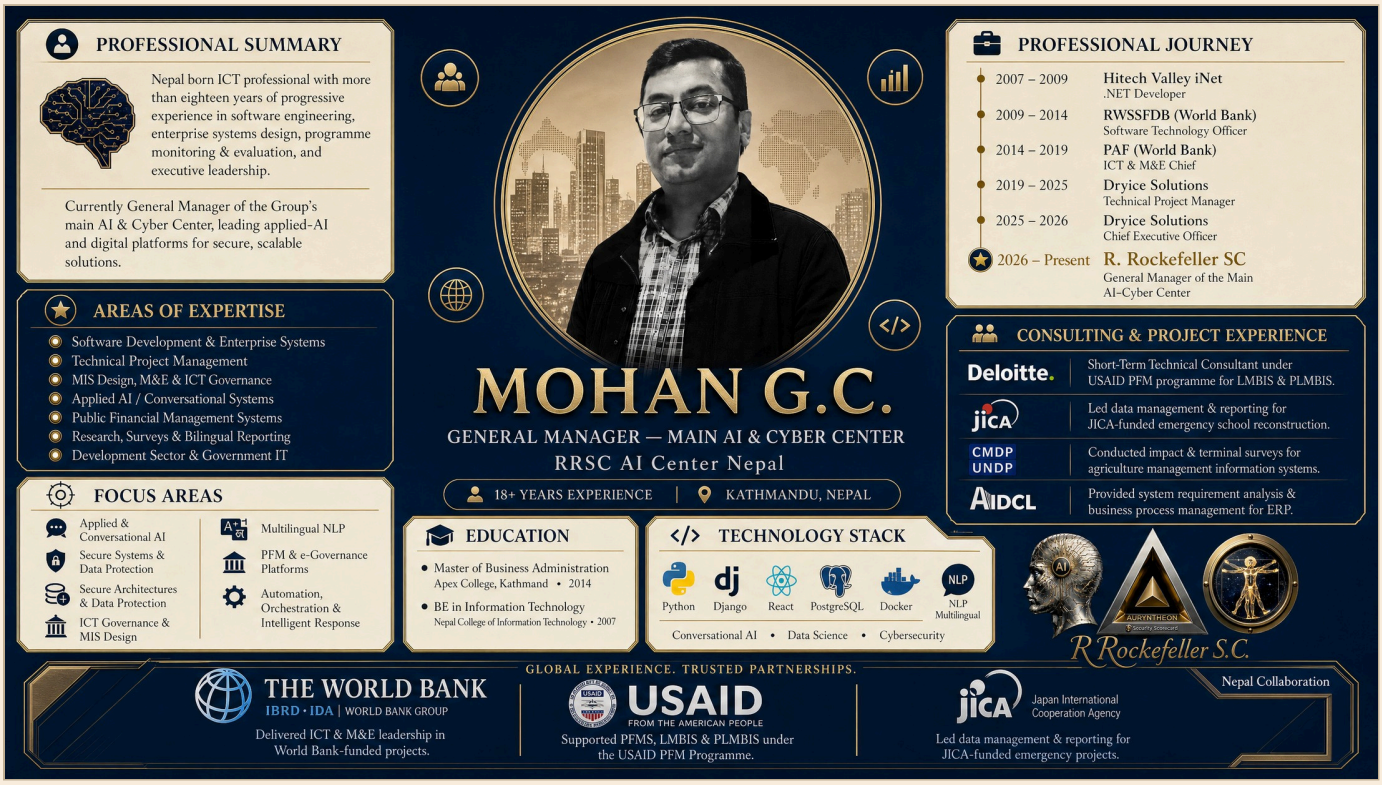

Mr Mohan G.C., General Director of the AI Center of Nepal: professional journey, areas of expertise and institutional experience.

Mr Mohan G.C. is General Director of the AI Center of Nepal and general manager of the Department's main AI and cyber centre. He is an ICT professional with more than eighteen years of experience across software engineering, enterprise systems design, programme monitoring and evaluation, and executive leadership.

The journey

He began as a .NET developer in 2007. From 2009 to 2014 he was Software Technology Officer of the Rural Water Supply and Sanitation Fund Development Board, and from 2014 to 2019 chief of ICT and of monitoring and evaluation at the Poverty Alleviation Fund, both programmes of the Government of Nepal funded by the World Bank. He joined Dryice Solutions in 2019 as Technical Project Manager and became its Chief Executive Officer in 2025.

Consulting and projects

For Deloitte he served as technical consultant under the USAID Public Financial Management programme on the budgetary information systems LMBIS and PLMBIS. He led data management and reporting for the emergency school reconstruction funded by JICA, and conducted impact and terminal surveys on agriculture management information systems for a UNDP programme.

Formation

Master of Business Administration, Apex College, Kathmandu, 2014. Bachelor of Engineering in Information Technology, Nepal College of Information Technology, 2007.

With him, the Director

Mr Dharmaraj Budhathoki, Director, is co-founder of the operating entity, with the company since 2012. His record covers the full software development lifecycle, from requirement analysis to national rollouts and ERP implementations for government bodies and for programmes of USAID and the World Bank.

APPLIED AI

In recent years he has directed his work toward civic technology enabled by AI, with conversational systems and multilingual natural language processing.

CHAPTER 13 · THE PRESIDENT OF THE CENTER

Santosh Bhusal, President of the Center

Founder of the operating entity of the centre, chief of its AI technical labs, and the person who holds the institutional representation of Lalitpur within the Group.

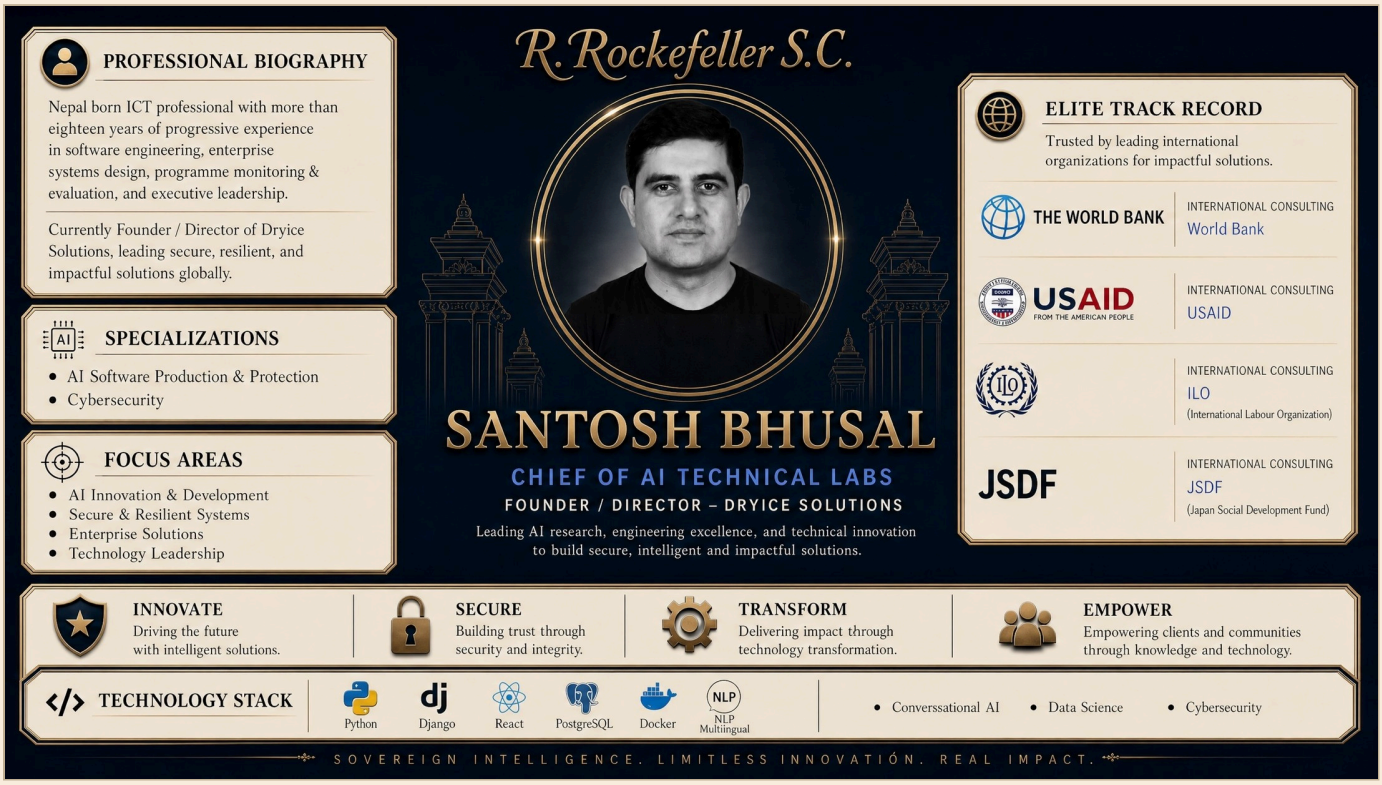

Mr Santosh Bhusal, President of the AI Center of Nepal and chief of the AI technical labs.

Mr Santosh Bhusal is President of the AI Center of Nepal and, in the structure of the Department, Chief of Staff of the technical centres and chief of its AI technical labs. He is the founder and director of Dryice Solutions, the operating entity of the centre, and he holds the institutional representation of the centre within the Group.

Formation and record

He holds a Bachelor's degree in Information Technology and a Master's in Business Studies, and brings more than seventeen years of international experience across the ASEAN region in system design, enterprise solutions and business process transformation.

International consulting

His track record is in international consulting for the World Bank, USAID, the International Labour Organization and the Japan Social Development Fund: ICT strategy, digital transformation and the implementation of enterprise systems across government and international projects, with expertise in ICT governance, system integration, cybersecurity and infrastructure management.

What he leads

AI software production and its protection, and cybersecurity: the two specialisations his profile records. He leads the international client engagement of the centre, and with Mr Mohan G.C. he signed the instrument of 5 May 2026 by which the centre was integrated into the Group.

The stack of the labs

Python, Django, React, PostgreSQL and Docker, with multilingual natural language processing, for conversational AI, data science and cybersecurity.

HIS MANDATE

Leading AI research, engineering excellence and technical innovation to build secure, intelligent and impactful solutions.

CHAPTER 14 · THE SENIOR LEADERSHIP

The senior leadership team

The five people who direct the centre of Lalitpur: the general direction, the presidency, the direction of delivery, the infrastructure and the databases.

The AI Center of Nepal: the senior leadership team.

The senior leadership of the centre joins three directors to the two specialists on whom every platform physically rests. Together they coordinate a senior management team of fifteen and a floor of more than fifty certified engineers.

Mohan G.C., General Director

More than eighteen years of ICT leadership with the Government of Nepal, across the Ministry of Finance, the Financial Comptroller General Office and the ministries of industry and of agriculture, and with the World Bank, IFAD, USAID, JICA and the Japan Social Development Fund.

Santosh Bhusal, President of the Center

He anchors the institutional governance of the centre and its relationships with multilateral donors, among them JICA, the Japan Social Development Fund and KOICA.

Dharmaraj Budhathoki, Director

Senior leader within the senior management team, coordinating the daily operational delivery of the centre across AI, cybersecurity and digital.

Udaya Neupane, Network and System Infrastructure

He leads network architecture, system infrastructure and security grade design, with a record in secure network architecture: firewalls, VPNs and data centre networking.

Upahar Kumar Joshi, Database Administrator

He leads database engineering, data governance and the structured data architecture on which the e-governance platforms are built.

IN FIGURES

Fifteen senior managers and experts. More than fifty certified engineers. More than 5,000 sq. ft. in Lalitpur.

CHAPTER 15 · THE TECHNICAL LEADERSHIP

The technical leadership team

Technical leads, project managers and database administrators: the operational backbone of the centre.

The AI Center of Nepal: the technical leadership team.

Four people carry the technical delivery of the centre from the first requirement to acceptance: two technical leads who are also project managers, and two database administrators.

Sarju Bista, Technical Lead and Project Manager

An expert in .NET, with a record across enterprise management information systems and HR systems. She coordinates the technical delivery of several projects at once across the AI, cybersecurity and digital programmes.

Sunil Adhikari, Technical Lead and Project Manager

An expert in PHP and Laravel, focused on the deliverables of institutional and governmental clients: complete project lifecycles and scalable application architecture.

Upahar Kumar Joshi, Database Administrator

Senior administrator on Oracle and Microsoft SQL Server. He leads database engineering and data governance for mission critical systems with high availability and regulatory adherence.

Sushil Kumar Sah, Database Administrator

Oracle expert. He co-leads database engineering and the structured data architecture, with a focus on enterprise grade Oracle deployments for national and government databases: performance tuning and disaster recovery.

Why two database administrators sit in the leadership

Every national platform the centre has delivered rests on a database that must stay available, recoverable and auditable: treasury, labour, health logistics, land records. The administrators answer for that, and they are certified to do so.

IN ONE LINE

From the first requirement to acceptance, the same people answer for the system.

CHAPTER 16 · THE SENIOR KEY RESOURCES

The senior key resources

The thirteen senior professionals recognised under the instrument of 5 May 2026: the operational core of the centre, with more than two hundred professional years between them.

SENIOR KEY RESOURCE	FUNCTION	PROFILE
Mohan G.C.	General Director	Executive leadership; eighteen years and more; World Bank and USAID programme record
Santosh Bhusal	President of the Center	Seventeen years and more of enterprise and business process transformation across the ASEAN region
Dharmaraj Budhathoki	Director	Full lifecycle delivery leadership; co-founder of the operating entity
Sarju Bista	Technical Lead, Project Manager	Fifteen years and more across enterprise MIS, HR systems and .NET platforms
Sunil Adhikari	Technical Lead, Project Manager	End to end project lifecycles for enterprise and government; scalable application architecture
Udaya Neupane	Network and System Infrastructure Expert	Fifteen years and more of secure network architecture: firewalls, VPNs, data centre networking
Sushil Kumar Sah	Database Administrator	Certified administrator; national and government databases; performance tuning, disaster recovery
Upahar Kumar Joshi	Database Administrator	Certified administrator; mission critical systems with high availability and regulatory adherence
Archana Awale	Quality Assurance Lead	Fifteen years and more; test strategy, automation, defect management on national and international projects
Manoj Kumar Mahato	Senior Java Developer	Fourteen years and more; Java and Spring, microservices, secure high performance systems
Monish Manandhar	Senior Backend Developer	Sixteen years and more; scalable backends, API development, database architecture, cloud
Yam Bahadur Limbu	Senior Developer	Eleven years and more; PHP and Laravel, RESTful APIs, system integration
Bijay Joshi	Frontend Developer, UI and UX Specialist	Thirteen years; responsive, user centred web applications

The thirteen Senior Key Resources are the operational core of the centre. Together they bring more than two hundred professional years, and each of them writes or runs production systems today.

Why the list matters

An institution that commissions a system is entitled to know who will build it. The register names the people, their function and the depth of their record: technical leads who have carried national platforms, certified database administrators for mission critical estates, a quality assurance lead with fifteen years across government and international projects, and senior engineers on Java, backend, PHP and interface design.

How the bench is used

For any new centre the senior positions are staffed through a combination of Group secondment, the Lalitpur bench and local recruitment. Capability is developed first where it exists and is then transferred, under institutional ownership, to where it is required. It is the method by which Lalitpur was integrated and by which the Emirati centre is being established.

Certified individuals

Across the senior bench the engineers hold individual professional certifications: Oracle Certified Professional, Microsoft Certified IT Professional, Cisco CCNP and CCNA, and Red Hat RHCE.

IN ONE LINE

Thirteen senior professionals, more than two hundred professional years, one production floor in daily delivery.

CHAPTER 17 · THE DEVELOPMENT TEAM

The development team

Four of the engineers who write the Department's code every day: the people an institution meets when its system is being built.

The development team of the AI Center of Nepal.

They are part of the thirteen Senior Key Resources of the centre, and between them they cover the languages in which the Department's platforms are written.

Manoj Kumar Mahato

Senior Java Developer, with more than fourteen years on Java and Spring, microservices and secure, high performance systems. He leads enterprise grade Java engineering and the structured backend services for multilateral and governmental clients.

Monish Manandhar

Senior Developer, with more than sixteen years in scalable backends, API development, database architecture and cloud.

Yam Bahadur Limbu

Senior Developer, with more than eleven years on PHP and Laravel, RESTful APIs and system integration.

Bijay Joshi

Frontend Developer and UI and UX Specialist, with thirteen years of responsive, user centred web applications. He leads the engineering of the applications that people actually use.

How the team works

Inside the ISO 9001:2015 quality system of the centre: requirement specification and design documents, configuration management, test strategy and defect management, vulnerability assessment and penetration testing before go live, formal acceptance.

What they have built

The platforms of the delivery record in Part VIII: public finance, labour management, health logistics, land records.

AROUND THEM

Technical leads and project managers, certified database administrators, network and infrastructure, quality assurance: the register of the previous chapter names them all.

CHAPTER 18 · THE MIDDLE EAST

The first horizon is the Middle East

The Gulf states have made artificial intelligence a matter of state, and they ask for it under national control. It is the region where the Department is building its second centre.

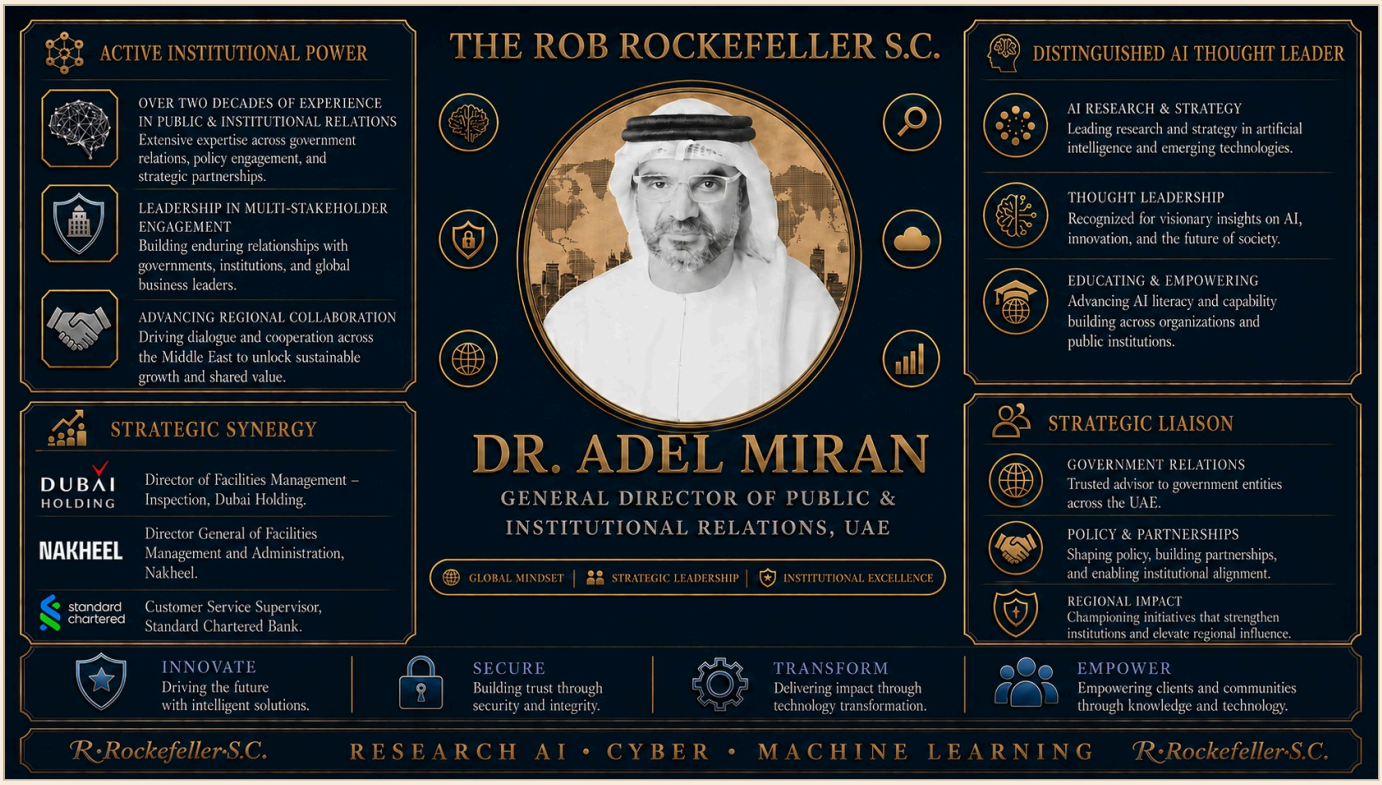

Dr Adel Miran, General Director of Public and Institutional Relations for the Middle East.

The Department looks first to the Middle East. The Emirates have G42 and the investment platform MGX; Saudi Arabia has Humain, backed by the Public Investment Fund; Qatar has established Qai under the Qatar Investment Authority. In February 2026, at the World Governments Summit, G42 announced a national agent factory for healthcare, education, legal services and energy.

What the region asks for

Agents, under national control, at scale. Data that stays in the country, and suppliers able to say what is inside the system they deliver. These are the requirements the architecture of Part VI was written for: the data behind a boundary the model never crosses, the system on servers the institution controls, the keys in the client's hands.

People before buildings

The Department is present in the region first through people: institutional relations for the Middle East are held by Dr Adel Miran with Mr Ahmed Almuaini, and the next chapter presents the team. The second AI & Cyber centre is in constitution in the Emirates.

The Emirates first

Within the region the Department begins from the Emirates: the institutional relations team is resident there, the Presidency anchors it in Abu Dhabi, and the second AI & Cyber centre is being constituted there. The four chapters that follow are given to it.

Saudi Arabia and Qatar

In the Kingdom of Saudi Arabia and in the State of Qatar the presence is one of representation only. The representation team is formed by two country Presidents: H.H. Prince Abdul Majeed bin Saud bin Muhammad Al Saud, President for Saudi Arabia, and Sheikh Tameen Bin Mohammed Al Thani, President for Qatar. There are no offices and no operations in the two countries at present.

THE GULF IN DATES AND FIGURES

13 billion dirhams for an AI native government in Abu Dhabi by 2027. Since **14 September 2024** the Saudi data protection law is fully enforceable.

CHAPTER 19 · THE EMIRATES

The team in the Emirates

In the Gulf nothing of consequence moves without standing, protocol and continuity. The presence in the Emirates was built first as a team of institutional relations; the technical centre follows it.

The organisation chart for the United Arab Emirates: the Presidency, the Chief Executive Officer and the four resident executives.

Every senior figure of the team carries a second mandate inside the Emirati establishment. That principle opens the doors that engineering alone does not.

Dr Adel Miran

General Director of Public and Institutional Relations for the Middle East. He is the bridge between the Department and the Emirati institutions: government relations, sovereign counterparties and protocol. He holds the mandate concurrently with a senior directorate of administration at Dubai Holding, and he is a recognised voice on artificial intelligence in the region.

Mr Ahmed Almuaini

Director of Branch Operations and Institutional Coordination, the coordinator at the Emirati federal level. With Dr Miran he holds institutional relations for AI in the structure of the Department.

What the chart records

The Presidency and its lineage, the dual mandate of Dr Adel Miran at Dubai Holding, the continuity of Ms Dolly Tabet with the Al Hamed family, the background of Mr Mohamed Al Hosani at the national energy company, and the federal coordination of Mr Ahmed Almuaini.

The resident leadership

Ms Dolly Tabet, Deputy CEO for the Emirates, brings more than twenty seven years of service to the Al Hamed family. Mr Mohamed Al Hosani, Executive Manager, comes from the headquarters of ADNOC. The Presidency is held by H.E. Sheikh Khalifa Bin Khalid Al Hamed; the company held with the Presidency is scoped to sustainability, and the AI and cyber perimeter is being constituted now.

WHAT FOLLOWS THE TEAM

A permanent AI & Cyber centre, in constitution with a local technical partner, engineered to the standards of Lalitpur and drawing on its bench while the resident team is formed.

CHAPTER 20 · THE EMIRATI TEAM

The team in the Emirates, one by one

Four resident executives under the Presidency. Each of them carries a second mandate inside the Emirati establishment: that is the principle on which the presence was built.

GENERAL DIRECTOR, PUBLIC AND INSTITUTIONAL RELATIONS · MIDDLE EAST

Dr Adel Miran

CONCURRENT STANDING

Senior directorate of administration and facilities management at Dubai Holding, the diversified global investment group of the Ruler of Dubai; previously at Nakheel.

FOR THE DEPARTMENT

Government relations, sovereign counterparties and the protocol without which nothing of consequence moves in the Gulf. A recognised voice on AI in the region.

DIRECTOR, BRANCH OPERATIONS AND INSTITUTIONAL COORDINATION

Mr Ahmed Almuaini

CONCURRENT STANDING

Member of a leading Emirati family with deep institutional and legal ties; doctorate from the University of Aberdeen.

FOR THE DEPARTMENT

The coordinator at the Emirati federal level: the mechanism by which standing becomes process. With Dr Miran he holds institutional relations for AI.

DEPUTY CEO · EMIRATES

Ms Dolly Tabet

CONCURRENT STANDING

French executive with more than twenty seven years of continuous service to the Al Hamed family.

FOR THE DEPARTMENT

The operational continuity at the centre of the daily engagement with the family's broader perimeter.

EXECUTIVE MANAGER · EMIRATES

Mr Mohamed Al Hosani

CONCURRENT STANDING

Senior background at the headquarters of ADNOC, the national energy company, with more than twenty years of corporate strategy and people development at the federal energy level.

FOR THE DEPARTMENT

One of the institutional addresses in the Gulf that carry the most weight.

The architecture of the Emirati presence rests on one principle: every senior figure carries an additional institutional mandate outside the Group, within the governance of the Emirates, in sovereign investment, in the federal ministerial structures or in the national energy operator. The panel above records the four resident executives, their standing and what each does for the Department.

What the team does for AI

It channels the capabilities of the Department, the sovereign AI software, the cyber practice and the digital platforms delivered through Nepal, into the highest levels of regional governance and enterprise. It is the permanent platform through which the Department engages the federal institutions and the sovereign actors of the region.

Dr Adel Miran and the public conversation on AI

Dr Miran brings more than two decades of experience in public and institutional relations: government relations, policy engagement and strategic partnerships, and the building of lasting relationships with governments, institutions and business leaders. His profile records three lines of work on artificial intelligence itself: research and strategy on AI and emerging technologies, recognised thought leadership on AI, innovation and the future of society, and the advancement of AI literacy and capability across organisations and public institutions.

Strategic liaison

For the Department he is a trusted adviser in the relations with government entities across the Emirates, shapes policy dialogue and partnerships, and champions initiatives that strengthen institutions and their regional influence.

From standing to process

In an environment in which access is governed by protocol and continuity, relations are not enough without an operations directorate that turns them into meetings, mandates and signed instruments. That is the function Mr Almuaini holds at the federal level.

IN ONE LINE

Every senior figure of the team carries a second mandate inside the Emirati establishment.

CHAPTER 21 · THE PRESIDENCY

The Presidency in the Emirates

The presence in the Emirates is presided over by H.E. Sheikh Khalifa Bin Khalid Al Hamed. It is the institutional anchor at the summit of Abu Dhabi, and it accompanies the Department as it constitutes its centre.

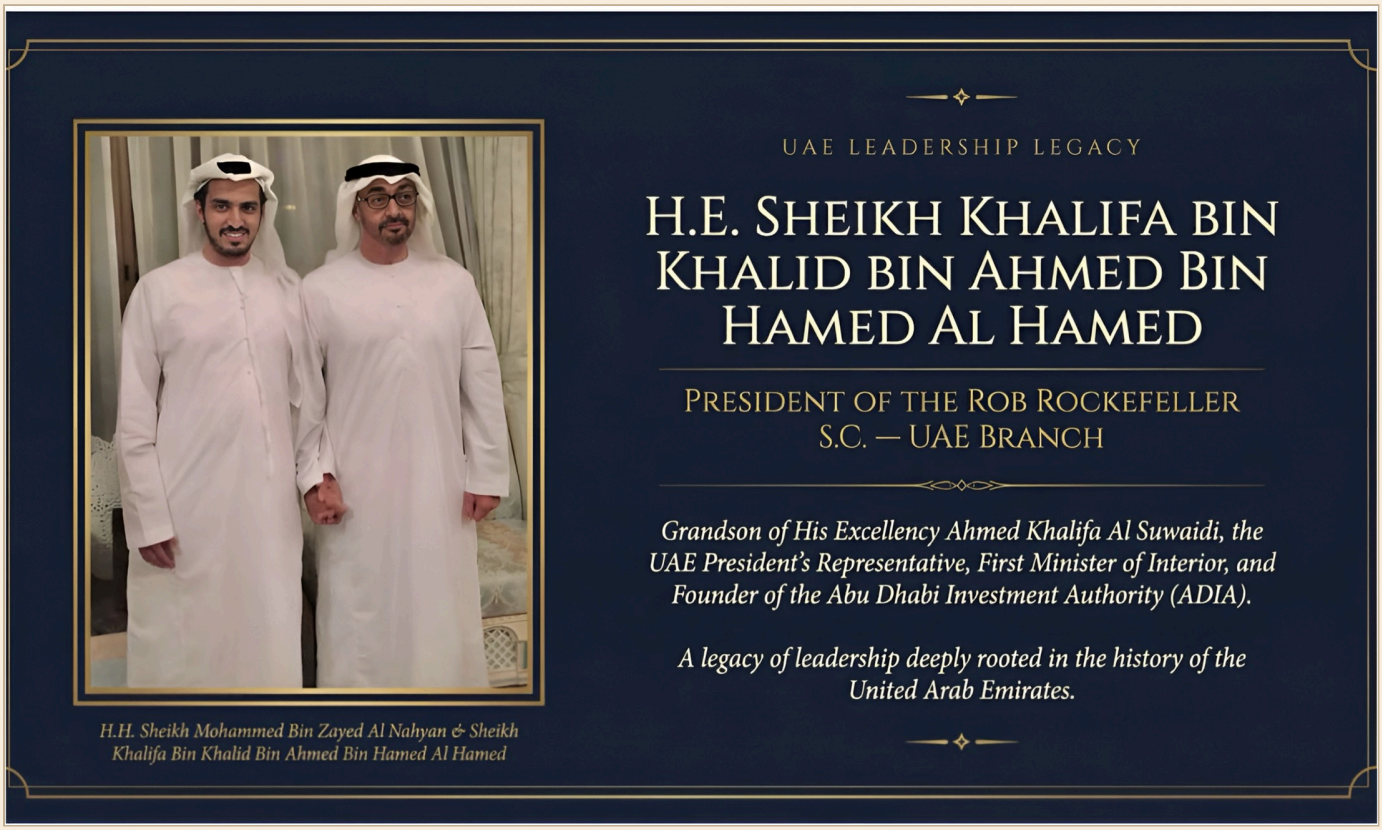

H.E. Sheikh Khalifa Bin Khalid Bin Ahmed Bin Hamed Al Hamed, President for the United Arab Emirates.

H. E. Sheikh Khalifa Bin Khalid Bin Ahmed Bin Hamed Al Hamed, of the Al Hamed family of Abu Dhabi, presides over the Group's presence in the United Arab Emirates. He is the grandson of H.E. Ahmed Khalifa Al Suwaidi, a statesman of the founding generation of the Emirates, the closest institutional collaborator of the founding father Sheikh Zayed, and the founder of the Abu Dhabi Investment Authority, the sovereign fund that defined for the world what patient sovereign capital means.

Not a ceremonial patron

As President he presides over the Emirati institutional presence, anchors it within the sovereign perimeter of Abu Dhabi, and represents the continuity between the generation that invented that model and the generation that is now redeploying it into sovereign technology.

A permanent platform

Around the Presidency stands the resident leadership presented in the previous chapter. Together they form the permanent platform through which the Department engages the federal institutions and the sovereign actors of the region: the institutional and diplomatic channel through which it operates, and intends to operate, in the Gulf.

The perimeter, stated precisely

The company the Group holds in Abu Dhabi with the Presidency is dedicated to environmental protection and sustainability, another vertical of the Group, and is distinct from the AI & Cyber Department. The Presidency and its support accompany the Group across both dimensions. For artificial intelligence and cyber, the Emirati company is being constituted now.

IN ONE LINE

The institutional component stands, at the summit of Abu Dhabi. The technical component is what the Department now adds.

CHAPTER 22 · THE EMIRATI CENTRE

The centre in constitution in the Emirates

The second AI & Cyber centre of the Department: a permanent operational seat that joins the institutional relations already in place to resident engineering, for the Gulf and for the African corridor.

IN PLACE

The institutional component

The Presidency and the resident institutional relations team, fully constituted and active.

BEING ADDED

The technical component

A permanent operational centre with resident engineering, in constitution with a local technical partner.

WHAT IT WILL DELIVER

For Gulf institutions

AI software from zero, multi agent systems behind a security gateway, and the cyber practice that protects both.

The step now under way in the Emirates is the establishment of a permanent operational centre through which the Department's engineering capability is channelled into the region. It integrates the institutional and diplomatic component already in place with a resident technical and productive one, engineered to the same standards, under the same quality system, and drawing on the same bench as Lalitpur.

Three poles, one network

The Department's network is organised on three poles: Lalitpur for engineering production, Europe for sovereign positioning, and the Emirates as the regional hub for digital, AI and cybersecurity initiatives across the Middle East and the broader African perimeter. With the new operational centre the Emirati hub adds a production dimension to its mandate of representation: resident engineering, and joint delivery with the local partner for Gulf institutional clients.

What the centre brings to the Gulf

The capability to produce any artificial intelligence software from zero, for any sector, end to end and under sovereign control: a capability that very few institutions possess, that is already operational and certified in Lalitpur, and that is the core of what the Department brings to the Emirates. To it the centre adds the new line of Part VI: multi agent systems in which the institution's data never reaches the model.

Why the region asks for both at once

Government and enterprise programmes in the Emirates increasingly ask for both at once: the intelligent system and the infrastructure, or the machine, that embodies it. The centre is the address for that demand across the Gulf Cooperation Council and the African corridor: sovereign AI software and the platforms that carry it, engineered, secured and operated under one governance and one quality system.

With a local technical partner

The AI and cyber perimeter in the Emirates is being constituted now, with prospective local technical partners able to work alongside the Department's engineering and to develop in the Emirates what it builds elsewhere.

How it is built

By the same phased method as every centre of the Department: definition, formalisation, foundation, execution and scale, each closed by a decision gate. The senior positions are staffed through secondment, the Lalitpur bench and local recruitment, and capability is transferred to resident personnel as it is formed.

What Part VI means for an Emirati institution

A ministry, a regulator or a bank that commissions agents from the centre keeps its data behind a boundary whose keys it holds, on servers it controls. For estates already standardised on Microsoft, the same design is delivered on that stack, inside the sovereign cloud. Abu Dhabi's objective of a government that is AI native by 2027 means agents inside public bodies: the architecture is built so that they can be placed on sensitive files with the controls a regulator expects, tokenization, human approval and a single audit trail.

THREE PERIMETERS OF COVERAGE

The Gulf. The Emirates first, and from there the Gulf Cooperation Council.

The wider Arab perimeter. Tunisia, Libya, Egypt and the Maghreb.

Africa. The continental network of the Group.

CHAPTER 23 · ACADEMIC PARTNERSHIPS

Academic partnerships for AI

The Department's relations are institutional in the Gulf and academic across the Atlantic. Dr Michele Vincenti, Associate Dean for Graduate Studies at University Canada West in Vancouver, directs its global academic partnerships in AI and cybersecurity.

Dr Michele Vincenti, Director of Global Academic Partnerships for AI and Cybersecurity, and Associate Dean at University Canada West.

Dr Michele Vincenti is Director of Global Academic Partnerships for Artificial Intelligence and Cybersecurity. He holds the mandate together with his academic position in Vancouver, where he is Associate Dean, Graduate Studies, at University Canada West: the first to hold that office, to which he was appointed in 2025 after chairing the department of leadership and people management.

Formation and record

He holds a PhD in Human and Organizational Systems from Fielding Graduate University and an executive MBA from Royal Roads University, and brings more than thirty years of experience in the financial sector, as a consultant and as an executive of financial institutions. He is the founder and president of Alvana Business Consulting, and he works in English and Italian.

What the mandate covers

The academic vertical of the Department in AI and cybersecurity: relations with universities and graduate schools, the formation of talent for the centres, and the dialogue between academia, industry and institutions.

Relations as a system

With the institutional relations team in the Emirates, the academic directorate completes the way the Department presents itself to the world: to governments and sovereign counterparties through Abu Dhabi, to universities and the professions through Vancouver.

A DUAL MANDATE

In the university. Associate Dean, Graduate Studies, University Canada West, Vancouver.

In the Department. Director of Global Academic Partnerships, AI and Cybersecurity.

CHAPTER 24 · AN ACADEMIC VOICE ON AI

An academic voice on AI

University Canada West publicly congratulated Dr Vincenti on his appointment. What he brings to the Department is a university's graduate school, and a record of public work on artificial intelligence and leadership.

Dr Michele Vincenti, and the public congratulations of University Canada West on his appointment.

The appointment of Dr Michele Vincenti as Director of Global Academic Partnerships for Artificial Intelligence and Cybersecurity was greeted publicly by his university, University Canada West, a business and technology oriented university in Vancouver founded in 2005.

At the European Parliament

In 2025 he represented the university at a conference hosted at the European Parliament, where he presented his work on artificial intelligence and leadership development, and on the potential and the limits of AI in deeply human domains.

Teaching and supervision

At the university he teaches investments, strategy, leadership, management consulting and corporate finance, and he has guided more than two hundred students through their final thesis or consulting project. As Associate Dean he leads the graduate programmes and their relations with employers, alumni and business leaders.

Between Canada and Italy

He speaks English and Italian and understands French and Spanish. His advisory work spans North America and Europe.

Why it matters to the Department

Centres are made of engineers, and engineers are formed in universities. The academic directorate gives the Department a channel to graduate schools for the formation of talent, and an academic voice on the responsible use of AI.

THE UNIVERSITY

University Canada West, Vancouver. Dr Vincenti is its first Associate Dean, Graduate Studies, appointed in 2025.

CHAPTER 25 · SICILY

Sicily: the European operational base

The Department's European operational base will be in Sicily. It is the third pole of delivery, after the production floor of Lalitpur and the centre in constitution in the Emirates, and the one from which European institutions will be served.

POLE ONE · OPERATIONAL SINCE 5 MAY 2026
Lalitpur
The production floor: codebase production, training pipelines, deployment automation, quality assurance.

POLE TWO · IN CONSTITUTION
The Emirates
The institutional relations team, and the second AI & Cyber centre for the Middle East.

POLE THREE · IN FORMATION
Sicily
The European operational base: AI software development and research inside the European Union.

The Department is establishing its European operational base in Sicily: a centre for AI software development and research through which European institutions and companies will be served from inside the European Union. The work is in progress, and this chapter states the role of the base and the method by which it is being built.

Why a base inside the Union

European clients increasingly require that the systems they commission are engineered, deployed and supported under European jurisdiction. The AI Act, the data protection regulation, the rules on the security of networks and on the operational resilience of the financial sector all presume a supplier able to answer for a system from within the Union. A base in Sicily gives the Department that address.

What Europe lacks

In the reading of the Department, Europe has excellent institutions for each single dimension of sovereign artificial intelligence: academic networks for research, the joint undertaking for high performance computing, national centres of excellence for applications. What it does not have is a permanent seat that integrates the four conditions of end to end sovereign engineering: original architecture design, training at production scale, deployment security, and vertical engineering across the sectors of strategic interest. The Sicilian base is conceived as an operational answer to that gap, working with the existing European institutions rather than competing with them.

Why Sicily

The island sits at the centre of the Mediterranean, between the European market and the Department's two other horizons, the Middle East and Africa. A European base there faces all three.

What the base will do

Produce AI software for European institutional and industrial clients, deliver on European soil the multi agent systems of Part VI, and carry out cybersecurity engineering, under the same quality system as Lalitpur.

How it will be built

By the method the Department applies to every new centre: in phases, each closed by a decision gate. Definition of perimeter and scope; formalisation of governance and intellectual property; foundation, with licensing, recruitment and first programmes; execution; and scale under a separate instrument. No phase presupposes the next. From inception the base is integrated with Lalitpur, which contributes senior engineering capacity during the foundation phase; as the Sicilian seat matures, the centre of gravity moves toward it, in proportion to recruitment and training.

European programmes

A seat inside the Union is also the condition for taking part in European programmes for research and digital capacity, such as Horizon Europe and the Digital Europe Programme, alongside European partners.

Milan and Sofia

Milan remains the administrative coordination centre for Italy. In Sofia the Group maintains a cooperation of a purely scientific nature with a centre of competence of the Bulgarian Academy of Sciences.

THE BASE IN BRIEF

Role. European operational base of the Department.

Work. AI software development and research; multi agent systems; cybersecurity engineering.

Method. Five phases, each closed by a decision gate; Lalitpur as the bench during foundation.

Status. In formation. This report states the role and the method, and claims nothing that is not yet in place.

CHAPTER 26 · THE EIGHT STAGES

From zero to deployment: the eight stages

What the production of AI software from zero actually consists of, stage by stage. The sequence is the discipline under which every institutional system of the Department is built.

	STAGE	WHAT IS DONE
1	Institutional framing and sovereignty assessment	Before a line of code is written: the jurisdiction and regulatory perimeter, the classification of the data, the infrastructure on which the system will run, the audit and reporting obligations, and the exit position. The output is the sovereignty specification against which every later decision is tested.
2	Corpus engineering	The training corpus is engineered, not inherited: licensing with documented provenance for every source, cleaning and deduplication at scale, domain corpora built with the institution's own experts, deliberate language coverage including non Latin scripts.
3	Tokenisation and representation	The vocabulary through which the model reads its language is built for the target languages and domains, with terminology, numerals, dates and identifiers treated as first class objects. Token efficiency translates directly into inference cost.
4	Architecture authoring	The model family is designed rather than selected: parameter budget and the fraction evaluated per query, class of attention and context length, depth and width, routing and gating. Decisions are calibrated at sub scale before production compute is committed.
5	Training engineering	A distributed systems discipline: parallelism across data, tensor and pipeline dimensions, memory and throughput engineering, checkpointing and failure recovery, curriculum and schedules, real time instrumentation.
6	Evaluation and ablation	Three levels of measurement: standard held out, domain evaluation built with the institution's experts, adversarial evaluation with the cyber pillar. Every architectural claim is supported by a documented ablation; regression suites accompany the system for life.
7	Adaptation, alignment and integration	Instruction and preference adaptation on curated data; retrieval against the institution's documents with access control preserved; behavioural policy set to the institution's own obligations; the surrounding platform of interfaces, workflows and dashboards.
8	Deployment, security and operation	On the institution's own infrastructure or sovereign cloud: containerised, reproducible, versioned, with model artefacts protected at rest, in transit and at inference and keys held in hardware security modules. Then monitoring, retraining and transfer to the institution's staff.

The expression building AI is used today for work of radically different depth, from the configuration of a commercial interface to the authoring of an architecture and the training of a model on original corpora. The distinction determines who owns the resulting system, where the institution's data resides, what happens when a foreign provider alters its terms, and whether the system can run at all inside a sovereign perimeter.

Why the order matters

Each stage fixes something the next depends on. The sovereignty specification of stage 1 decides where the system may run; the corpus of stage 2 decides what the model can know; the tokeniser of stage 3 decides what every later query will cost. Stage 4 is the one almost no institution outside a small international set performs at all.

Explained simply

A model is first of all the text it has read. If an institution cannot say what that text was, it cannot answer for the model years later. That is why provenance is recorded source by source, and why the result of each stage is a document as well as a piece of software.

The quality wrapper

All eight stages run inside the ISO 9001:2015 quality management system of the engineering centre: requirement specification and design documents, configuration management, test strategy and defect management, vulnerability assessment and penetration testing before go live, and formal acceptance.

EIGHT STAGES, ONE PERIMETER

Corpus, tokeniser, architecture, training, evaluation, adaptation, platform, deployment. No external dependency is inserted along the way.

CHAPTER 27 · STACK AND MODEL CLASSES

The stack, the model classes, the platforms

What the Department builds with, what it builds, and the forms in which institutional intelligence is delivered.

LAYER	WHAT THE DEPARTMENT BUILDS WITH
Model engineering	Python and the production deep learning toolchain; distributed training frameworks; custom data pipelines; experiment tracking and ablation infrastructure; quantisation and distillation for constrained targets.
Application engineering	Django, React, .NET, Java and the Spring ecosystem, PHP and Laravel, mobile and desktop clients; microservice and API architectures; enterprise integration.
Data engineering	Oracle, PostgreSQL, MySQL and SQL Server at national scale; high availability, disaster recovery and Data Guard configurations; certified database administration.
Infrastructure	Containerised deployment and DevSecOps; on premise clusters, government cloud and sovereign cloud; secure network architecture; monitoring and lifecycle automation.
Security	Proprietary cryptographic stack including the 4096-bit algorithmic work of the Scientific and Technical Lead; zero trust architecture; HSM key custody; vulnerability assessment and penetration testing.

The upper register records the stack, layer by layer. The lower register records the classes of model the Department engineers and the use an institution makes of each. Between the two lies the point that matters most to a client: the model is never delivered alone.

One deliverable

Around every model the Department builds the system that carries it: the e-governance and public financial management platforms, the sector management information systems, the dashboards and the operational workflows. The intelligence and the platform are engineered by the same institution, under the same quality system, and delivered as one deliverable.

Sized to the problem

The Department trains from one billion parameters upward and sizes each model to the institutional problem rather than to fashion. For machines in the field, perception and conversational models are compressed by quantisation and distillation so that they run on constrained hardware.

Explained simply

A parameter is one adjustable number inside a model. One billion parameters is a small model by today's standards and is often enough for a precise institutional task. What matters to the institution is how many of them must be evaluated for each query, because that is what it pays for in hardware and energy.

From models to agents

The conversational, document and analytical classes are the ones an agent uses when it works inside a process. Part VI explains how agents are composed and secured.

WHAT IS HANDED OVER

The model, the platform that carries it, the documentation and the trained operators.

MODEL CLASS	INSTITUTIONAL APPLICATION
Conversational and assistant models	Vertical assistants operating on the institution's own documents and procedures, with access control preserved and answers attributable to sources.
Multilingual language models	Language coverage engineered for the jurisdiction, including non Latin scripts; translation, summarisation, classification and extraction.
Document and process intelligence	Extraction, classification and reconciliation across registries, filings, claims, procurement files, case records and financial instruments.
Analytical and forecasting models	Time series and structured data models for budgetary, logistical, energy and operational planning.
Perception and embedded models	Vision, speech and sensor fusion models sized for constrained hardware: autonomous systems, robotics fleets, inspection, industry.
Private institutional models	The Group's own model portfolio, operated for institutional clients through its centres, at the Group's sole discretion.

CHAPTER 28 · THE FIVE TIERS

The five tiers of the market

Procurement frequently treats as equivalent five activities that differ by orders of magnitude in depth, ownership and exposure. This is the taxonomy the Department applies, and the place where the new agent line has to be located with care.

TIER	ACTIVITY	WHAT THE INSTITUTION ACTUALLY HOLDS
I	API integration and resale	An interface to a foreign provider's model. The institution holds a contract, not a system: the model, the weights, the roadmap, the pricing and the terms remain with the provider, and the data leaves the perimeter at every query.
II	Prompt and agent construction	Orchestration built on the same foreign model: workflows, prompt libraries, agent chains. Genuine engineering value at the workflow layer, but the intelligence remains external, and a change of provider policy can invalidate the work overnight.
III	Fine tuning of open weights	An adapted version of a third party's model. The institution holds an artefact it can host; it does not hold the architecture, the corpus, the tokeniser or the training pipeline, and it inherits every property of a base model it did not build.
IV	Continued pre-training and adaptation at scale	Substantial capability and meaningful model change. The architecture is still someone else's: the institution operates a derivative, and its independence lasts as long as the upstream licence.
V	Sovereign engineering from zero	Architecture authored, corpus engineered, tokeniser constructed, model trained from zero, cryptography proprietary, platform built, deployment inside the institution's own perimeter. The institution holds a system it owns, can audit, can operate and can continue without its builder.

The register is the taxonomy the Department applies when it is asked what distinguishes it from the market. Its sovereign build sits at Tier V. The consequences of the tier at which a system is built are operational, not theoretical, and each of them is a question a procurement authority is entitled to ask before signature.

Four questions before signature

Where does the data go at inference time, and under whose jurisdiction does it then sit. What happens to the deployed system if the upstream provider changes its terms, its pricing, its availability in the territory or its model behaviour. Can the system run without external connectivity, which is the ordinary condition of defence, critical infrastructure and classified deployments. And what does the institution own at the end: a running system with its documentation and its trained operators, or a dependency it must renew.

No tier is disparaged

Tiers I to IV have legitimate and often excellent applications. For many commercial problems they are the correct answer, and the Department itself integrates commodity components where sovereignty is not at stake. When the requirement is sovereign, the tier is not a preference. It is the requirement.

Where agents sit in this taxonomy

Tier II describes agents as the market usually builds them: orchestration resting on one foreign model, with the data sent to it at every query. Its two weaknesses are named in the register: the intelligence stays external, and a change of provider policy can invalidate the work overnight. The multi agent architecture of Part VI is designed against exactly those two weaknesses. The model receives tokens, never data, and it is a replaceable component: changing it does not touch the agents, the security or the identity layer. The orchestration engine and the security gateway are written by the Department as the client's own code.

Where the model itself must be sovereign

Where the institution must own the model as well, the sovereign build of Tier V applies, and the resulting model can take its place behind the same gateway. Chapter 39 sets the two lines side by side.

THE TEST

Where the orchestration runs, what the model receives, who approves, what is recorded, and what the institution owns at the end.

CHAPTER 29 · THE CYBER PRACTICE

The cyber practice: the second pillar

Cyber is not an accessory of the artificial intelligence practice. It is the second pillar, of equal weight, exercised across the full perimeter of the discipline.

DOMAIN	WHAT THE DEPARTMENT PERFORMS
Offensive security	Penetration testing of networks, applications, mobile and physical perimeters; red and purple teaming.
Vulnerability and exposure management	Continuous assessment, attack surface discovery, remediation engineering, re-testing to closure.
Security architecture	Zero trust design, segmentation, identity and privileged access, PKI, HSM design and key custody.
Cryptographic engineering	The proprietary stack including the 4096-bit work; key lifecycle; post quantum migration planning.
Detection and response	Security operations centre design and operation; threat hunting; incident response; forensics.
Application and product security	Secure development lifecycle, threat modelling, code analysis, DevSecOps pipelines.
Cloud and infrastructure security	Sovereign and hybrid cloud hardening, container security, isolation, backup and recovery.
Operational technology	Industrial control and SCADA: protocol aware segmentation, passive monitoring.
Third party and supply chain risk	Continuous rating of vendor ecosystems, supply chain detection and response.
Governance, risk and compliance	Audit readiness: information security and AI management systems, sectoral regulation.
Human layer and capability building	Awareness, role based training, crisis exercises, formation of the client's own team.
Security of artificial intelligence	Model artefacts, agentic systems, adversarial testing, integrity of data and training pipelines.

FORM ONE
Assess
Penetration test, architecture review, compliance gap, third party exposure baseline.

FORM TWO
Build
Zero trust design, cryptographic infrastructure, identity, secure pipelines, hardening.

FORM THREE
Run
Monitoring, detection engineering, continuous rating of the institution and its suppliers.

FORM FOUR
Respond
Forensics, containment, recovery, reporting to board and regulator.

FORM FIVE
Transfer
Formation of the client's own resident team: the institution able to continue alone.

The register records the twelve domains of the practice: offensive, defensive, architectural, cryptographic and regulatory. The band beneath it records the five forms in which the practice is engaged, from a single assessment to the formation of the client's own team.

The cryptographic core

The pillar carries the Group's cryptographic infrastructure: proprietary encryption architectures including the 4096-bit algorithmic work of the Scientific and Technical Lead, zero trust models, and the deployment security regime that protects model artefacts at rest, in transit and at inference, with keys held in hardware security modules under split custody.

A strategic partnership of the Group

The practice is consolidated by the Strategic Alliance signed in 2025 with SecurityScorecard, the global reference in cybersecurity ratings, executed directly between Mr Valentino Loforese and Dr Aleksandr Yampolskiy, its Chief Executive Officer and co-founder. It is a partnership of the Group, not a unit of the Department, and it gives the Department an instrument of measurement the market recognises.

Measure and remediate

Rating institutions measure; engineering institutions build. The Department sits on both sides of that line: it can measure an exposure with a recognised instrument and remediate it at the level of the architecture.

TWO PILLARS, ONE INSTITUTION

We engineer the intelligence, and we engineer its security. Neither is a service line attached to the other.

CHAPTER 30 · WHAT AN AGENT IS

What an agent is, and why institutions ask for it now

The word agent is used loosely, and institutions are being offered very different things under it. This chapter fixes the meaning used in the rest of the report, in plain language, before the architecture is described.

ONE

A model answers

It receives a question and returns text. It does nothing else, and it remembers nothing of the institution's work.

TWO

An agent does a job

Software that uses a model to carry out one defined task: it reads what it is allowed to read, uses the tools it has been given, keeps the state of the work and prepares a result for a person.

THREE

A multi agent system runs a process

Several narrow agents, each with one responsibility, composed by an orchestration engine and placed behind one security gateway.

A language model, taken alone, answers. Ask it to draft a letter and it drafts one; it does not know which file the letter belongs to, it cannot open the register, and it will not remember tomorrow what it wrote today. An agent is the software built around the model to turn an answer into a task: it is given a scope, a set of tools, a memory of the work in progress and a set of permissions, and inside those limits it prepares a result. In the words of the Department's specification, it is a specialised worker with a defined scope.

What an agent is not

An agent in this sense is not a chatbot with a new name, and it is not robotic process automation, which repeats fixed steps and stops when the screen changes. It reasons over the content of the file in front of it, within limits the institution sets, and it leaves the decision to a person.

From one agent to several

One general agent asked to do everything is hard to test and harder to trust. The Department defines agents by responsibility: records, documents, compliance, analysis. Each has a defined scope, defined tools and defined permissions. An orchestration engine decides which agent does what, in which order, with which tools and which memory, and combines them into a process the institution already recognises: find the file, draft the document, check it against the rules, answer the question with sources.

Why institutions ask for it now

The demand is visible in the programmes that set expectations in the regions where the Group works. In February 2026, at the World Governments Summit, G42 announced a national agent factory for healthcare, education, legal services and energy. Abu Dhabi has committed 13 billion dirhams to a digital strategy for 2025 to 2027 whose objective is a government that is AI native by 2027. Institutions in these markets will expect agents, under national control, at scale.

Why they hesitate

An agent with credentials and tools is an unattended user account with a reasoning engine attached. It can read files and act on systems, and that is a risk unless every input is filtered, every output is checked, every action that matters is approved by a person, and everything is recorded. Gartner expects more than 40 per cent of agentic projects to be cancelled by the end of 2027, and inadequate risk control is one of the three reasons it gives.

What the Department therefore built

Not an agent product placed on top of somebody else's platform, but an architecture in which the controls are the layer through which the agents run. The specification was written for engineers and says what the system is. The chapters that follow say why the design matters, tier by tier, and assume no technical background.

IN ONE SENTENCE

The architecture puts the client's data behind a boundary the AI model never crosses, runs the whole system on infrastructure the client controls, and makes the model a replaceable part.

CHAPTER 31 · THE ARCHITECTURE

Three tiers, explained from the beginning

The system is built as three layers stacked on top of each other. Each layer has one job. The middle layer is where the sovereignty of the data is guaranteed.

TIER 1	**Access and identity**	Who may enter, with which rights. Multi factor authentication, roles, sessions that expire. Connected to the identity system the institution already has.
TIER 2	**Orchestration and AI security gateway**	Agents, tools, memory. De-identification and tokenization before anything reaches the model. Injection defence, data loss prevention, human approval, audit. This is the data sovereignty boundary.
TIER 3	**Inference**	The language model, on the institution's servers or external. It receives tokens only and never sees a real name, account or address. Answers are checked and re-identified inside the perimeter.

Imagine a building with three floors. On the ground floor a reception desk checks who is entering, with which badge and with which rights. On the first floor are the offices where the work is organised: tasks are assigned, documents are prepared, and every document that has to be sent upstairs is first stripped of names, account numbers and anything that identifies a person or a client, and replaced with codes. On the top floor a very capable analyst reads the coded documents, does the thinking and sends back an answer. The analyst never sees a real name. The first floor puts the real names back before the answer reaches the person who asked. The specification calls it a sandwich, because the security sits in the middle, between access and intelligence.

Tier one: access and identity

Nobody reaches the system without being identified. The specification requires multi factor authentication, roles that determine what each person may do, and sessions that expire. The tier connects to the identity system the institution already has, so that the rules governing its other systems govern this one.

Tier two: orchestration and the security gateway

This is the heart of the system and the part the Department writes. The orchestration engine decides which agent does what, in which order, with which tools and which memory. The security gateway, in the same tier, does five things: it removes and tokenizes personal and sensitive data before anything reaches the model; it defends against attempts to manipulate the agents through their inputs; it applies data loss prevention rules; it requires a person's approval for actions that matter; and it records everything. This tier is the data sovereignty boundary: on one side is the institution's data, on the other side are tokens.

Tier three: inference

The language model. It can run on the institution's own servers, in which case nothing leaves the building at all, or it can be an external model, in which case it receives only tokenized data. The answer returns through the gateway, which checks it for safety and for groundedness, that is, whether it is actually supported by the documents, and re-identifies it inside the perimeter. Keys are managed by the client, with hardware security modules where required.

What the institution sees

For the people who use it the system is a single workplace: they sign in as they do today, they ask, they review what the agents prepared, they approve. The three tiers are invisible to them and visible to the auditor.

WHY THREE AND NOT ONE

Because each tier can be checked, replaced and audited on its own. The model can change without touching the security. The security can be strengthened without touching the agents. The identity system can be the client's own. A single block would have to be rebuilt every time one part changed.

CHAPTER 32 · THE BOUNDARY

Why the data never leaves, and why that is true

Many vendors say the data never leaves. In this architecture the sentence is true for a precise reason: the model receives tokens, not data, and re-identification happens only inside the client's perimeter.

WHAT THE OFFICER WRITES	Prepare the renewal letter for Mr Karim Haddad, file 2291, account 4471 0092.
WHAT THE MODEL RECEIVES	Prepare the renewal letter for `PERSON_01`, file `CASE_01`, account `ACCOUNT_01`.
WHAT THE OFFICER READS	The finished letter, with the real name, file and account restored inside the perimeter.

Illustration with invented data. The correspondence between tokens and real values is kept in a vault inside the client's perimeter, under keys the client holds.

The claim that a system keeps data inside the institution is made by almost everyone in the market. It is usually true in a weak sense: the data is stored in a region the client chose, under a contract that promises not to use it. It remains readable by the provider's systems, and it leaves the institution's control the moment it is sent to the model. This architecture makes the claim true in a strong sense, and it is worth explaining exactly how.

Tokenization, step by step

Before any text reaches the model, the gateway reads it and finds every element that identifies a person, a company, an account, a location or a case: names, identifiers, addresses, amounts tied to a person, references. Each element is replaced with a token, a code that means nothing outside the gateway. The model receives the text with the tokens in it and can reason about it perfectly well: it does not need to know whom a token stands for in order to draft a letter, check a document against a rule or summarise a file. When the answer comes back, the gateway replaces the tokens with the real values, inside the perimeter, before the answer reaches the person who asked.

What this changes

It changes who has to be trusted. In the usual arrangement the institution must trust the model provider with its data and rely on contracts. Here it does not have to trust the model with anything, because the model never sees anything worth protecting. This is why an external model can be offered as an option without weakening the guarantee.

The rest of the gateway

Tokenization is the first control, not the only one. The gateway also defends against prompt injection, that is, documents or messages constructed to make an agent do something it should not; it applies data loss prevention rules to everything that goes out; it requires human approval for actions with consequences; and it checks every answer before it is shown. Each control is a separate function that can be audited on its own.

THE SENTENCE, MADE PRECISE

The model receives tokens. The vault that maps tokens to data stays inside the perimeter. The keys are the client's. Re-identification happens only inside. Anyone can check each of the four statements.

CHAPTER 33 · TWO TRACKS

One architecture, two tracks

A proprietary track, in which every layer is built by the Department and installed on the client's own servers, and a Microsoft aligned track for institutions that already run on that stack. The security model is the same on both.

TRACK B · THE REFERENCE DESIGN

The proprietary track

FOR WHOM
Institutions that require the highest degree of sovereignty and independence, including classified and disconnected environments.

ORCHESTRATION
Engine developed in house by the Department.

IDENTITY, GATEWAY, OBSERVABILITY
Built with proprietary components.

WHERE IT RUNS
On the client's own servers, including fully disconnected environments, or on a secure cloud provided for the client.

MODEL
An open weight model installed locally, or an external model that receives tokens only. Replaceable.

DEPENDENCY
None on any single external vendor.

TRACK A · FOR ESTATES ALREADY IN PLACE

The Microsoft aligned track

FOR WHOM
Institutions already standardised on Microsoft 365, Entra, Azure and Purview, with trained staff and signed contracts.

ORCHESTRATION
Implemented on the Microsoft Agent Framework.

IDENTITY, GOVERNANCE, CONTENT SAFETY
Delivered by the Microsoft services the institution already operates.

WHERE IT RUNS
On Microsoft's sovereign cloud.

MODEL
Within the estate the institution has chosen. Replaceable.

DEPENDENCY
Confined to an estate the institution has already chosen and contracted.

SAME ON BOTH TRACKS The three tiers, the tokenization boundary, human approval and the single audit trail. An institution can begin on one track and move to the other, because the security model does not change.

The specification offers two ways of deploying the same architecture, and the point deserves a clear explanation because it is easily misread. The proprietary track is the reference design. The Microsoft aligned track is the same design fitted to an estate the client already owns. Neither is an integration of somebody else's product sold as the Department's own.

The proprietary track

The orchestration engine is developed in house. The identity layer, the security gateway and the observability are built with proprietary components, and the whole system runs on the client's own servers, including fully disconnected environments, or on a secure cloud provided for the client. The track removes any dependency on a single external vendor and is intended for institutions that require the highest degree of sovereignty and independence.

The Microsoft aligned track

For institutions already standardised on the Microsoft enterprise stack the architecture is implemented on the Microsoft Agent Framework, with identity, governance and content safety delivered by the services they already operate, and deployment on Microsoft's sovereign cloud. The track exists because those tools are already in place and the institution wants to adopt quickly without replacing them.

Why two tracks is the right design

Institutions do not start from the same place. A ministry that has spent years building on Microsoft, with trained staff and signed contracts, will not replace its identity system to adopt an agent platform. Offering it the same three tier architecture on the stack it already runs means it can start now. An institution that must not depend on any external vendor, or that operates classified or disconnected environments, needs the proprietary track.

What was verified

The second track relies on the Microsoft Agent Framework being generally available. It is: version 1.0 reached general availability at the beginning of April 2026, and at the Build conference of 2 and 3 June 2026 the agent harness and the multi agent orchestration patterns reached stable release.

THE READING TO KEEP

The code is the Department's, the servers are the client's, the model is replaceable.

CHAPTER 34 · THE AGENTS AT WORK

The four agents at work

The specification defines four kinds of agent for public bodies. The same four apply, with different rules, to private institutions: banks, corporate groups, family holdings.

AGENT ONE

Records agent

FOR A MINISTRY
Reads the institution's archives and registers, finds the right file, keeps the record of what was consulted.

FOR A BANK
Reads the client file, the account history and the correspondence, under the same access rules the bank already applies to its people.

AGENT TWO

Documents agent

FOR A MINISTRY
Drafts letters, decisions and reports from templates and from the file, in the institution's language, for a person to review.

FOR A CORPORATE GROUP
Drafts contracts, board papers and regulatory filings from the group's own precedents, never from the public internet.

AGENT THREE

Compliance agent

FOR A MINISTRY
Checks a file against the applicable rules before a decision is issued and flags what is missing.

FOR A BANK
Runs the checks that anti money laundering, know your customer and sanctions rules require, on tokenized data, and escalates to a person every case that needs judgement.

AGENT FOUR

Analysis agent

FOR A MINISTRY
Reads across thousands of files to answer a question the institution has, with sources.

FOR A FAMILY HOLDING
Reads across the group's positions, agreements and reports and answers questions about exposure, obligations and deadlines, with nothing sent to any external service.

Each agent is a specialised worker with a defined scope, and the orchestration engine combines them into a process. The panel shows what each does, first for a public body and then for a private institution.

A process, not a chat

Consider a request that reaches a ministry. The records agent finds the file and notes what it consulted. The documents agent drafts the reply from the institution's own templates. The compliance agent checks the draft against the applicable rules and flags what is missing. An officer reads the result and approves it. Four narrow steps, each with an output the institution already measures: a file found, a document drafted, a check completed, a decision taken by a person.

The same chain in a bank

The records agent assembles the client file. The compliance agent runs the required checks on tokenized data and escalates what needs judgement. The documents agent drafts the report. A compliance officer decides.

Why the same design serves both

A ministry and a bank differ in their rules, not in their structure. Both hold files about people, both must apply rules to those files, both must be able to show afterwards what was done and why. The architecture separates the rules, which are configured per institution, from the mechanism, which is the same. This is why the same system can serve a ministry, a bank, a corporate group or a family holding with only the rules changed.

Narrow by design

Narrow agents are easier to test, easier to trust with autonomy and easier to restrict when something goes wrong. Adding a fifth agent later adds an agent, not a platform: it inherits the gateway, the identity rules and the audit trail that are already there.

The scope, stated plainly

The specification is titled for public sector bodies. Nothing in the architecture is specific to the public sector. The title describes the first intended reader, not the limit of the design.

THE MEASURE OF VALUE

The time and the error rate of outputs the institution already measures, before and after.

CHAPTER 35 · THE PERSON IN THE LOOP

The person stays in the loop, and the agents are secured

The agent prepares, the person decides, the gateway records both. Behind that sentence stands a security discipline the Department already held before the architecture was written.

STEP ONE	STEP TWO	STEP THREE	STEP FOUR	STEP FIVE
The agent prepares	**The gateway checks**	**The person decides**	**The action runs**	**The gateway records**
Inside its scope, with its tools and its permissions.	Inputs filtered, outputs checked, data loss rules applied.	Sees the action, the data it touches and the chain that led to it.	Only after approval, where the institution's rules require one.	The preparation and the decision, in one audit trail.

None of the four agents issues a decision, sends a document or changes a record without a person's approval where the institution's rules require one. The specification treats human approval as a control, not as a limitation. This is the arrangement regulators expect and the one institutions accept.

Approval that is practicable

Approval is only real if the person can see what is being approved. Because every action passes through one tier, the person approving sees the action, the data it touches and the chain that led to it.

A discipline already held

The Presentation of August 2026 set out the security of artificial intelligence as a discipline in six dimensions, recorded in the register below. The controls it names for agentic autonomy are the ones the gateway applies.

Why agents raise the stakes

An AI system introduces exposures that conventional security programmes were not designed to see, and they grow sharply once models are given tools, autonomy and access to institutional systems. A model that only answers can be wrong. An agent that acts can be wrong inside the institution's records. The Department addresses the exposures as engineering problems, not as policy statements.

One record for the auditor

Every action of every agent passes through the orchestration tier and is recorded there. An auditor, a regulator or the institution's own compliance function reads one log, not one per application, and finds in it both what the agent prepared and what the person decided.

Why the intersection is rare

Securing an AI system requires knowing how it was built. The Department builds artificial intelligence from zero, and it attacks and defends what it builds, together with what others have built, in one perimeter and with the same engineers.

IN ONE LINE

The agent prepares, the person decides, the gateway records both.

EXPOSURE	THE ENGINEERING ANSWER
The model artefact	Weights and behavioural parameters are the entire training investment in a file that can be copied. A layered cryptographic regime protects them at rest, in transit and at inference, with split key custody and integrity verification at each call.
The training pipeline	Corpus poisoning and contaminated third party data are attacks against the model's future behaviour. The control is provenance: every source documented, every transformation reproducible.
The inference boundary	Instruction injection through documents, web pages, emails and tool outputs is the characteristic attack on assistant systems. The answer is architectural: separation between content and instruction, retrieval that preserves the user's own access rights, model text treated as untrusted input downstream.
Agentic autonomy	An agent with credentials and tools is an unattended user account with a reasoning engine attached. Least privilege per task, scoped and short lived credentials, human confirmation for irreversible actions, complete action logging, and a kill path that works.
Model behaviour	Adversarial testing of the model as a system: jailbreak and evasion campaigns, extraction and inversion attempts, guardrail regression suites at every release, with results recorded as evidence.
The deployment estate	Inference servers, vector stores, orchestration frameworks and edge devices are ordinary infrastructure with extraordinary value: hardening, segmentation, monitoring, patching and recovery.

CHAPTER 36 · EFFICIENCY

Where the efficiency of the design comes from

One gateway for every agent and every model. One place to audit. A model that can be swapped without touching anything else. Agents that do one job each. The efficiency is structural, not a matter of tuning.

Efficiency in a system of this kind is not measured only in speed. It is measured in how much has to be built once and reused, how much has to be changed when something changes, and how much has to be checked to know that the system is behaving. On all three counts the design is efficient for reasons that follow from its structure.

Built once, used by everything

The security gateway is a single component through which every agent and every model passes. Tokenization, injection defence, data loss prevention, approval and audit are implemented once and apply to every agent added afterwards. An institution that later adds a fifth or a tenth agent does not add a fifth or tenth security layer; the new agent inherits the existing one. In the usual arrangement, where each application carries its own AI integration, each one carries its own controls, its own gaps and its own audit.

What an institution avoids

The alternative is familiar to any institution that has bought AI features application by application: one integration inside the document system, another inside the case system, a third inside the help desk, each with its own supplier, its own data path and its own way of logging. Every one of them has to be assessed, contracted and audited separately, and none can be replaced without touching the application around it.

Changed in one place

Language models improve every few months. Here the model is a component behind the gateway, and replacing it, whether with a newer open weight model installed locally or with a different external service, does not touch the agents, the security or the identity layer. The institution keeps pace with the state of the art at the cost of a component swap, not a rebuild. The same holds for the identity system and for the observability tools.

Checked from one point

Every action of every agent passes through the orchestration tier and is recorded there. Governance, audit and observability are monitored from a single point. An auditor, a regulator or the institution's own compliance function reads one log, not one per application. This is what makes the system explainable.

Agents that do one thing

Each agent has a defined scope, defined tools and defined permissions. The orchestration engine composes them; it does not ask one general agent to do everything.

Explained simply

Think of a building with one guarded entrance instead of a guard at every office door. Adding an office does not add a guard; changing the locks is done once; and the visitors' book is a single book.

THE EFFICIENCY, IN THREE LINES

Security built once, inherited by every agent. Model swapped without a rebuild. One audit trail for the whole system. These are properties of the architecture, and they hold on both tracks.

CHAPTER 37 · ECONOMICS AND RISK

What the design means in money, risk and compliance

For the institution that adopts it, the architecture changes the shape of the cost, removes a category of risk, and turns compliance from a promise into a mechanism.

THE SHAPE OF THE COST

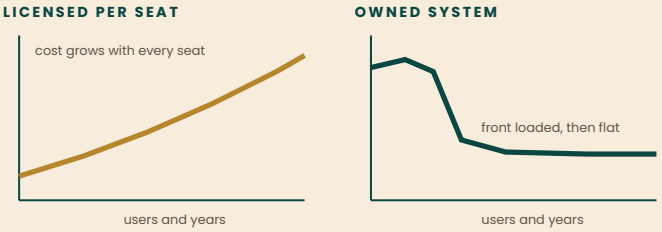

Schematic of the shape only, not to scale. No figures are implied.

RULES THAT ARE NOW ENFORCEABLE

14 Sept 2024
SAUDI PERSONAL DATA PROTECTION LAW FULLY ENFORCEABLE

EUR 35m or 7%
EU AI ACT: CEILING OF FINES FOR THE MOST SERIOUS VIOLATIONS

Sources: Saudi Data and Artificial Intelligence Authority; Regulation (EU) 2024/1689, Article 99.

This report puts no figures on the cost of a deployment, because the figures depend on each institution. What can be said with precision is how the cost is shaped, which risks disappear, and what the compliance position becomes.

The shape of the cost

In the licensed model an institution pays per user, per month, for as long as it uses the product, and at the end owns nothing. The model provider, the platform provider and the integrator each take a share, and each can reprice. Here the institution pays for the design and the construction of its own system, then for its maintenance and its infrastructure. The orchestration and the gateway are its property. The model, on the proprietary track, is an open weight model that costs nothing to license and runs on hardware the institution already budgets for. There is no fee per seat and no data transfer fee, because the data does not travel. The cost is front loaded and then flattens; it does not grow with the number of people using the system.

The risk that disappears

An institution whose intelligence layer sits on a single external platform is exposed to that platform's decisions: price, terms, continuity, jurisdiction. The proprietary track removes the dependency entirely; the Microsoft aligned track confines it to an estate the institution has already chosen. In both cases the model is replaceable, so no single model provider can hold the institution. The tokenization boundary removes a second risk: a breach or a misuse on the model side cannot expose the institution's data, because the model never had it.

Compliance as a mechanism

Regulators in the regions where the Group operates have moved from principles to enforceable rules. The Saudi Personal Data Protection Law has been fully enforceable since 14 September 2024. The European Union's AI Act is in force, with fines that reach 35 million euros or 7 per cent of worldwide turnover for the most serious violations. Financial regulators in the Gulf and in Europe require operational resilience and control over outsourced technology. Against rules of this kind a contract that promises good behaviour is weak evidence. A gateway that can show, for every request, that personal data was tokenized before leaving the perimeter, that a person approved the action and that the answer was checked, is strong evidence. The question changes from whether the institution trusts its provider to a log it can produce.

The fine, explained

Seven per cent of worldwide turnover means that a group with revenues of one billion euros is exposed to a fine of up to 70 million. The ceiling is whichever of the two amounts, 35 million or 7 per cent, is higher.

FOR THE READER IN A HURRY

Cost: owned, front loaded, flat. Risk: no platform can hold the institution, no model can leak its data. Compliance: demonstrated by mechanism, not asserted by contract.

CHAPTER 38 · WHY AGENT PROJECTS FAIL

Why agent projects fail, and where this design differs

Gartner expects more than 40 per cent of agentic AI projects to be cancelled by the end of 2027, for three reasons. The architecture is built against each of the three.

40%+	3	~130
AGENTIC AI PROJECTS CANCELLED BY THE END OF 2027, GARTNER FORECAST OF JUNE 2025	REASONS: ESCALATING COSTS, UNCLEAR BUSINESS VALUE, INADEQUATE RISK CONTROLS	GENUINE VENDORS AMONG THE THOUSANDS THAT CLAIM AGENT CAPABILITY

REASON GIVEN BY GARTNER	HOW IT ARISES	HOW THE DESIGN ANSWERS
Escalating costs	A stack of licences whose price grows with usage; each new use case needs its own integration and its own controls.	Security and orchestration are fixed at construction. The most volatile component, the model, can be an open weight one with no licence fee. Adding an agent adds an agent, not a platform.
Unclear business value	An agent deployed as a general assistant: nobody can say which process it improved.	Agents are defined by responsibility inside a process. Each output is one the institution already measures: a file found, a document drafted, a check completed, an answer with sources.
Inadequate risk controls	An agent that reads files and acts on systems without filtered inputs, checked outputs, approvals and records.	That list is a description of tier two. The controls are not added after the pilot; they are the layer through which the pilot runs.

In June 2025 Gartner published a prediction that has since been widely cited: over 40 per cent of agentic artificial intelligence projects will be cancelled by the end of 2027. The reasons it gave were escalating costs, unclear business value and inadequate risk controls. The register sets each reason against the feature of the design that answers it.

The second finding

Gartner added an estimate that matters as much as the cancellation rate: of the thousands of vendors claiming to sell agentic systems, only about 130 were selling the real thing, the rest having relabelled chatbots, assistants and process automation as agents. An institution comparing offers is, most of the time, comparing relabelled products, and the word agent on a brochure says nothing about what is behind it.

The test that separates the two

Where the orchestration runs, what the model receives, who approves, and what is recorded. The Department's specification answers the four questions on every page, and an institution can put the same four questions to any supplier. If the orchestration runs on the vendor's platform, if the model receives the institution's text as it is, if approval is a setting rather than a control and the record belongs to the vendor, the product is an assistant under another name.

A prediction, read correctly

The figure is a forecast made in 2025, not a measurement of 2026, and it is quoted here as such. Its value lies in the three reasons, which describe how projects fail rather than how many.

The consequence for a client

An institution evaluating agent systems is, in practice, choosing which of the three failure modes it is most exposed to. This architecture is one of the few in which all three have been designed out rather than managed afterwards. That is a claim about design, and it can be checked tier by tier.

IN ONE LINE

The controls are the layer through which the agents run.

CHAPTER 39 · TWO LINES, ONE DOCTRINE

Two lines of delivery, one doctrine

The sovereign build and the multi agent architecture answer different requirements. This chapter sets them side by side and states why they are consistent.

WHAT THE INSTITUTION REQUIRES	THE LINE THAT ANSWERS	WHAT IT HOLDS AT THE END
It must own the model itself: architecture, corpus and weights	Sovereign build from zero, eight stages	A system it owns, can audit, can operate and can continue without its builder
It must put agents to work on sensitive files and keep the data under its own control	Multi agent architecture, proprietary track	Orchestration and gateway as its own code, on its own servers, with a replaceable model behind a boundary whose keys it holds
It already runs on the Microsoft estate and will not replace it	Multi agent architecture, Microsoft aligned track	The same three tiers, boundary, approval and audit, fitted to the estate it already owns
It operates classified or disconnected environments	Either line, deployed inside the perimeter	A system that runs without external connectivity
It needs assurance on systems it did not build	Security and assurance, from the cyber pillar	Cryptographic architecture, penetration testing, red teaming and continuous rating

The Department now offers two lines of artificial intelligence work, and a careful reader of the Presentation of August 2026 will ask how they fit together. That Presentation states the doctrine of the sovereign build: every line of code written in house, models trained from scratch, no dependency on the external interfaces of commercial providers. The multi agent architecture admits an open weight model, or an external one, behind its gateway. The two statements are consistent, and this chapter says why.

Sovereignty has two objects

An institution can need sovereignty over the model, over the data, or over both. The sovereign build gives it the model: architecture, corpus, weights, documentation, the whole system owned and operable without its builder. The multi agent architecture gives it the data boundary: whatever model sits in tier three, the institution's data never reaches it, the keys never leave the client, and the code that orchestrates and secures the agents is the client's own.

The model as a commodity component

The August Presentation already recorded that the Department integrates commodity components where sovereignty is not at stake. In the multi agent architecture the model is such a component by construction: it is untrusted by design, it receives tokens and nothing else, and it can be replaced without touching the agents, the security or the identity layer. What made agent construction a Tier II activity in the market, data sent to a foreign model and work hostage to a provider's policy, is what the boundary and the replaceable model remove.

What does not change

Authorship. In both lines the Department writes what it delivers: the architecture and the models in the first, the orchestration engine, the security gateway and the agents in the second. On the Microsoft aligned track the same design is implemented on an estate the institution has already chosen, and the dependency is confined to that estate.

The two lines together

Because the model is replaceable, tier three can host any model, including one engineered from zero under the first line. An institution can start with agents on an open weight model and commission its own model later, without rebuilding anything above it.

ONE DOCTRINE

Sovereignty by construction, not by contract: of the model where the institution must own it, of the data boundary wherever agents touch its files.

CHAPTER 40 · THE MARKET

How large the demand for sovereign systems is

The market for artificial intelligence under national or institutional control is measured in tens of billions and grows at around 20 per cent a year. The demand comes from regulation, from geopolitics and from the institutions that hold the most sensitive data.

SOVEREIGN AI MARKET, US$ BILLION

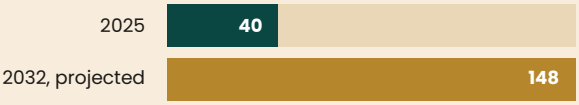

3.7x
GROWTH OF THE MARKET IN SEVEN YEARS

20.6%
COMPOUND ANNUAL GROWTH RATE

Source: MarketsandMarkets, Sovereign AI Market, global forecast to 2032.

DATA LEAKAGE AND COMPLIANCE AS FIRST DRIVER

Source: IDC InfoBrief sponsored by Cohere, The state of sovereign AI adoption, May 2026.

The figures in this chapter come from the public reports of market research firms and from a survey published in 2026. They are quoted as published, with their sources in chapter 47, and they should be read as orders of magnitude rather than as precise measurements: definitions of sovereign artificial intelligence vary between reports.

Size and growth

MarketsandMarkets values the sovereign AI market at 40 billion US dollars in 2025 and projects 148 billion by 2032, a compound annual growth rate of 20.6 per cent. North America is the largest region and Asia Pacific the fastest growing; growth in Europe is shaped by regulation; and the Gulf states are named among the sources of multibillion dollar sovereign programmes. Other firms publish different totals, some larger, and all describe growth at or above 20 per cent a year.

The growth rate, explained

A compound rate means that each year's growth is applied to the total of the year before. At 20.6 per cent a year a market multiplies by about 3.7 in seven years, which is how 40 billion becomes 148.

Who wants it, and why

An IDC survey of more than 500 senior decision makers at enterprises with revenues above one billion dollars, in Canada, the United States, the United Kingdom and Germany, was fielded in April and May 2026 and published by Cohere. Concern about data leakage and compliance is the first driver in every sector surveyed. The same survey found that one in three leaders had difficulty describing sovereign AI in their own words, and that only 13 per cent reported wide awareness of it inside their organisations. The demand exists; the understanding of what it requires does not yet.

The public sector and the Gulf

The Abu Dhabi government has committed 13 billion dirhams under its digital strategy for 2025 to 2027 to make its government AI native by 2027. Saudi Arabia has created a national champion, Humain, backed by the Public Investment Fund. Qatar has established Qai under the Qatar Investment Authority. The Emirates have the investment platform MGX and G42, which in February 2026 announced an agent factory for healthcare, education, legal services and energy at national scale. These programmes define what institutions in the region will expect: agents, under national control, at scale.

WHAT THE NUMBERS SAY TOGETHER

A large and fast growing market, driven by the institutions with the most sensitive data, in which most buyers cannot yet define what they need. The institutions that can explain the mechanism, not just the promise, will be the ones listened to.

CHAPTER 41 · WHO DELIVERS THIS TODAY

The organisations in the field, group by group

Every element of the architecture exists somewhere in the market. The question is who delivers all of them, as one system, on the client's own infrastructure, with an engineering organisation behind it.

GROUP	WHAT THEY DELIVER	CLIENT'S OWN CODE	CLIENT'S OWN SERVERS	MODEL REPLACEABLE
Hyperscalers	Sovereign cloud region, agent framework, security services	No	No	Within their catalogue
Sovereign model companies	Models and platforms installable on premises	No	Yes	Their models
Palantir	Complete platform, including disconnected environments, on open models	No, platform	Yes	Yes
National champions	National infrastructure and programmes	No	National	National models
Platform vendors, integrators	Configured combination of third party products	Configuration only	Sometimes	Sometimes
The Department's architecture	**Orchestration, gateway and agents built for the client**	**Yes**	**Yes, including disconnected**	**Yes, any**

To say how many organisations can deliver a system of this kind, the market has to be divided into the groups that actually exist, because they offer different things under the same words. Five groups cover the field.

The hyperscalers

Microsoft, Google and Amazon each offer a sovereign cloud and, increasingly, an agent framework. Their sovereignty is a region and a contract: the data stays in a country, on the provider's systems. It is a real improvement, and it is what the Microsoft aligned track uses. It is not the client's own infrastructure, and the security controls are the provider's services, not a boundary the client owns.

The sovereign model companies

Mistral and LightOn in France; Cohere and Aleph Alpha, which on 16 September 2026 signed a definitive agreement to combine, with headquarters in Toronto and Berlin; Domyn in Italy; Prem in Switzerland. They sell models and platforms that can be installed on premises or in disconnected environments, and this is valuable: they are the engines a proprietary deployment can use. They do not, as a rule, design and build the client's orchestration, gateway and agents as the client's own code. Some are tied to a hyperscaler: since July 2026 Mistral's models run in fully disconnected environments through Microsoft's Azure Local.

Palantir

The closest analogue. Its platform runs in classified and disconnected environments, and on 29 June 2026 it announced with NVIDIA an engine for running open models in sovereign environments for United States agencies. It is a complete system with real controls. It is also a proprietary platform on which the client builds and remains, concentrated on the United States government and its allies.

The national champions

G42 in the Emirates and Humain in Saudi Arabia build national infrastructure: data centres, national models and, in G42's case, an agent factory. They operate at the scale of a country's programme. They are not organisations that a single ministry, bank or corporate group engages to design its own system.

Platform vendors and integrators

Several hundred companies sell agent platforms, and consultancies and system integrators assemble them with a hyperscaler's services. Gartner's estimate of about 130 genuine vendors applies here. An integrator's deliverable is a configured combination of other companies' products: the client owns the configuration, not the engine.

HOW TO READ THE GRID

The comparison rests on what each group publicly says it delivers, not on tests. No claim is made that any competitor's system is insecure.

CHAPTER 42 · HOW MANY

The answer, stated carefully

Few, not none. Each element of the architecture is available somewhere. The complete combination, delivered by an engineering organisation as the client's own system, is held by a small number of organisations, and outside the United States by fewer still.

AVAILABLE TO MANY

What is common

Open weight models that can be installed locally.
Tokenization before inference, a pattern with published implementations.
Agent frameworks from every hyperscaler and from hundreds of vendors.
Sovereign cloud regions in every major market.

AVAILABLE FROM FEW

What is rare

Orchestration engine and security gateway written as the client's own code.
The whole system on the client's own servers, including disconnected environments.
A model that is replaceable and untrusted, behind a boundary whose keys the client holds.
An engineering organisation large enough to carry it, present in the region.

The question is whether it is true that few organisations offer this type of system. The answer has to be given carefully, because an exaggerated claim would be found out by the first technical reader, and an understated one would misrepresent the position. The answer is: few, not none, and the reason is the combination, not any single component.

What is common

Open weight models that can be installed locally are common and improving. Tokenization and de-identification of personal data before inference is a known pattern with published implementations. Agent frameworks are available from every hyperscaler and from hundreds of vendors. Sovereign cloud regions exist in every major market. Nobody should claim that any of these, taken alone, is rare.

What is rare

An organisation that designs and writes the orchestration engine and the security gateway as the client's own code, installs the whole system on the client's own servers including disconnected environments, treats the model as a replaceable and untrusted component behind a boundary the client holds the keys to, and does this with an engineering organisation able to carry it: a production floor of more than fifty certified engineers in Nepal working under an ISO 9001:2015 quality system, delivery experience for the World Bank, USAID, UNDP and the International Labour Organization, and an institutional relations team resident in the Gulf.

Who else

In the review of chapter 41, the only organisation delivering a comparable complete system in disconnected environments is Palantir, on its own platform, for the United States government and its allies. The model companies supply engines. The hyperscalers supply regions. The integrators supply configurations. The national champions supply countries.

Where this leaves an institution in the Gulf or in Africa

An institution in Riyadh, Abu Dhabi, Cairo or Lagos that wants a system of its own, on its own servers, with its own keys, and an engineering organisation present in the region to build it and stand behind it, has a short list. It can build a large internal team over years. It can adopt a hyperscaler's sovereign region and accept the dependency. It can license a platform built for another government. Or it can commission the system from one of the few organisations that build it as the client's own. That is the position, checked against what the market publicly offers as of September 2026.

THE CLAIM, AS IT CAN BE DEFENDED

Not that nobody else can do this. Rather: the elements are available to many; the complete system, built as the client's own and delivered on the client's own infrastructure by an engineering organisation with presence in the region, is available from few.

CHAPTER 43 · THE DELIVERY RECORD

The delivery record

Fourteen years of government grade and donor funded systems: the evidence behind the doctrine, and the class of system on which agents will work.

The institutional clients and funders behind the fourteen year delivery record of the engineering centre.

Behind the doctrine sits evidence a counterparty can inspect. For fourteen years the engineering centre has delivered systems for the World Bank, USAID through the Global Health Supply Chain Program, UNDP, the International Labour Organization, IFAD, GIZ, KOICA, JICA, the Japan Social Development Fund, the Swiss Embassy and Swiss Cooperation, and the Government of Nepal.

Systems a state runs on

Public finance through LMBIS and the Treasury Single Account; ILMIS, a labour management platform of thirteen modules; the eLMIS health logistics system across 58 sites; the Road Accident Information Management System across 65 sites; land records; provincial Chief Minister dashboards; an integrated tax system; centralised core banking; electronic building permits.

Why donor funded work is a hard school

Systems financed by multilateral donors are procured by competitive tender, delivered against written specifications, audited by the funder and accepted in writing by the government that will use them. They must run for years on the infrastructure of the public administration and be handed over to its staff. A centre that has worked this way for fourteen years has the habits an institutional client looks for: documents before code, tests before go live, transfer before exit.

WHAT THE LIST DEMONSTRATES

Not a portfolio of prototypes: systems through which a government collects taxes, pays salaries, manages land records and delivers services.

CHAPTER 44 · THE REFERENCE PORTFOLIO

The reference portfolio

Nine completed institutional engagements, each documented by an executed reference letter or a contractual completion document.

ENGAGEMENT	CLIENT OR FUNDER	REFERENCE
ILMIS: Integrated Labour Management Information System, thirteen modules, deployed on government data centre and G-Cloud with VAPT validation	Department of Labour and Occupational Safety, Government of Nepal	Letter ref. 688; contract DoLOS/2079/80/01, completed June 2023
Chief Minister's Dashboard, Pradesh 1: Provincial Support Programme	PwC AG under the Swiss Agency for Development and Cooperation	Services Procurement Agreement of 16 November 2022
Digital health profile of two municipalities, linked to social health protection	GIZ	Contract 83408972, July 2022 to December 2023
eLMIS rollout and support across fifty eight sites, levels 1 and 2	Chemonics International under the USAID Global Health Supply Chain Program	Reference letter of the Technical Director, 9 March 2022; contract GHSC-PSM-NPL-09
Office Automation System for the Embassy of Switzerland in Nepal	Federal Department of Foreign Affairs, Swiss Confederation	Contract executed 3 February 2020
Grievance Redress Management software, web based with SMS gateway, nationwide training	Department of Roads, Government of Nepal; World Bank funded	Letter of 19 December 2014
E-Commission Blueprint and ICT Roadmap; Case Management and Reporting System	National Human Rights Commission of Nepal; UNDP	Executed references, 2023 to 2024
PayPenny online money booking system	Connecting Creations Sdn. Bhd., Malaysia	Completion certification of 2 March 2021
National IT Roadmap of Nepal	Department of Information Technology, Government of Nepal	Executed reference, 2014

The register records each engagement with its client, its funder and the position of the reference. The documents themselves are reproduced in original facsimile in Part II of the Presentation of August 2026.

What a reference proves

A reference letter is signed by the institution that received the system, after the system went into use. It states what was delivered, under which contract and when. For a procurement authority it is the difference between a capability that is described and one that has been accepted.

Inside government infrastructure

Several of these systems run inside government data centres and government cloud, after vulnerability assessment and penetration testing. The eLMIS rollout under the USAID Global Health Supply Chain Program covered fifty eight sites at two levels of the supply chain.

Reading the register

The engagements span a decade, from the National IT Roadmap of 2014 to the labour management platform completed in 2023 and the references of the National Human Rights Commission in 2023 and 2024. They cover a national government, a provincial government, two bilateral cooperation agencies, a United Nations programme, a programme funded by the World Bank, a programme of USAID and a private client abroad.

What the systems have in common

Each of them holds records about people or public money, each applies rules to those records, and each had to satisfy someone other than its builder: a ministry, a donor, an auditor. ILMIS alone has thirteen modules. The grievance redress system for the Department of Roads combined a web platform with an SMS gateway and nationwide training of the officers who use it.

Why this matters for agents

Agents work on exactly this class of system: registers, case files, budget and labour platforms. An organisation that built such systems knows where the records are, how access is governed and what an auditor will ask.

THE RECORD IN ONE LINE

Fourteen years, nine executed references, national systems in production.

CHAPTER 45 · CREDENTIALS

Credentials and compliance perimeter

The quality credential under which every deliverable is produced, and the standards the Department builds into. What is held is stated as held; what is in formation is stated as in formation.

Official designation	R. Rockefeller S.C. AI Center — Nepal · Global Innovation Hub
Location	Lane-2, Kumaripati, Lalitpur, Nepal
Premises	5,000+ sq. ft. of dedicated operational space
Engineering team	50+ certified IT engineers and developers
Senior leadership	15 top managers and senior experts
Cumulative experience	200+ professional years across the senior team
Quality certification	ISO 9001:2015 (Quality Management Systems)
Operational history	14+ years of institutional track record
Activation date	5 May 2026

Credentials of the AI Center of Nepal and the centre at a glance.

QUALITY MANAGEMENT

ISO 9001:2015, certificate 77763/A/0001/UK/En issued by United Registrar of Systems under UKAS accreditation. Scope: Providing Software Development Services. Held by the engineering centre since 2017.

CRYPTOGRAPHIC CUSTODY

FIPS 140-3 Level 3 hardware security modules for key custody; Common Criteria EAL 4+ as implementation class.

POST QUANTUM READINESS

Alignment with the NIST programme; hybrid classical and post quantum key wrapping; EuroQCI compatibility perimeter.

ENGINEERING CREDENTIALS

Oracle Certified Professional, Microsoft Certified IT Professional, Cisco CCNP and CCNA, Red Hat RHCE across the senior bench.

REGULATORY PERIMETER

The EU AI Act, GDPR, NIS2, DORA, the Cyber Resilience Act and eIDAS in Europe; the Saudi Personal Data Protection Law; the NIST AI Risk Management Framework for North American counterparties.

AI MANAGEMENT SYSTEM

ISO/IEC 42001 is being implemented on the same lifecycle artefacts. It is reported as in formation, not as held.

Compliance frameworks are often presented as friction. Read correctly they are a specification, and they describe the system the Department already builds: documented data provenance, technical documentation adequate for assessment, logging and traceability, human oversight designed into the workflow, and an operator who can answer for the system because the operator built it.

HELD, AND IN FORMATION

Certified: the quality management system. In formation: the AI management system of ISO/IEC 42001.

CHAPTER 46 · THE RESEARCH PROGRAMME

The research programme, stated as it stands

One page, and only one. The programme is research conducted by the Group for its own account. It is not a service of the Department, and nothing in this report depends on it.

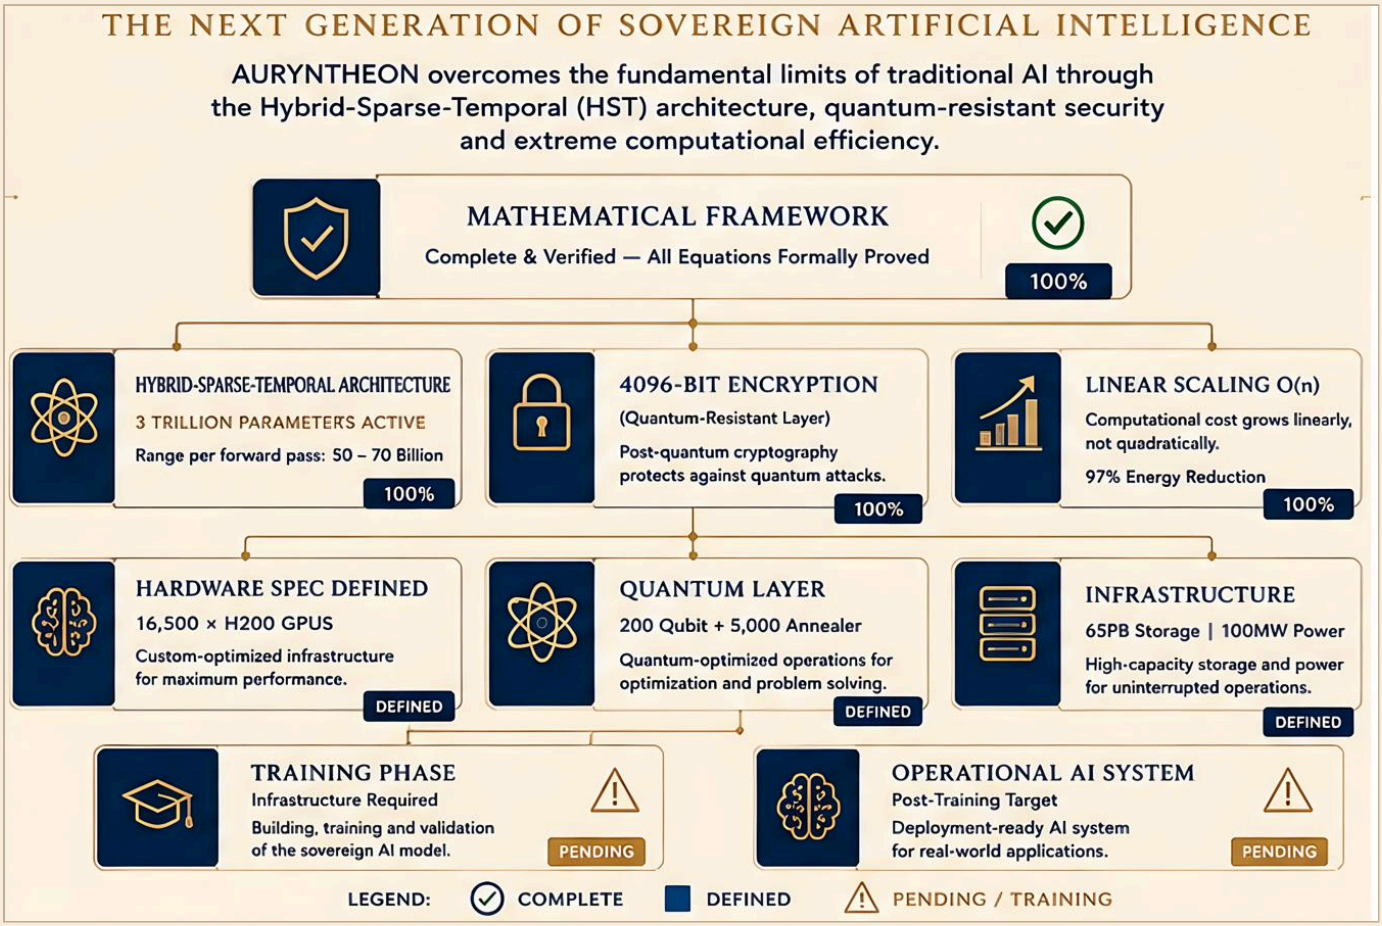

The status of the research programme, component by component: complete, defined, pending.

In parallel with its production work the Group conducts a proprietary research programme on advanced AI architectures, designated AURYNTHEON, under the scientific authority of Mr Umberto Colella.

Where it stands

The mathematical framework is complete and proved analytically, which places the programme at TRL 3 on the European scale, reported conservatively. The design is dimensioned at three trillion parameters, of which fifty to seventy billion are evaluated for each query, with linear scaling in the length of the text and a 4096-bit quantum resistant encryption layer. The infrastructure the training phase requires is defined and not yet acquired. The training phase and the operational system are pending and are expressly not claimed.

Explained simply

Three trillion is the size of the whole model; fifty to seventy billion is the part that works on any one question, about two per cent. The smaller that share, the less hardware and energy each answer costs.

Why it is mentioned at all

Because it shows the depth at which the scientific authority of the Department works, and because an honest account of the Department has to say what it researches as well as what it delivers. The production lines of Parts V and VI do not wait for it.

THE RULE

What is mature is evidenced, what is in formation is declared as such, and what belongs to a later phase is not claimed in advance.

CHAPTER 47 · VERIFIED FACTS

What was checked, and what was found

Every figure and every statement about third parties in this report was checked against a public source in September 2026. This chapter lists them, so that any reader can repeat the check.

Statements about the Department come from its Presentation of August 2026, from the technical specification of the multi agent architecture and its analysis of September 2026, and from the official figures of the Group. Statements about third parties come from the sources below.

Microsoft Agent Framework

Version 1.0 generally available from the beginning of April 2026: Microsoft dates general availability to 2 April, and its release post is dated 3 April. At the Build conference of 2 and 3 June 2026 the agent harness and the orchestration patterns reached stable release; the harness and hosted agents have since reached general availability. Sources: Microsoft Agent Framework blog; InfoQ, August 2026.

Gartner prediction

Over 40 per cent of agentic AI projects cancelled by the end of 2027, for escalating costs, unclear business value or inadequate risk controls; about 130 genuine vendors among thousands. Source: Gartner press release, 25 June 2025. It is a forecast of 2025 and is quoted as such.

Sovereign AI market

USD 40.0 billion in 2025, USD 148.0 billion in 2032, compound annual growth of 20.6 per cent for 2026 to 2032. Source: MarketsandMarkets, press release of 19 August 2026 and report page. Other firms publish different totals.

IDC survey

More than 500 senior decision makers at enterprises above one billion dollars in revenue, in Canada, the United States, the United Kingdom and Germany, April to May 2026; 82 per cent of financial services respondents put data leakage and compliance first; one in three cannot define sovereign AI; 13 per cent report wide awareness. Source: IDC InfoBrief sponsored by Cohere, document EUR254834226-IB, May 2026. The survey does not cover the Gulf.

Palantir and NVIDIA

Engine for running NVIDIA Nemotron open models in sovereign environments, for United States government agencies and critical infrastructure, including classified and disconnected environments. Announced on 29 June 2026. Source: Palantir press release.

Cohere and Aleph Alpha

Planned combination announced in April 2026; definitive agreement signed on 16 September 2026; headquarters in Toronto and Berlin; subject to regulatory approval. Source: joint press release.

Mistral and Microsoft

Expanded partnership announced in July 2026: Mistral models available in cloud, cloud connected and fully disconnected environments through Azure and Azure Local. Source: joint announcement.

G42 agent factory

Announced in February 2026 at the World Governments Summit, for healthcare, education, legal services and energy, at national scale. Source: The National, Abu Dhabi.

Gulf national programmes

G42 and the investment platform MGX in the Emirates; Humain in Saudi Arabia, backed by the Public Investment Fund; Qai in Qatar, under the Qatar Investment Authority. Source: CNBC, 24 May 2026.

Abu Dhabi digital strategy

13 billion dirhams for 2025 to 2027; government to be AI native by 2027. Announced on 21 January 2025. Source: Abu Dhabi Media Office.

Saudi PDPL and EU AI Act

The Saudi Personal Data Protection Law has been fully enforceable since 14 September 2024, under the Saudi Data and Artificial Intelligence Authority. The EU AI Act provides fines of up to 35 million euros or 7 per cent of worldwide annual turnover, whichever is higher, for the most serious violations. Sources: SDAIA; Regulation (EU) 2024/1689, Article 99.

Sovereign model companies

Deployment options as described by the companies themselves. These are vendor descriptions, not independent tests.

WHAT IS NOT CLAIMED

No figure is given for the cost of a deployment, because it depends on the institution. No claim is made that any competitor's system is insecure. The multi agent architecture is presented as specified; no client deployment is cited in this report. The research programme is reported at TRL 3 and its training phase is not claimed.

CLOSING

Closing considerations

What this report has documented, and the position from which the Department operates.

The report has documented an institution of two pillars. The first builds artificial intelligence software from zero, for any sector, inside one perimeter and under one quality system, and hands the result to the institution that commissioned it: secured, documented, operable and owned. The second is a complete cyber practice that protects what the first builds, together with what others have built.

What was added in 2026

A second line of AI work: the sovereign multi agent architecture. Three tiers, one security gateway through which every agent and every model passes, a tokenization boundary that keeps the institution's data away from the model, two deployment tracks, and four classes of agent that prepare work for a person to decide. It was designed against the three reasons for which agent projects fail, and it extends to private institutions what was first written for public bodies.

Where the strength lies

Not in size, and not in capital. In authorship: the Department writes what it delivers, and only the author can transfer a system, answer for its behaviour, secure it where it is actually exposed and continue it when the environment changes. In independence: the Group has grown on partnerships, without external equity dilution and without debt. In the capacity to refuse: no disclosure of its architecture, no pilots before validation, no claims it cannot evidence. And in a record that can be inspected.

What the Department is

A concentrated institution, not a large one: modest in headcount by the standards of the integrators, deep by the standards of its size. Very few institutions build artificial intelligence from zero at all. Fewer still also secure it, and fewer still can place agents on an institution's most sensitive files without the data ever reaching the model.

For whom this matters

For an institution that holds sensitive files and is being offered agents: a ministry, a regulator, a bank, a corporate group, a family holding. The report gives it the questions to put to any supplier, and the Department's answers to each of them.

IN ONE LINE

Authorship, independence, the capacity to refuse, and a record that can be inspected.

CONTACTS

Executive office	Mr Valentino Loforese, Chief Executive Officer and Majority Shareholder, R. Rockefeller S.C.
Scientific and technical authority	Mr Umberto Colella, Chief of Staff and Lead AI System Architect
Team email	team@therobrockefeller-sc.com
Web	therockefeller-sct.com
Registered offices	8 The Green, Ste R, Dover, DE 19901 · 1309 Coffeen Ave, Ste 1200-C, Sheridan, WY 82801
Registration	EIN 93-4657669 · Delaware LLC Reg. No. 6714331