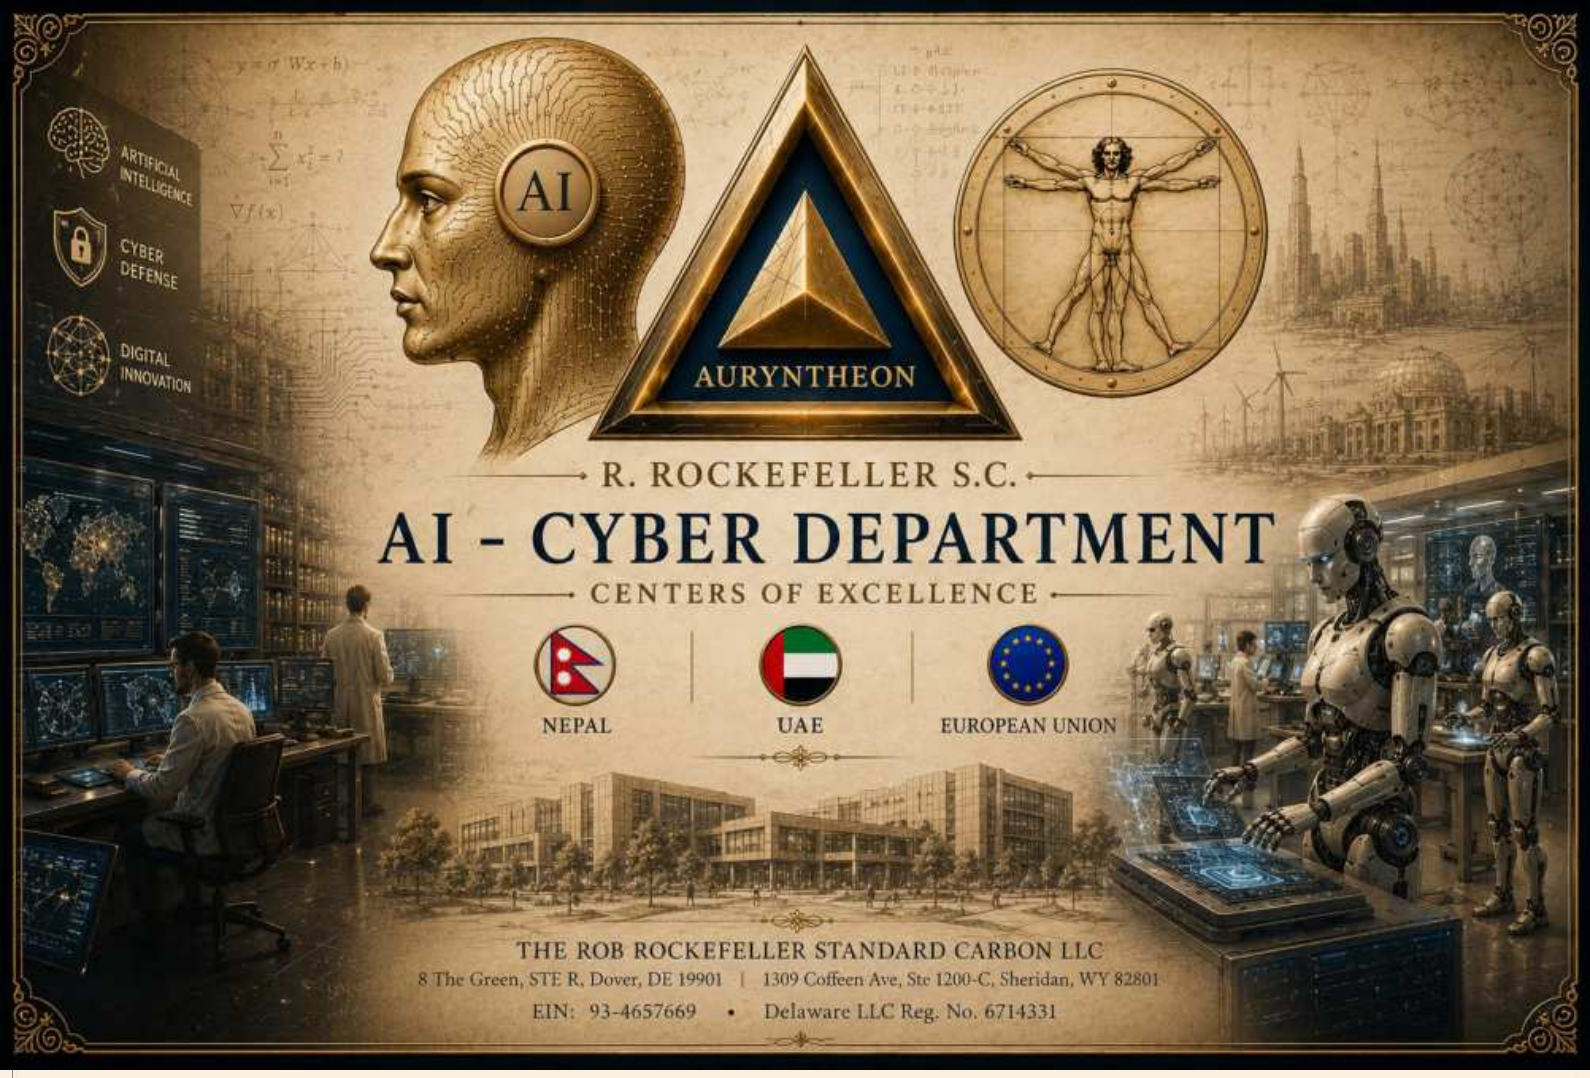

AI & CYBER DEPARTMENT PRESENTATION

COMPLETE EDITION · ISSUE v1.0 · AUGUST 2026

The two pillars of the Department, in full. First pillar: the production of artificial-intelligence software from zero, for any sector, under sovereign control. Second pillar: the complete cyber portfolio, offensive, defensive, architectural, cryptographic and regulatory, delivered across government, defence, energy, industry, finance, health, telecommunications, robotics and critical infrastructure. With the technology base, the engineering execution and the operating poles of Nepal, Italy, Bulgaria and the United Arab Emirates.

ISSUED BY THE OFFICE OF THE CHIEF EXECUTIVE OFFICER

FOR INSTITUTIONAL REVIEW

"Artificial Intelligence will liberate humanity, and Humans will finally become masters of their own time."

MR. VALENTINO LOFORESE · CHIEF EXECUTIVE OFFICER AND MAJORITY SHAREHOLDER

PURPOSE AND SCOPE

Purpose of the Present Presentation

ISSUING INSTITUTION	**The Rob Rockefeller Standard Carbon LLC** A Delaware limited liability company · Office of the Chief Executive Officer
ADDRESSEES	Institutional counterparties, governmental and enterprise clients, technology partners and prospective local partners of the Group
EDITION	Complete Edition · Issue v1.0 · August 2026

Subject

Presentation of the AI & Cyber Department of The Rob Rockefeller Standard Carbon LLC, documenting in full the two pillars on which the Department rests. The first is the production of artificial-intelligence software from zero, for any sector, end to end and under sovereign control, with the doctrine, the technology base and the deployment security architecture that support it. The second is the complete cyber portfolio: offensive security, vulnerability and exposure management, security architecture, cryptographic engineering, detection and response, application and cloud security, operational technology and critical infrastructure, third-party risk, governance and compliance, capability building, and the security of artificial intelligence itself, applied across every sector the Department serves. The Presentation documents further the certified engineering execution centre in Nepal, the European poles in Italy and Bulgaria, the Augusto Innovation Hub programme of Milazzo, the QUASAR cooperation in Sofia, and the institutional architecture of the Group in the United Arab Emirates. The Presentation closes with the measure of the institution and the position from which the Department operates.

Issuing institution

The Rob Rockefeller Standard Carbon LLC · a Delaware limited liability company · Chief Executive Officer and Majority Shareholder: Mr. Valentino Loforese · Italian Branch, Milan: Mr. Francesco Cuccurullo, Italian Branch Administrator.

Scope note

The present Presentation is confined, deliberately, to the perimeter of the AI & Cyber Department: its two pillars and the poles through which they are exercised. The Group operates further strategic verticals; they are not treated here. The corporate position of the Group in the United Arab Emirates touching an entity of another vertical, the sustainability joint venture presided over by H.E. Sheikh Khalifa Bin Khalid Al Hamed, is stated, once and with precision, in Section XVIII, together with the Group's openness, eventually and separately, to a dialogue in that dimension as well. That position is stated once, for completeness, and is not developed further here.

FOREWORD

A Note on the Present Presentation

The Presentation introduces itself in three movements: the institution writing; how the Group reads the Emirati opportunity, and why it has chosen to write; and the architecture of the document and the reading discipline it presupposes.

THE PRESENT DOSSIER IS THE institutional dossier of the AI & Cyber Department of The Rob Rockefeller Standard Carbon LLC. It does not seek to persuade by rhetoric: it documents, on evidence and for the institutional uses of the Group, what the Group is, what it holds, where it operates, what it is building in Europe, and what it now intends to build in the United Arab Emirates. Every claim in this document is either verifiable today, on the certifications and executed references reproduced in Part II, or stated expressly as being in formation.

I. The institution writing

The Rob Rockefeller Standard Carbon LLC (“RRSC” or “the Group”) is a Delaware-domiciled, multi-vertical institutional platform, anchored to royal and ministerial partnerships at sovereign level, governed under a unified shareholder structure and an irrevocable unified governance, and built to date entirely on partnership-funded growth: without external equity dilution and without debt. It operates through twenty-four active centres across seventeen sovereign jurisdictions on five continents, a network in continuous expansion. Its institutional posture rests on partnerships at the highest level: in the United Arab Emirates, H.E. Sheikh Khalifa Bin Khalid Al Hamed, of the Al Hamed ruling family of Abu Dhabi and grandson of H.E. Ahmed Khalifa Al Suwaidi, the founder of the Abu Dhabi Investment Authority, presides over RRSC UAE; in the Kingdom of Saudi Arabia the Group operates under a formal Royal Partnership with H.H. Prince Abdul Majeed bin Saud bin Muhammad Al Saud of the Saudi Royal House; in the State of Qatar the institutional anchor is Sheikh Tameen Al Thani of the Al Thani Royal Family.

Within this platform, the AI & Cyber Department is the flagship, and it stands on two pillars of equal weight. The first is the 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, exercised across the whole discipline, from offensive testing and security operations to cryptographic engineering, operational-technology protection and regulatory assurance, for institutions in every sector. Both are exercised under the technical authority of Mr. Umberto Colella, Chief of Staff and Lead AI System Architect, through the Group's own R. Rockefeller S.C. AI Center of Lalitpur, Nepal, the Global Innovation Hub; anchored in Europe on the Italian Branch of Milan and Brindisi and on the Sofia scientific cooperation with a European-Union-funded centre of competence of the Bulgarian Academy of Sciences; consolidated by the Strategic Alliance signed in 2025 with SecurityScorecard, the global leader in cybersecurity ratings; and now extending to the Augusto Innovation Hub of Milazzo, Sicily, where the Group is technology partner and technical leader of one of the most significant sovereign-innovation campuses in development in Southern Europe.

II. The Group and its counterparties

The Group's presence in the United Arab Emirates is institutional before it is commercial: a royal anchor within the sovereign perimeter of Abu Dhabi, and a resident senior leadership whose members hold concurrent mandates inside the Emirati sovereign-wealth, federal ministerial and national-energy structures. That architecture stands, fully constituted and active, and it is the diplomatic and institutional channel through which the Group operates in the Gulf. What the Group now intends to add is the technical and operational component: a permanent operational centre through which its engineering capability is channelled into the region, established in partnership with a local partner of the right calibre and, decisively, of the right sector.

FOREWORD · CONTINUED

The Counterparties, and How to Read

Through its representatives and technical management, the Group has received and is evaluating several partnership proposals for the Emirates; none of them originates from the artificial-intelligence and robotics sector. The Department is presented here as it is, to the institutions that engage it: governmental administrations and agencies, enterprise and industrial clients, scientific and defence counterparties, and the technology institutions with which it works. The present Presentation places before the reader, completely and on evidence, what the Department is in Nepal, in Italy, in Bulgaria and in the Emirates, what it engineers, how it secures what it engineers, and the standard of proof it applies to its own statements. It documents a position. It does not solicit one.

III. The architecture of the Presentation

The Presentation is organised in two Parts. **Part I** opens with the Statement of the AI Partner: a structured account, in seven points, of what the Group brings, the status of its technology base, its people, its operating model, its governance, its commercial doctrine and the durability of the value. It proceeds through the Group's specialty in the production of artificial-intelligence software from scratch, documented in its own right across the production doctrine of eight stages, the engineering stack and model classes, the five tiers at which the market operates and the comparative landscape of the institutions able to build from zero, the technical positions the Group has taken in depth, the regulatory and cryptographic perimeter it builds into, and the state of the sovereign-AI market as at August 2026, and then through the doctrine of sovereign AI, the proprietary technology base, the deployment security architecture and the cybersecurity vertical, and then through the operating poles of the Group, presented each in its own Section: Nepal, the engineering execution centre, certified and referenced across fourteen years; Italy, the European hub of Milan and Brindisi and the Augusto Innovation Hub of Milazzo; Bulgaria, the Sofia AI and Cyber hub within the European defence and quantum perimeter; and the United Arab Emirates, presented in three Sections covering the institutional architecture, the Emirati AI and Cyber Hub, and the corporate position and the present opening. Part I closes with the strategic global context, the Group's statement of separation on the People's Republic of China, the executive briefing on the sovereign question, the measure of the Group's strength, and the closing considerations. **Part II** reproduces, in original facsimile, the institutional verification materials of the engineering execution centre: the ISO 9001:2015 Certificate of Registration and ten executed reference letters and contractual completion documents.

The reading discipline. The discipline is the one the Group applies to itself: technology readiness is reported conservatively, on the EU Technology Readiness Level scale; what is mature is documented; what is in formation is stated as such; and what the Group, by method, does not claim in advance, it does not claim here. The Group prefers to under-state and over-deliver than the converse.

A note on the protected substance. The protected substance of the Group's proprietary sovereign AI architecture, including its mathematical specification and its architectural documentation, is not disclosed, under any instrument, to any counterparty, and does not form part of any review perimeter. The programme is presented, here as everywhere, at capability level; what a partnership draws on is the programme's outputs, not access to the architecture.

Nepal: we are the engineering. Italy: we are the European hub, and Milazzo is what we will build next.
Bulgaria: we are inside the European defence and quantum perimeter. The Emirates: we are the institution.
What is missing in the Emirates is one thing only: the operational centre.
The Group is willing to evaluate strategic cooperations to establish it, and to extend the same method elsewhere.

CONTENTS

Detailed Index of the Presentation

PART I · THE INSTITUTIONAL STATEMENT AND THE OPERATING POLES

PART II · SUPPORTING DOCUMENTS

ON REQUEST.....

Note. Page numbers refer to the consolidated Presentation. The Presentation is issued as a single controlled document; individual Sections are not designed for circulation in isolation.

READER'S GUIDE

How to Read this Presentation

The Presentation is designed for reading in two distinct passes, calibrated to the institutional posture of the reviewer. The two passes are not sequential: a senior decision-maker may consult only the first; a technical advisor may begin directly from the second.

First pass · Institutional review · for principals and senior decision-makers

#	DOCUMENT	PURPOSE
1	**The Statement of the AI Partner**	The capabilities of the Group and the position from which it engages its counterparties
2	**Sections XVI to XVIII · The Emirates**	The institutional architecture, the Emirati AI and Cyber Hub, and the present opening
3	**Section XII · The Augusto Innovation Hub**	The Group's next European flagship: what will be built at Milazzo
4	**Section XV · The Sovereign Question**	Concentrated executive briefing on the global context and the strategic stakes
5	**A Statement of Separation · The China Perimeter**	What the Group does and does not accept in China, and how the separation is enforced
6	**The Measure of the Institution**	Where the strength of the Group actually lies, stated so that it can be tested
7	**Closing Considerations**	What the Presentation has documented, and the position from which the Department operates

Second pass · Technical and operational review · for technical and financial advisors

#	DOCUMENT	PURPOSE
1	**The Production Doctrine and the Engineering Stack**	How AI software is built from zero, stage by stage, and with what
2	**The Five Tiers and the Comparative Landscape**	Where the market operates, who builds from zero, and how capabilities coincide
3	**The Engineering in Depth and the Compliance Perimeter**	The technical positions taken, and the regulatory and cryptographic frameworks built into
4	**Section I · The Doctrine of Sovereign AI**	What sovereign AI means in the Group's engineering practice
5	**Section II · Technology Base**	The proprietary architecture programme: posture, validation position, conditions of disclosure
6	**Section III · Deployment Security**	Threat model, layered cryptographic regime, custody model, standards alignment
7	**Sections VI to X · Nepal**	The engineering execution centre: certification, references, leadership, senior roster
8	**Section XIV · Global Context**	The architectural convergence of the Chinese national programme and the Western vacuum
9	**Part II · Supporting Documents**	Original facsimiles of the verification materials

THE STATEMENT

The Statement of the AI Partner

The capabilities that The Rob Rockefeller Standard Carbon LLC brings to its institutional cooperations and client programmes, set out in seven points, after a clarification the Group considers necessary before entering the body of the Statement.

The present Statement is not a proposal in the promotional sense and does not seek to persuade: it documents, for the institutional uses of a cooperation or a client programme, why the Group is in a position to do what such an engagement would assign to it. It is organised in seven points: what the Group brings; what the status of its technology base actually is; who the people are; how the work proceeds; how it is governed; the commercial doctrine under which it operates; and why the value endures. Each point is answered on evidence.

Preliminary Clarification On the nature of the Group's principal technology asset.

The flagship asset of the Group's artificial-intelligence vertical is a frontier sovereign AI architecture of globally strategic class, held under the institutional sovereignty of the Group, designed and authored by our Chief of Staff and Lead AI System Architect. It is not a product, not a platform, and not a research project of the kind that is examined through a written questionnaire. The number of individuals in the world today qualified to assess such an asset on its technical merits can be counted on the fingers of one hand. This is a statement of fact about the field, not a posture, and it carries two consequences for the reader of this Presentation.

The first is that an asset of this class is not verifiable through external review of documents. Even were the entirety of the internal documentation, mathematical specification and engineering codebase to be disclosed in writing, which it will not be, the disclosure would not place the recipient in a position to render a technical judgement on it. An asset of this class is assessed by working alongside the principal author, over time, within a confidentiality framework. The second is that the asset is not patentable in any meaningful sense, and the protection regime under which it is held is, accordingly, the trade-secret regime applicable to assets of this class at the frontier of the field, supplemented by copyright on the implementation and by the contractual confidentiality architecture that surrounds every engagement with the technology.

One further clarification is owed to the reader, and it concerns scope. The architecture programme is a strategic reserve of the Group in the most literal sense: it is cited in this Presentation because it establishes, better than any other single fact, the caliber of the institution the partner would join. The Group holds, within this programme, an original and proprietary mathematical formulation whose stability, convergence and sparsity properties have been demonstrated analytically: a result of extreme rarity, and, as documented in Section XIV, precisely the element that the Chinese national programme, for all its scale, does not possess. The programme is in an active phase of research and evolution: advanced institutional dialogues for its development are under way within the European Union perimeter, with programmes of the North Atlantic alliance, and with other European centres, concerning in particular the training phase; in the European defence sector, a formal scientific cooperation has been executed in Sofia with a European-Union-funded centre of competence of the Bulgarian Academy of Sciences, and that cooperation is now advancing to its second phase. In parallel, the Group holds its own private plans for the self-financing of that phase, so that the choice between the institutional path and the private one is made by the Group on the criterion of independence alone. The programme is not part of the perimeter of any cooperation contemplated by this Presentation: the Group conducts it alone, for its own account, and is not seeking partners for it; the involvement of a third party, if it were ever to occur, would be a separate and sovereign decision of the Group. For a prospective partner it is added value: the demonstration of the level at which the institution operates. In documents of the present class, the asset is therefore described at capability level; its mathematics and its architecture are not disclosed, in any form, to any counterparty.

THE STATEMENT · CONTINUED

Points 1 and 2

1 · Contribution of RRSC

1.1 The technology axis. RRSC brings to its cooperations one converging technology axis, proprietary to the Group's perimeter: end-to-end sovereign artificial-intelligence engineering. This is the capability to design, train, secure, deploy and operate AI systems for strategic verticals under conditions of sovereign control, without third-party commercial APIs, foreign pre-trained models or extraterritorial cloud exposure, exercised through the R. Rockefeller S.C. AI Center of Lalitpur, Nepal, the Group's Global Innovation Hub, and the wider engineering network. The Group is one of the very few institutions in the world able to produce AI software from scratch, for any sector; this production capability is the substance of what it brings, with the architecture programme of Section II standing behind it as a strategic reserve, outside the cooperation perimeter. Around this axis sits the cybersecurity vertical of Section IV: cryptographic infrastructure, deployment security, penetration testing and supply-chain security discipline.

1.2 Existing validation and institutional recognition. The engineering execution centre operates under ISO 9001:2015 certification (United Registrar of Systems, certificate 77763/A/0001/UK/En, current cycle valid through 22 August 2026, scope Providing Software Development Services). The delivery record spans fourteen years of government-grade and donor-funded systems for the World Bank, USAID, UNDP, GIZ, the Swiss Agency for Development and Cooperation, and the Government of Nepal, documented in Sections VII and VIII and reproduced in original form in Part II. The cybersecurity vertical is anchored on the Strategic Alliance signed in 2025 with SecurityScorecard, the global leader in cybersecurity ratings, executed directly between Mr. Valentino Loforese and SecurityScorecard's CEO and co-founder Dr. Alexander Yampolskiy.

1.3 Engineering execution capacity. Engineering implementation is carried out through the Lalitpur centre: more than fifty certified IT professionals, thirteen Senior Key Resources formally recognised as members of the RRSC global professional network, cumulative experience exceeding two hundred professional years, and a multi-year delivery record on government-grade systems including the eLMIS rollout across fifty-eight sites under the USAID Global Health Supply Chain Program. The Gulf institutional leadership, presented in Section XVI, anchors the Group in the Emirates. The Italian Branch, in Milan and Brindisi under Mr. Francesco Cuccurullo, anchors the European perimeter.

2 · Status of the Technology Base

We report technology readiness conservatively, on the EU Technology Readiness Level scale. We prefer to under-state and over-deliver than the converse.

ASSET	TRL	EVIDENCE POSITION
Sovereign AI architecture (integrated system)	TRL 3	Analytical proof of the architectural claims complete: the proprietary mathematical formulation demonstrates stability, convergence and sparsity properties in advance of empirical scale-up; component-level validation by peer-reviewed precedent at sub-scale
Deployment security architecture	TRL 3	Analytical and experimental proof of concept; engineering specification complete at the engineering level
Engineering execution capacity (sovereign AI software and platforms)	Mature	Certified under ISO 9001:2015; referenced across fourteen years of institutional delivery; in daily production

There are no live pilot implementations of the integrated sovereign AI system at this date, and the Group does not introduce pilots before the prototype phase has been closed and the architectural claims empirically validated at scale. The introduction of pilots before validation would compromise both the credibility of the programme and the interests of any partner involved. This discipline is a feature of the Group's method, not a limitation of it.

THE STATEMENT · CONTINUED

Points 3 and 4

3 · Team and Expertise

3.1 Scientific and Technical Lead. Mr. Umberto Colella, Chief of Staff and Lead AI System Architect of the Group, is the Scientific and Technical Lead of the sovereign AI programme and the originating author of its architecture. He holds Bachelor's and Master's degrees in Software Engineering from the University of Technology Sydney, with Master's specialisation in Machine Learning and a final thesis on a proprietary 4096-bit encryption algorithm. Through the University's competitive international exchange programme he completed selective engagements at Google, Sony Corporation Tokyo and Microsoft Corporation Redmond, followed by cryptographic engineering of institutional cross-border settlement protocols at Euroclear. A substantial part of his subsequent record concerns dual-use and national-security programmes for countries of the North Atlantic alliance perimeter, held under confidentiality and, in part, classification obligations; his full profile is provided in Section V at the level of competence area rather than specific deliverable.

3.2 Engineering execution team. The thirteen Senior Key Resources of the Lalitpur centre, formally recognised as members of the RRSC global professional network under the corporate instrument of 5 May 2026, constitute the operational foundation of the engineering capacity available to any partnership; their condensed profiles are provided in Section X, within a facility staffed by more than fifty certified IT engineers and fifteen senior managers.

3.3 The institutional leadership. Beyond the engineering bench, the Group's institutional leadership completes the operational perimeter on which its cooperations draw: in the Emirates, a royal presidency and four senior resident executives, each carrying an additional institutional mandate inside the Emirati establishment (Section XVI); in Europe, the Italian Branch leadership and the Bulgarian directorate (Sections XI and XIII). For any new centre, senior technical positions are staffed through a combination of Group secondment, the Lalitpur engineering bench, and local recruitment through the channels of the local partner.

4 · Operational Model

4.1 The three-pole delivery architecture. For the sovereign technology programmes, delivery is organised on a three-pole architecture: Lalitpur for engineering execution, codebase production, training pipelines, deployment automation and quality assurance, with scientific direction through Mr. Umberto Colella; Europe for sovereign positioning, anchored on the Italian Branch, on the Sofia cooperation, and, in formation, on the Augusto campus; and the Emirates for institutional channelling of the Middle East and African perimeter. A new operational centre in the Gulf completes this architecture: the production and representation seat through which the Group's AI and cybersecurity capabilities are delivered into the region.

4.2 The phased method for a new centre. The Department establishes a centre in phases, each closed by a decision gate, and no phase presupposes the next. Definition: the perimeter and the scope of the centre are fixed. Formalisation: governance, intellectual-property perimeter and executive annexes are recorded. Foundation: the centre is established and licensed, the resident unit is recruited and trained, and the first programmes commence, drawing on the global bench while local capability is built. Execution: programmes in delivery. Scale: expansion governed by a separate instrument. It is the method by which Lalitpur was integrated, by which Milazzo is being built, and by which the Emirati centre is now being established.

THE STATEMENT · CONTINUED

Points 5, 6, 7 · Closing Considerations

5 · Governance and Working Model

Proportionate to the scale of the undertaking, a partnership of this class operates on a two-tier structure, expandable at the foundation phase. A Steering Committee with parity representation of the Group and the partner: strategic direction, scope, milestone validation and external positioning, on a quarterly cadence with extraordinary sessions on request of any member. A Technical and Programme Committee for day-to-day execution: the Group's technical leadership and the partner's programme direction. No unilateral external positioning concerning the partnership is made by either party without Steering Committee approval; communications discipline is treated as an institutional matter, not an administrative one. External scientific and technical communications run through the Scientific and Technical Lead; institutional and contractual communications run through the programme direction.

6 · Commercial Doctrine

The doctrine of the Group is uniform across all twenty-four centres and admits no exception: RRSC does not deploy cash into partner ventures, and does not ask its partners to capitalise RRSC. The Group contributes technology, engineering execution, quality infrastructure and institutional network; the partner contributes what it holds: sector capability, local operational presence, jurisdictional credibility and market access. The structure of each cooperation is defined jointly, on the contributions each party actually holds, and the intellectual property of each party remains its own, with joint deliverables regulated case by case.

7 · Application Logic and Sustainability

7.1 Target domains. The capability addresses a specific class of problems on which current architectures are structurally constrained: sovereign AI infrastructure for national administrations and public-sector procurement bodies requiring independence from foreign commercial APIs; robotics and autonomous systems, where embedded intelligence, multilingual conversational AI and the protection of model artefacts operating on machines in the field are procurement requirements; energy-constrained AI deployment, where the per-query energy profile of sparse architectures provides decisive value; long-context reasoning applications in scientific, engineering, legal and regulatory analysis; and critical-infrastructure deployments within the security disciplines of Section III.

7.2 Sustainability of the value. The economic logic is structured so that value is generated across multiple paths, not contingent on a single outcome: component-level value, since the validated components each have independent commercial application; the durable trade-secret regime protecting the architectural know-how; the scientific standing produced by peer-reviewed publication as the programmes mature; the trained resident team, itself a strategic asset of any centre the Group establishes; and, at each phase boundary, a fully evidence-based decision basis in the hands of both parties.

Closing Considerations of the Statement

The Statement has recorded, with the discipline the Group applies to itself, where the work is mature: the engineering execution capacity certified and referenced across fourteen years, the operational and certification infrastructure, the institutional perimeter of the Group in Europe and the Gulf, and the analytical mathematical foundation of the sovereign AI programme, a result no other actor in the field possesses. It has recorded, with equal clarity, where matters stand in formation: the new Emirati operational centre, the Augusto build-out and its expansion, the second phase of the Sofia cooperation. And it has recorded what the Group, by method, does not claim in advance: the empirical validation at scale of the integrated architecture, which belongs to the prototype phase of the programme. No technical engagement on the protected substance of the architecture takes place with any counterparty of the Group, in any phase of any partnership.

THE SPECIALTY

The Production of AI Software from Scratch

The core business of the Group, and the reason its counterparties come to it: the engineering of artificial-intelligence software from zero, for any sector, within a single institutional perimeter. Very few institutions in the world do this. The Group does it in daily production, under certification, and it is expanding.

Producing AI software from scratch 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 deploying and operating the whole inside the client's sovereign perimeter. It requires, simultaneously, 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 Group holds all three, in the configuration set out below.

ARCHITECTURAL AUTHORITY

Umberto Colella

Chief of Staff and Lead AI System Architect. Author of the Group's proprietary architecture. University of Technology Sydney; selective engagements at Google, Sony Tokyo and Microsoft Redmond; cryptographic engineering at Euroclear; a record in dual-use and national-security programmes of the North Atlantic perimeter. Qualities exceptionally rare in today's world of AI engineering.

THE PRODUCTION FLOOR

50+ Certified Engineers

The R. Rockefeller S.C. AI Center of Lalitpur, Nepal, the Group's Global Innovation Hub: 5,000+ sq. ft. of dedicated operational space, more than fifty certified IT engineers and developers, and fifteen senior managers and experts in daily production.

THE NETWORK BENCH

100+ Professionals

Across the Group's AI, cybersecurity and digital engineering network: engineers, architects, database and infrastructure specialists, quality assurance and programme management, in Nepal, Italy, Bulgaria and the Gulf. A bench in continuous expansion.

ACCUMULATED DEPTH

200+ Professional Years

Cumulative senior experience across the thirteen Senior Key Resources of the engineering centre, recognised as members of the RRSC global professional network: from national e-governance platforms to enterprise MIS, secure infrastructure and applied conversational AI.

CERTIFIED QUALITY

ISO 9001:2015

UKAS-accredited certification of the engineering centre, scope Providing Software Development Services, in its third triennial cycle. Every deliverable is engineered under a quality management system institutions recognise, with VAPT validation and full lifecycle discipline.

VERIFIABLE RECORD

14+ Years, Institutional

Delivery for the World Bank, USAID, UNDP, GIZ, the Swiss Agency for Development and Cooperation, JICA, KOICA, IFAD, ILO and the Government of Nepal. Ten executed references are reproduced in original facsimile in Part II of this Presentation.

GLOBAL REACH

24 Centres · 17 Nations

A platform operating across five continents, built on partnership-funded growth without external equity dilution and without debt, and expanding: Europe, the Gulf, Asia and Africa, with new centres in structuring.

CYBER DISCIPLINE

SecurityScorecard Alliance

Strategic Alliance signed in 2025 with the global leader in cybersecurity ratings, executed directly between Mr. Valentino Loforese and SecurityScorecard's CEO and co-founder. Penetration testing, red teaming, supply-chain security and continuous cyber-risk rating across the Group's deployments.

THE SPECIALTY

AI Software from Scratch

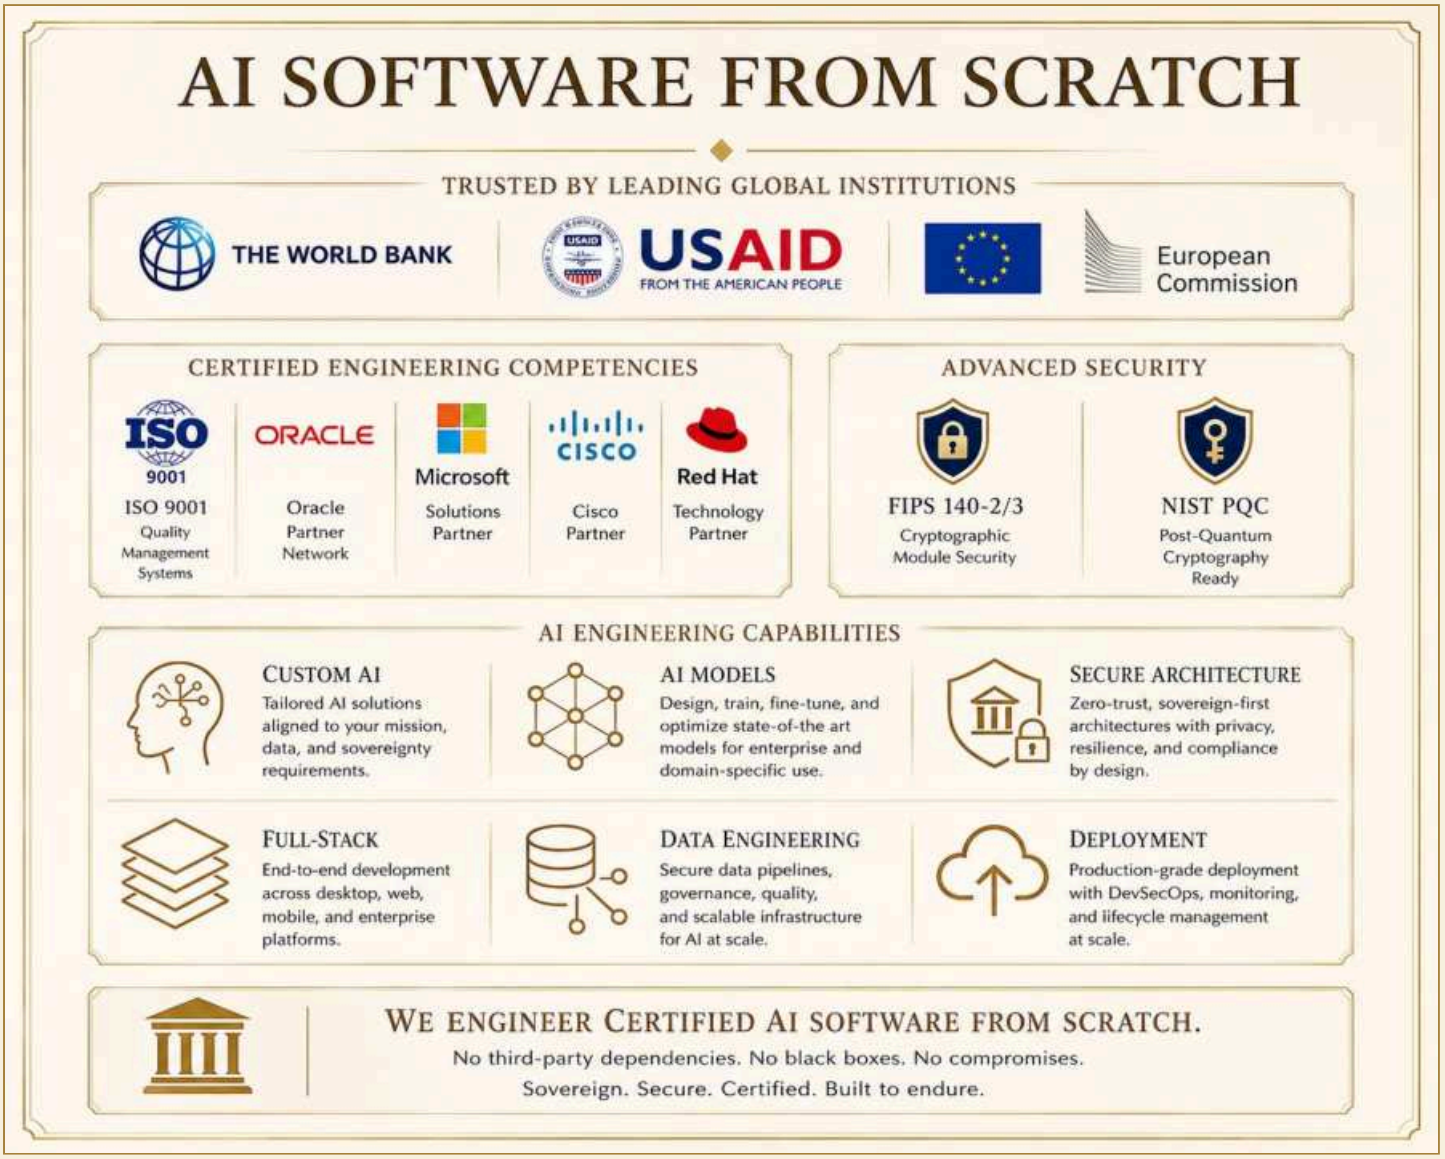

WHAT THIS PANEL RECORDS

The engineering proposition of the Department on a single panel. The upper band records the institutions for which the engineering centre has delivered: the World Bank, USAID and the European Commission. The left panel records the certified engineering competencies held by the team: ISO 9001 quality management, the Oracle Partner Network, Microsoft Solutions Partner, Cisco Partner and Red Hat Technology Partner. The right panel records the advanced security credentials: FIPS 140-2 and 140-3 cryptographic module security, and readiness for the NIST post-quantum cryptography standards. The central band sets out the six engineering capabilities the Department exercises: custom AI aligned to the client's own mission and sovereignty requirements, the design and training of models, secure zero-trust architecture, full-stack development across desktop, web, mobile and enterprise platforms, secure data engineering at scale, and production-grade deployment with DevSecOps and lifecycle management.

THE NETWORK

The Global Presence of the Group

WHAT THIS PANEL RECORDS

The operating footprint of the Group on a single panel. The map records the centres of the network across five continents: Dover in Delaware; the European nodes of Milan, Brindisi and Fermo, of Sofia, and of Istanbul and Bursa; the Gulf presence of Dubai and Riyadh; the African corridor from Tunis, Tripoli and Cairo through Accra, Abidjan, Port Harcourt, Brazzaville, Kinshasa and Nairobi to Johannesburg; and the Asian pole of Lalitpur, with New Delhi, Mumbai, Singapore and Shanghai. The right-hand panel records the institutions with which the Group works in the AI and cyber perimeter and the accreditations attaching to that work. The lower band records the figures of the network: twenty-four active centres, seventeen nations, more than two hundred and fifty engineers, and a footprint in continuous expansion. It records further the AI and cyber centres of the United Arab Emirates, Nepal and the European Union, and the engineering centres of Egypt, Turkey and Italy, each under the certifications set out beneath them.

THE PRODUCTION DOCTRINE · PART ONE

From Zero to Deployment · The Eight Stages

What the production of artificial-intelligence software from scratch actually consists of, stage by stage, in the Group's engineering practice. The sequence below is the discipline under which every institutional system of the Group is built.

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 is not academic. It determines who owns the resulting system, where the data of the institution resides, what happens when a foreign provider alters its terms, and whether the system can be operated at all within a sovereign perimeter. The Group's practice occupies the deepest position of that range, and the present Section documents it in the order in which the work is actually performed.

Stage 1 Institutional framing and sovereignty assessment

Before a line of code is written, the engagement is framed institutionally: the jurisdiction and the regulatory perimeter within which the system will operate; the classification of the data it will process and the constraints attaching to each class; the infrastructure on which it will ultimately run, whether the institution's own data centre, a national cloud or a controlled facility; the audit, retention and reporting obligations it must satisfy; and the exit position, meaning what the institution owns and can continue to operate should the relationship end. This stage produces the sovereignty specification against which every subsequent decision is tested.

Stage 2 Corpus engineering

A model is, in the first instance, its corpus. The Group engineers the training corpus rather than inheriting it: acquisition and licensing with documented provenance for every source; cleaning, normalisation and deduplication at scale; domain corpora assembled with the institution's own subject-matter experts; language coverage engineered deliberately, including the non-Latin scripts and the morphologies that generic corpora under-represent; and, where data is scarce, controlled synthetic augmentation whose generation path is itself documented. Provenance discipline is not bureaucracy: it is what allows an institution to answer, years later, the question of what its model was trained on.

Stage 3 Tokenisation and representation

The vocabulary through which a model reads its language is an architectural decision, and one that inherited tokenisers make on the institution's behalf, usually in favour of English. The Group constructs the tokenisation and representation layer for the target languages and domains: morphology-aware segmentation, treatment of institutional and technical terminology, numerals, dates and identifiers as first-class objects, and the measurement of token efficiency per language, which translates directly into inference cost and context capacity for every query the institution will ever run.

Stage 4 Architecture authoring

The model family is then designed rather than selected: the parameter budget and the fraction of it evaluated per query; the class of the attention mechanism and the context length it must sustain; the depth and width schedule; the routing and gating structure where sparsity is used; and the integration points at which the deployment security of Section III will operate. Design decisions are calibrated at sub-scale, on the scaling behaviour observed in controlled runs, before any commitment of production compute. This is the stage that almost no institution outside a small international set performs at all, and it is the stage on which everything downstream depends.

THE PRODUCTION DOCTRINE · PART TWO

Training, Validation, Deployment and Life

Stage 5 Training engineering

Training at production scale is a distributed-systems discipline before it is a scientific one. The work comprises the parallelism strategy across data, tensor and pipeline dimensions; the memory and throughput engineering that determines whether a run completes or stalls; checkpointing and failure recovery, because hardware fails and a run that cannot resume is a run that is lost; the training curriculum and the schedule of learning rates and data mixtures; and the instrumentation through which the run is observed in real time rather than diagnosed afterwards. The Group trains from one billion parameters upward, sizing the model to the institutional problem rather than to fashion.

Stage 6 Evaluation and ablation discipline

A system that is not measured is not engineered. Evaluation proceeds on three levels: standard held-out measurement for comparability; domain evaluation constructed with the institution's own experts, which is the only measurement that predicts operational behaviour; and adversarial evaluation, in which the model is attacked deliberately, in cooperation with the cybersecurity vertical of Section IV. Every architectural claim is supported by a documented ablation: the component is removed, the system is remeasured, and the contribution is recorded. Regression suites accompany the system for its entire life, so that no later improvement silently degrades an earlier guarantee.

Stage 7 Adaptation, alignment and integration

The trained base is then brought to institutional use: instruction and preference adaptation on curated data; retrieval integration against the institution's document estate, with access control preserved at the retrieval layer so that a model never surfaces to a user what that user could not lawfully open; behavioural policy engineered to the institution's own obligations rather than to a foreign provider's terms of service; and the surrounding platform, the interfaces, workflows, dashboards, integrations and administrative functions through which the intelligence actually reaches the people who use it. The Group builds that platform too: this is the difference between delivering a model and delivering a system.

Stage 8 Deployment, security and operation

Deployment is performed on the institution's own infrastructure or within its sovereign cloud perimeter: containerised, reproducible, versioned, with the cryptographic regime of Section III applied to the model artefacts at rest, in transit and at inference, and with key custody held under hardware security modules on the split-custody model. Operation follows: monitoring of quality and drift, incident procedure, retraining cadence, and the documented lifecycle under which the institution's own staff progressively assume control, with knowledge transfer, administrator and user documentation, and training of trainers as contractual deliverables.

The quality wrapper

The eight stages are executed inside the ISO 9001:2015 quality management system of the engineering centre, under full software-development-lifecycle discipline: requirement specification and system design documents, configuration management, test strategy and defect management, vulnerability assessment and penetration testing before go-live, and formal acceptance. The reference documents reproduced in Part II of this Presentation are the record of that discipline applied, repeatedly, to national systems in production.

Corpus, tokeniser, architecture, training, evaluation, adaptation, platform, deployment.
Eight stages, one institutional perimeter, no external dependency inserted along the way.

THE ENGINEERING STACK

The Stack, the Model Classes and the Platforms

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

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 deployment targets.
APPLICATION ENGINEERING	Django, React, .NET, Java and the Spring ecosystem, PHP and Laravel, mobile and desktop clients; microservice and API architectures; enterprise integration across institutional estates.
DATA ENGINEERING	Oracle, PostgreSQL, MySQL and SQL Server at national scale; petabyte-class pipelines; structured-data architecture, performance tuning, high availability, disaster recovery and Data Guard configurations; certified database administration.
INFRASTRUCTURE	Containerised deployment and DevSecOps; on-premise clusters, government cloud and sovereign cloud environments; secure network architecture, firewalls, VPNs and data-centre networking; 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; supply-chain security under the SecurityScorecard alliance.

Model classes the Group engineers

CLASS	INSTITUTIONAL APPLICATION
Conversational and assistant models	Vertical assistants for administrations and enterprises, operating on the institution's own document estate and procedures, with access control preserved and answers attributable to sources.
Multilingual language models	Language coverage engineered for the jurisdiction, including non-Latin scripts and under-served morphologies; translation, summarisation, classification and extraction across institutional corpora.
Document and process intelligence	Extraction, classification and reconciliation across administrative documents: 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, integrated into the management information systems the Group also builds.
Perception and embedded models	Vision, speech and sensor-fusion models sized for constrained hardware, for autonomous systems, robotics fleets, inspection and industrial environments, under the artefact-protection regime of Section III.
Private institutional models	The Group's own model portfolio, operated for institutional clients through the Group's centres under the custody regime of Section III, at the Group's sole discretion.

Around every model the Group 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 documented in Sections VII and VIII and evidenced in Part II. The intelligence and the platform are engineered by the same institution, under the same quality system, and delivered as one deliverable.

THE MARKET, READ PRECISELY

The Five Tiers of AI Software Production

Institutional procurement frequently treats as equivalent five activities that differ by orders of magnitude in depth, ownership and exposure. The taxonomy below is the one the Group applies when it is asked what distinguishes it from the market.

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 of service 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 itself 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, which is a material improvement; it does not hold the architecture, the corpus, the tokeniser or the training pipeline, and it inherits every property, and every latent liability, of a base model it did not build.
IV	**Continued pre-training and adaptation at scale**	Substantial capability: large-scale training infrastructure, real data engineering, meaningful model change. The architecture, however, is still someone else's; the institution operates a derivative, and its independence lasts as long as the upstream licence.
V	**From-scratch sovereign engineering**	Architecture authored, corpus engineered, tokeniser constructed, model trained from zero, cryptography proprietary, platform built, deployment inside the institution's own perimeter, documentation and knowledge transfer delivered. The institution holds a system it owns, can audit, can operate and can continue without its builder. This is the tier at which the Group works.

Why the distinction is operational, not theoretical

Four consequences follow from the tier at which a system is built, and each of them is a question a procurement authority is entitled to ask 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 be operated at all in an environment without external connectivity, which is the ordinary condition of defence, critical-infrastructure and classified deployments. And, at the end of the engagement, what precisely does the institution own: a running system with its weights, its documentation and its trained operators, or a dependency it must renew.

Tiers I to IV have legitimate and often excellent applications, and the Group does not disparage them: for many commercial problems they are the correct answer, and the Group itself integrates commodity components where sovereignty is not at stake. But when the requirement is sovereign, the tier is not a preference. It is the requirement.

THE COMPARATIVE LANDSCAPE · PART ONE

Who, in the World, Builds from Zero

The set of institutions able to author an architecture and train a model from zero is small. The set able to do that and also deliver, secure and operate the resulting system inside a client's sovereign perimeter is smaller still. The categories below describe the landscape as the Group reads it.

A The frontier commercial laboratories

A handful of laboratories, predominantly in the United States, train frontier models from zero at a scale no other actor approaches, and their scientific contribution defines the field. Their delivery model, however, is the commercial interface and the hosted platform: the institution consumes intelligence as a service, under the provider's terms and within the provider's jurisdiction. Sovereign transfer, meaning the delivery of a system the client owns and operates without the provider, is generally not the commercial proposition, and the export, licensing and jurisdictional questions attaching to such systems are outside the client's control.

B The Chinese national programme

As documented in Section XIV, a coordinated national ecosystem trains from zero at scale and releases significant open-weight families. The engineering is formidable and the architectural direction is the one this Presentation also pursues. For European, Gulf and allied procurement, however, the jurisdictional question is decisive and, in a growing number of perimeters, dispositive; and the validation method remains empirical iteration at immense compute scale rather than prior analytical proof.

C The European independents

A small number of European companies genuinely train models from zero and offer deployment options within European perimeters. They are the closest category to the Group's own position and the Group regards them with professional respect. The differences are typically three: the integration of cryptographic deployment security into the architecture itself rather than around it; the breadth of vertical engineering, meaning the ability to deliver not only the model but the institutional platform that carries it; and the certified, referenced record of delivery into government infrastructure over more than a decade.

D The global systems integrators

The large integrators bring delivery capacity, government relationships and programme management at a scale no boutique matches. Their AI substance, however, is characteristically assembled on top of third-party foundation models: excellent Tier II to Tier III work, delivered with discipline, but leaving the institution with the dependency described in the preceding Section.

E The hyperscale cloud providers

The hyperscalers supply the infrastructure on which much of the field runs, and offer sovereign regions and contractual assurances of increasing sophistication. Sovereignty so constructed is contractual and geographic rather than architectural: it constrains where data sits, not who authored the system or who can continue to operate it.

F The national research institutes and academies

Scientific depth of the highest order, and the natural partners of an industrial programme, as the Group's own cooperation in Sofia illustrates. They are not, however, delivery organisations: they produce knowledge, not certified systems in production under service obligations.

THE COMPARATIVE LANDSCAPE · PART TWO

The Seven Capabilities, and Where They Coincide

The comparison that matters to an institution is not of size but of capability coincidence: how many of the seven capabilities below a single counterparty holds within one perimeter.

CAPABILITY	FRONTIER LABS	CHINESE PROGRAMME	EUROPEAN INDEP.	SYSTEMS INTEGRATORS	HYPER-SCALERS	RRSC
Architecture authored in-house	**Full**	**Full**	**Full**	Not typical	**Partial**	**Full**
Training from zero on engineered corpora	**Full**	**Full**	**Full**	Not typical	**Partial**	**Full**
Proprietary cryptography integrated at architecture level	Not typical	Not typical	**Partial**	Not typical	**Partial**	**Full**
Sovereign transfer: client owns and operates	Not typical	**Partial**	**Partial**	**Partial**	Not typical	**Full**
Certified delivery record in public institutions	Not typical	**Partial**	**Partial**	**Full**	**Partial**	**Full**
Vertical engineering: the platform, not only the model	Not typical	**Partial**	**Partial**	**Full**	**Partial**	**Full**
Independence from third-party commercial APIs	Not applicable	**Full**	**Full**	Not typical	Not typical	**Full**

The classification reflects the Group's reading of publicly stated positioning as at August 2026, at the level of category rather than of any individual institution, and concerns commercial delivery models rather than internal technical capability. It is offered as an analytical instrument for procurement, not as an assessment of any named counterparty.

The reading

Each category above is excellent within its own logic, and an institution choosing among them should choose on the requirement, not on the reputation. What the table isolates is coincidence: the frontier laboratories hold the first two capabilities and, by construction of their business model, not the fourth; the integrators hold the fifth and sixth and, characteristically, not the first two; the hyperscalers hold infrastructure and contractual sovereignty rather than authorship. The Group's proposition is not that it exceeds any of these categories at that category's own strength, which it does not claim and would not be credible. It is that the seven capabilities coincide within a single institutional perimeter of modest size, certified, referenced and operating in daily production.

The size point deserves candour, since a serious counterparty will raise it. The Group is not a hyperscaler and does not present itself as one: it is a focused institution of twenty-four centres, whose engineering depth is concentrated rather than distributed. For an institution that requires a system it will own, operate and audit within its own perimeter, concentration is an advantage: one architectural authority, one quality system, one accountable counterparty, and a client that finishes the engagement holding the system rather than a subscription.

APPLICATION AND ENGAGEMENT

The Verticals Served and the Forms of Engagement

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

The verticals

VERTICAL	SYSTEMS THE GROUP 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: the class of system documented in Sections VII and VIII and evidenced in Part II.
Defence and critical infrastructure	Systems engineered to operate without external connectivity, under the artefact-protection regime of Section III, within the dual-use discipline of the Group's technology doctrine.
Energy and industry	Operational intelligence, forecasting and optimisation across industrial estates and energy infrastructure, integrated with the institution's own control and reporting environments.
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, fleet-level model custody, and the integration of AI capability into robotics platforms engineered by the Group or by its partners.
Finance, telecommunications and enterprise	Document and process intelligence, risk and reconciliation systems, customer-facing multilingual assistants operating inside the institution's own perimeter.

The forms of engagement

FORM	WHAT THE INSTITUTION RECEIVES
Sovereign build	A complete system engineered from zero across the eight stages, delivered on the institution's own infrastructure, with source, documentation, trained operators and the exit position defined at Stage 1.
Joint centre	A permanent engineering centre established with a local partner or host institution, under joint governance, drawing on the Group's global bench during the foundation phase and transferring capability to resident personnel thereafter.
Programme delivery	Engineering capacity from the certified bench, integrated into the institution's own programme, under the Group's quality system and its delivery discipline.
Private model operation	The qualified operation of the Group's private models for institutional clients through the Group's centres, under the custody regime of Section III and at the Group's sole discretion.
Security and assurance	Cryptographic architecture, deployment security, penetration testing, red teaming and continuous cyber-risk rating under the SecurityScorecard alliance, whether or not the underlying system was built by the Group.

THE ENGINEERING IN DEPTH · PART ONE

Sparsity, Attention and the Economics of Inference

The technical positions the Group has taken, and why each of them is an institutional decision before it is a scientific one. A system's architecture determines what it costs to run, where it can run, and for how many years the institution can afford to keep it running.

T.1 Active parameters, not total parameters.

The figure the market quotes is the total parameter count. The figure that determines cost is the active fraction: the proportion of the model actually evaluated to answer a single query. In a dense architecture the two are identical, and every query pays for the entire model. In a sparsely activated architecture, organised as a set of specialised sub-networks under a learned routing function, a query is served by a small subset of the whole. The consequences are arithmetic rather than rhetorical: a system whose active fraction is a few per cent of its total capacity serves a query at a small fraction of the compute, the memory bandwidth and the energy of a dense system of equivalent capacity. Over the operational life of an institutional deployment, measured in millions of queries, that ratio is the difference between a system an administration can afford to operate on its own infrastructure and one it cannot.

Sparse routing is not free engineering. It introduces the problems the Group's architecture work exists to solve: keeping the load balanced across sub-networks so that capacity is not wasted or overflowed; ensuring the routing function remains stable through training rather than collapsing onto a handful of favoured paths; and engineering the parallelism so that the sub-networks distributed across devices do not turn every query into a communication storm. These are the problems on which an architecture is either authored or inherited, and inheriting them means inheriting someone else's answers to all three.

T.2 The complexity class of attention, and what long context really costs.

Conventional attention compares every position in a sequence with every other: the compute grows with the square of the length. Doubling the context quadruples the cost of processing it. For institutional work, where the natural unit is not a paragraph but a procurement file, a case record, a legislative corpus or a year of operational logs, that quadratic term is the binding constraint. The Group's architectural direction combines attention layers of linear complexity in the sequence length with a minority of conventional layers where global comparison genuinely earns its cost, producing systems that read very long documents at a cost that grows proportionally rather than quadratically.

A second, less discussed cost governs deployment: the cache of intermediate state that an autoregressive model maintains while generating. It grows with every token held in context and with every concurrent user, and in production it is usually the memory ceiling, not the compute, that determines how many officials can use a system at once on a given cluster. Architectural decisions taken at Stage 4 of the production doctrine, on the structure of that state and on how aggressively it can be shared and compressed, determine the concurrency an institution obtains from hardware it has already bought.

T.3 What this yields at the point of service.

Translated into procurement terms, the three levers above, active fraction, attention complexity and cache structure, determine the four numbers an institution should ask any vendor to state and to demonstrate on the institution's own hardware: the latency to first token under realistic load; the sustained generation rate per concurrent user; the number of concurrent users per accelerator; and the energy consumed per thousand queries. A vendor operating at Tiers I to III of the taxonomy cannot answer these questions about its own system, because the system is not its own.

THE ENGINEERING IN DEPTH · PART TWO

Training at Scale, Compression and the Edge

T.4 The parallelism problem.

A model of institutional scale does not fit on one device, and training it is therefore an exercise in dividing a single computation across many. The division is performed simultaneously along several axes: replication of the model across devices each processing different data; partition of individual tensor operations across devices within a node; partition of the layer stack into pipeline stages across nodes; and, for sparse architectures, distribution of the specialised sub-networks themselves. Each axis has a different communication profile, and the art consists in choosing the combination that keeps the accelerators computing rather than waiting. The measure of that success is the fraction of theoretical hardware capacity actually converted into training progress, and it is the single most honest indicator of a training organisation's competence. It is also the difference between a run that costs what was budgeted and one that costs three times as much.

Around the arithmetic sits the engineering that determines whether a run finishes at all: memory management through activation recomputation and optimiser-state sharding; numerical strategy in reduced precision, with the loss-scaling and stability discipline that reduced precision demands; checkpoint cadence calibrated against the observed failure rate of the cluster, since on any real installation devices fail during long runs and a run that cannot resume cleanly is a run that must restart; and continuous instrumentation, so that a divergence is caught in the hour it begins rather than in the week it is discovered.

T.5 Sizing before spending.

The Group's discipline is to calibrate before committing. The relationship between model capacity, corpus volume and achieved quality is regular enough to be measured at small scale and extrapolated: a ladder of controlled runs establishes, for the institution's own domain and languages, the capacity worth building and the corpus volume required to train it properly. An institution that skips this step characteristically buys a model too large for the data it possesses, and then pays for that mismatch on every query for the rest of the system's life.

T.6 Compression, distillation and the machine in the field.

Not every deployment target is a data centre. Systems that must run inside a vehicle, a robot, an inspection unit or a facility without external connectivity are constrained by memory, power and latency budgets fixed by physics rather than by procurement. The Group engineers for those targets by reducing the numerical precision of the deployed model, by distilling a smaller model that inherits the behaviour of a larger one on the domain that matters, and by pruning and restructuring where the architecture permits. The engineering question is never compression in the abstract but the measured behaviour after compression on the institution's own evaluation set, together with the deterministic fallback that governs what the machine does when the model is uncertain.

The security regime of Section III applies at this edge with particular force, because the machine is physically accessible to an adversary. Its models remain under the per-deployment cryptographic envelope; their behavioural parameters are verified at each inference call at a cost measured in the order of one to two per cent of wall-clock time; and tampering results in a refusal to serve rather than in silent misbehaviour. An intelligent machine that fails closed is an engineering choice, and it is the one the Group makes.

Architecture is not a preference exercised at design time.
It is the cost, the concurrency and the deployability of the system for its entire operational life.

THE COMPLIANCE PERIMETER · PART ONE

The Regulatory Landscape the Group Builds Into

Institutional AI is now a regulated activity. The Group engineers to the frameworks below by construction, and states them here because an institution's exposure is determined by the perimeter its supplier was built for.

R.1 The European Union: the AI Act as it now stands.

Regulation 2024/1689 entered into force on 1 August 2024 and applies in phases. Prohibited practices have applied since February 2025; the obligations on providers of general-purpose AI models since August 2025. Following the simplification package adopted by the European Parliament on 16 June 2026 and by the Council on 29 June 2026, the timetable now reads as follows: the transparency obligations of Article 50 apply from 2 August 2026; the extension of those obligations to systems already on the market, and the newly introduced prohibitions, from 2 December 2026; the obligations on stand-alone high-risk systems listed in Annex III from 2 December 2027, deferred from the original date of 2 August 2026; and those on high-risk AI embedded in regulated products under Annex I from 2 August 2028. The architecture of the Act, its risk tiers, its conformity assessment regime and the supervisory role of the AI Office, is unchanged, and the supervisory powers of the AI Office over general-purpose models have been extended rather than reduced. Penalties reach the higher of thirty-five million euro or seven per cent of worldwide annual turnover.

The operational consequence for an institution is that the deferral is a deferral of dates, not of work. What high-risk classification actually requires, a complete inventory of deployed systems, documented risk management, data governance with recorded provenance, technical documentation sufficient for a conformity assessment, logging, human oversight and post-market monitoring, is precisely what must be engineered into a system at the time it is built. Retrofitting it onto a system assembled on a foreign commercial interface, whose training data, model documentation and behavioural policy belong to a third party, is in many cases not possible at all. The Group's eight-stage production doctrine produces these artefacts as by-products of the method, not as a compliance exercise appended at the end.

R.2 The adjacent European perimeter.

Around the AI Act sit the frameworks that govern the infrastructure it runs on and the data it processes: the General Data Protection Regulation, for lawful basis, minimisation and cross-border transfer, which is where consumption of extraterritorial commercial interfaces is most often exposed; the NIS2 Directive, for the security and incident obligations of essential and important entities; the Digital Operational Resilience Act, for the financial sector and its critical technology providers, including the contractual and exit requirements attaching to them; the Cyber Resilience Act, for the security properties of products with digital elements placed on the Union market; and the eIDAS framework, for trust services and institutional signing identity. Each of them presumes an operator who knows what is inside the system. That presumption is only satisfiable when someone built it.

THE COMPLIANCE PERIMETER · PART TWO

Cryptographic Standards and the Governance of AI

R.3 The post-quantum transition, in its actual state.

The National Institute of Standards and Technology published the first three post-quantum standards on 13 August 2024: the module-lattice key-encapsulation mechanism of FIPS 203, and the digital signature standards of FIPS 204 and FIPS 205. A code-based key-encapsulation mechanism, HQC, was selected in March 2025 as a mathematically independent second option and its standard is in preparation; a further signature standard derived from FALCON is in development. In the United States, national security systems are governed by the CNSA 2.0 migration schedule, and a federal executive order of June 2026 sets deadlines for civilian high-value systems at 2030 for key establishment and 2031 for signatures. The direction of travel is therefore settled and dated.

The Group's position is engineered accordingly, and it is stated precisely. The symmetric primitives protecting model artefacts retain their security properties against quantum adversaries at the key sizes deployed; the exposure is at the key-wrapping layer, where classical asymmetric primitives are used. That layer is abstracted behind an interface supporting both classical and post-quantum key establishment, with hybrid operation on the migration path, so that migration is a configuration change rather than an architectural one. For an institution whose data must remain confidential for decades, this matters today rather than in 2030, because an adversary who records encrypted traffic now can decrypt it later: the harvest happens before the capability arrives.

R.4 The governance standards.

ISO/IEC 42001:2023	The first international management system standard for artificial intelligence: accountability, transparency, lifecycle risk management, model and data inventory, human oversight. Increasingly requested in enterprise and public procurement questionnaires alongside information-security certification.
THE 42000 FAMILY	Impact assessment, requirements for certification bodies, and the AI-specific security and privacy standards that extend the information-security corpus to models and their data.
ISO 9001:2015	Held by the Group's engineering centre since 2017, in its third triennial cycle, scope Providing Software Development Services: the discipline under which every deliverable of this Presentation is produced.
FIPS 140-3 · COMMON CRITERIA	Hardware security module custody at Level 3 and an implementation class of EAL 4+ for the cryptographic components of the deployment security architecture.
NIST AI RMF	The risk-management vocabulary through which North American institutional counterparties commonly frame AI assurance, mapped by the Group onto the same lifecycle artefacts.

The Group's own trajectory in this dimension is stated as it is: the quality management system is certified and in its third cycle; the AI management system discipline of ISO/IEC 42001 is being implemented on the same lifecycle artefacts, and the Group reports it as in formation rather than as held. That distinction is exactly the kind the Group applies to itself throughout this Presentation.

R.5 Why the perimeter is a capability, not a burden.

Compliance frameworks are frequently presented as friction. Read correctly they are a specification, and they describe, with unusual precision, the system the Group already builds: documented data provenance, technical documentation adequate for assessment, logging and traceability, human oversight designed into the workflow, security engineered into the artefact, and an operator who can answer for the behaviour of the system because the operator built it. An institution buying from Tier I to Tier III inherits a compliance position determined by a foreign provider's decisions. An institution buying a system engineered from zero holds its own.

THE MARKET AS IT STANDS, AUGUST 2026

Sovereign AI Has Become a Procurement Category

What was an aspiration in 2024 is, in 2026, a budget line and a procurement requirement across the G20 and the Gulf. The Group documents the shift here because it defines the demand its capability answers.

M.1 The shift.

Sovereign artificial intelligence, meaning nationally controlled capability across weights, compute, data pipelines and the people who produce them, has moved from declaration to line item. The rationales that governments state are consistent: dependence on a small number of foreign laboratories and an equally narrow set of accelerator suppliers is uncomfortable at the level of national risk; language, cultural context and regulatory obligation argue for nationally controlled training data and inference; and the deployment architectures now specified in tenders are concrete about data residency rather than aspirational about it. National programmes and nationally sponsored model families now exist across Europe, the Gulf, South and South-East Asia and East Asia, and public procurement in several jurisdictions is increasingly routed through them or through sovereign-hosted variants.

M.2 The Gulf, specifically.

The United Arab Emirates has moved further and faster than most. In May 2026 the UAE Cyber Security Council, together with the national telecommunications operator and a technology partner, launched a national Sovereign AI Platform for national-scale infrastructure, announced at ISNR in Abu Dhabi and open for onboarding: an environment for generative models, agents and analytics operating inside UAE-controlled infrastructure, governed by a sovereign AI security framework that validates and monitors models, agents and workflows before they are admitted into sensitive environments. In parallel, Abu Dhabi has set the objective of becoming the first fully AI-native government across its digital services by 2027, with full adoption of sovereign cloud for government operations and a large programme of AI solutions across public services. The Kingdom of Saudi Arabia and the State of Qatar are pursuing programmes of comparable ambition within their own institutional architectures.

M.3 What this demand actually requires of a supplier.

A validated, monitored, sovereign environment does not merely relocate a foreign model into a national data centre. Admission to environments of this class requires a supplier able to state what is inside the system: the provenance of the training data, the structure of the architecture, the behaviour of the model under adversarial conditions, the cryptographic protection of the artefacts, the logging and oversight design, and the exit position of the institution. Every one of those requirements is answerable by an institution that built the system from zero, and answerable only partially, or by reference to a third party's documentation, by an institution that did not.

This is the precise point at which the Group's capability and the market's requirement meet. The Group does not claim the compute scale of a national programme or of a hyperscaler, and does not need it: what these environments require of a technology partner is authorship, transferability, security and accountability at institutional grade. Those are the four things the Group was built to provide.

SECTION I

The Doctrine of Sovereign AI

Sovereignty by construction, not by certification: what "sovereign AI" means in the Group's engineering practice, and the end-to-end capability it produces.

I.1 What sovereign AI means.

NO THIRD-PARTY CODE	Every line of code is written in-house. Zero dependencies on external AI APIs of any commercial provider.
NO EXTERNAL PRE-TRAINED MODELS	AI models are trained from scratch on proprietary data and controlled infrastructure; no derivative training from external pre-trained models outside licensed control.
PROPRIETARY CRYPTOGRAPHY	An in-house, quantum-resistant cryptographic stack, integrated at every layer, never licensed or shared.
SOVEREIGN DEPLOYMENT	On-premise or sovereign-cloud deployment: everything runs within the client's secure perimeter, with zero-trust architecture verifiable to institutional and defence standards.

I.2 The end-to-end capability.

The substantive proposition rests on a capability that very few institutions in the world possess: the ability to engineer from scratch, end to end and within a single institutional perimeter, the complete stack of an artificial-intelligence system for any strategic vertical, under sovereign control of data, architecture, training, deployment and operation. The capability comprises original architecture design at the frontier; training at production scale, from one billion parameters upward, with the distributed-systems engineering and the ablation discipline appropriate to research-grade results; deployment security and cryptographic infrastructure (Section III); vertical engineering across the strategic domains of institutional interest, from defence and critical infrastructure to healthcare, energy, finance, public administration, telecommunications, aerospace, advanced manufacturing, robotics and autonomous systems; sovereign data engineering under the applicable data-protection disciplines; and integration, validation, certification and deployment under ISO 9001:2015 quality discipline. The Group does not adapt third-party architectures. It designs them.

I.3 Why the doctrine matters to a robotics partner.

A robot is a deployed model. Its intelligence lives as trained artefacts operating on machines in the field, on hardware the operator does not always control, in environments the operator cannot always secure. For a robotics company, the sovereign doctrine translates into three direct advantages: conversational and perception models that can be engineered from scratch for the partner's platforms, including multilingual natural language processing with full Arabic-language capability; a cryptographic protection regime for the model artefacts themselves, at rest, in transit and at inference on the machine (Section III); and independence from foreign commercial APIs, which is increasingly a condition of governmental and critical-infrastructure procurement in the Gulf as it is in Europe.

We engineer certified AI software from scratch. No third-party dependencies. No black boxes.
No compromises. Sovereign. Secure. Certified. Built to endure.

SECTION II

Technology Base · Technical Positioning

The Group's proprietary sovereign AI architecture programme: posture, design doctrine, validation position, and conditions of further disclosure.

II.1 Positioning.

The flagship technological asset of the Group's artificial-intelligence vertical is a proprietary frontier architecture programme for sovereign large-scale AI, authored and led by the Group's Chief of Staff and Lead AI System Architect, with a continuous documented research record dating from his Master's research at the University of Technology Sydney. The programme is a strategic asset of the Group in the most literal sense. It is conducted by the Group alone, for its own account: the Group is not seeking partners for the programme, and the programme is not part of the perimeter of the cooperation contemplated by this Presentation. It is documented in the present Section because it establishes, better than any other single fact, the caliber of the institution the partner would join, and because every centre of the Group, including the Emirati centre this Presentation contemplates, is, by construction, qualified to operate the Group's private models for institutional clients. Consistently with the Preliminary Clarification of the Statement and with the confidentiality discipline that governs assets of this class, the present Section describes the programme at the level of design doctrine and validation posture. The programme is designated AURYNTHEON. Its mathematical specification, its component structure and its engineering documentation are not disclosed: they are not released, under any instrument, to any counterparty.

II.2 Design doctrine.

The programme is built on an efficiency-first architectural thesis: that the next generation of large-scale artificial intelligence is defined not by parameter count alone but by the fraction of the model that must be evaluated to serve a query, by the complexity class of its attention mechanism over very long contexts, and by the integration of deployment security into the architecture itself rather than around it. The design principles of the architecture operate on multi-scale learned sparsity, on linear-complexity temporal attention for million-token contexts, on architectural upcycling of validated components, and on integrated deployment security in which the gating primitives of the architecture and the cryptographic verification of Section III are the same engineering object operating in dual function. The design direction is aligned, as documented in Section XIV, with the architectural class toward which the most advanced national AI programme in the world has converged; the alignment was reached in autonomy and in advance of that convergence.

II.3 The mathematical foundation.

The element that distinguishes the programme within its architectural class, globally, is the priority given, from inception, to the analytical mathematical solution of the architectural problem itself. The architecture rests on an original and proprietary mathematical formulation whose stability, convergence and sparsity properties have been demonstrated analytically; the mathematical proof precedes, and structurally enables, the subsequent empirical validation rather than emerging from it. The rarity of this result cannot be overstated. As Section XIV documents on public sources, the Chinese national programme, for all the scale of its compute, its data and its state coordination, operates in this architectural class by empirical iteration: its claims are validated through training outcomes, not through prior analytical proof. The formulation the Group holds is, on the institutional verifications conducted to date, precisely the element that programme lacks, and that no American or European actor possesses either. The specification is conserved under the trade-secret regime described in the Statement.

SECTION II · CONTINUED

Trajectory, Validation and Disclosure

II.4 Institutional trajectory of the programme.

The programme is in an active phase of research and evolution, and its institutional trajectory is already a matter of record. Advanced dialogues for its development are under way within the European Union perimeter, with programmes of the North Atlantic alliance, and with other European centres of competence, concerning in particular the training phase and the infrastructure that phase requires. In the European defence sector, a formal scientific cooperation has been executed in Sofia with a European-Union-funded centre of competence of the Bulgarian Academy of Sciences, covering quantum-resistant cybersecurity, sovereign AI and critical infrastructure; that cooperation is now advancing to its second phase and is presented in Section XIII.

In parallel with those dialogues, the Group maintains its own private corporate plans for the self-financing of the training phase, prepared internally and independent of any institutional instrument. Two paths are consequently under evaluation, the institutional and the private, and the Group will choose between them on a single criterion: which of the two preserves, in full, its independence and its exclusive control over the asset. All of this is conducted by the Group alone and in its own name; any future involvement of a third party would be a separate and sovereign decision of the Group, taken in its own time.

II.5 Validation position.

Theoretical validation, comprising the complexity-theoretic, sparsity-emergence and energy-per-operation analyses underpinning the architectural claims, is complete: the analytical proof of the critical functions places the programme at TRL 3 on the EU Technology Readiness Level scale, reported conservatively, with the deployment security architecture likewise at TRL 3. Component-level validation by reference, via peer-reviewed published precedent on the principal building blocks at sub-scale, is well established in the open literature. Empirical validation of the integrated system at scale is the explicit deliverable of the prototype phase of the programme; the Group does not claim it in advance.

II.6 Perimeter of disclosure.

The present Section, together with Sections I and III, constitutes the entire disclosure of the programme; no further release exists. The Architecture Technical Brief, the mathematical foundations, the complexity analysis and the validation programme mapping are not released, under any instrument, to any counterparty: they remain within the Group's engineering perimeter, under the trade-secret regime described in the Statement. A partnership draws on the programme's outputs and on the institutional trajectory recorded at II.4, not on access to the architecture.

II.7 The programme and the perimeter of this Presentation.

For clarity of perimeter, three statements. First, the cooperation contemplated by this Presentation rests on the Group's core business: the production of AI software from scratch, for any sector, documented in Sections I and VI to X, together with the cybersecurity vertical of Section IV. Second, the architecture programme is an exclusive undertaking of the Group: conducted alone, for the Group's own account, in institutional dialogues with programmes of the North Atlantic alliance and within the European Union perimeter; it may, in the future and at the Group's sole and sovereign discretion, involve a third party, but at the present time the Group works on it exclusively in its own name. It is cited as added value, not as an object of the partnership. Third, the outputs of the Group's model portfolio, private models operated under the custody regime of Section III, may be made available to institutional clients through the Group's centres, at the Group's sole discretion; the per-query efficiency doctrine of the architecture, sparsity and linear-complexity attention, is part of the value those outputs carry, including for embedded and edge intelligence. Access to the architecture itself is not available to any counterparty.

The architecture is the reserve. The outputs are the offering.
What the partnership sells is sovereignty, delivered.

TECHNOLOGY BASE

AURYNTHEON · Development Status Architecture

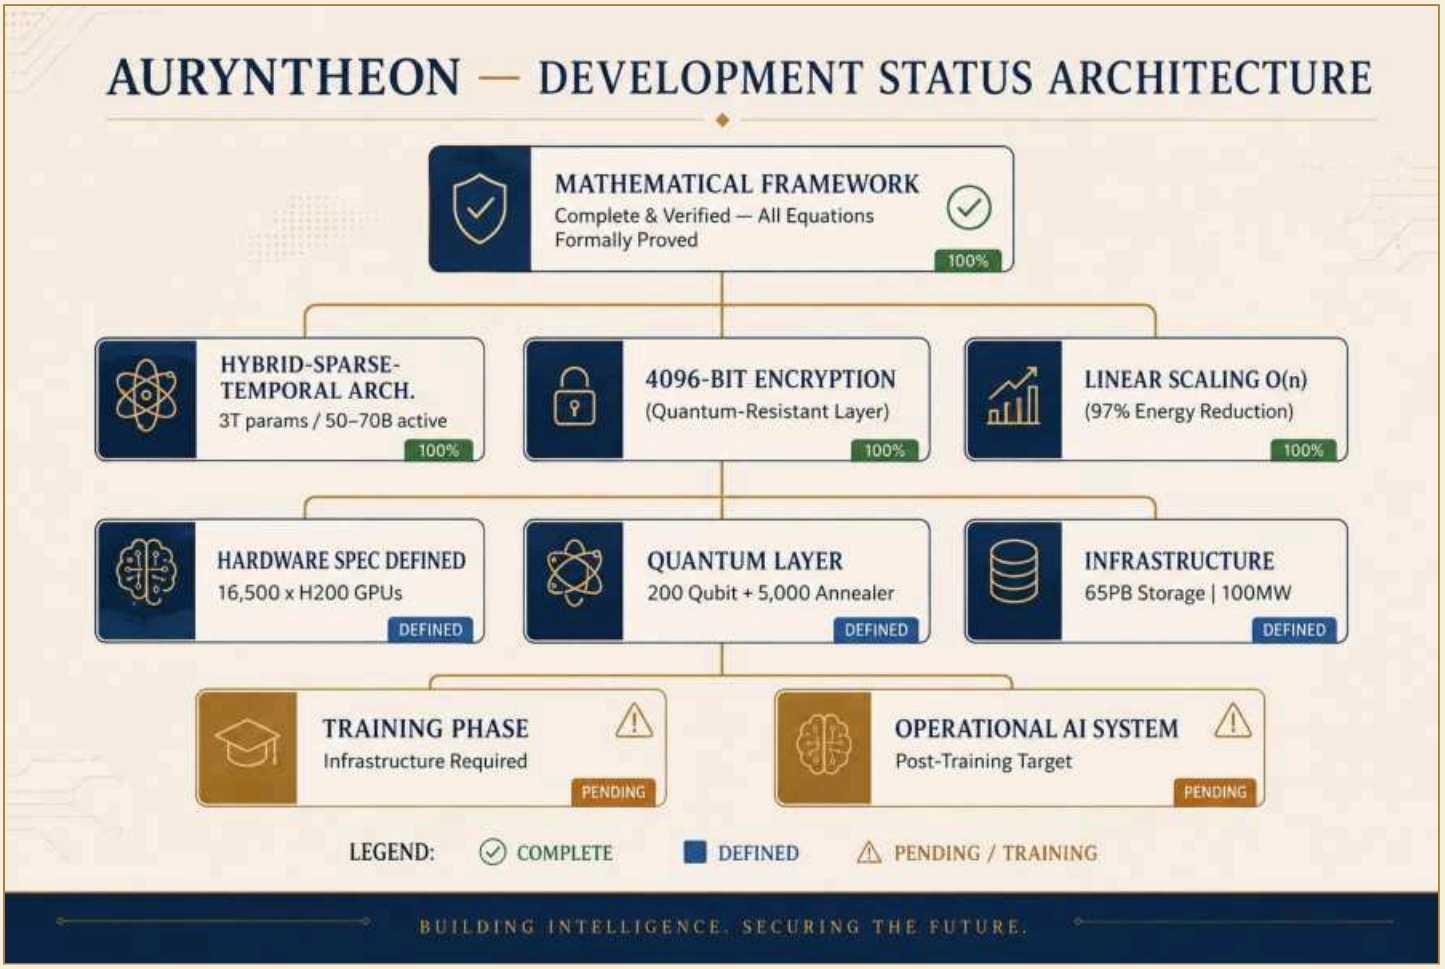

WHAT THIS PANEL RECORDS

The status of the Group's proprietary architecture programme, component by component, reported on the same conservative discipline used throughout the present Presentation. Complete and formally verified: the mathematical framework, with all equations proved analytically. Complete at the level of design: the hybrid sparse-temporal architecture, dimensioned at three trillion parameters with fifty to seventy billion active per query; the 4096-bit quantum-resistant encryption layer; and linear scaling in the sequence length, with the energy reduction that follows from it. Defined, and not yet acquired: the infrastructure the training phase requires, namely the GPU-class compute, the quantum layer, the storage estate and the power envelope. Pending, and expressly not claimed: the training phase itself, which awaits that infrastructure, and the operational system that is its post-training target. This is the position described in Section II and in Section XVIII.4: the analytical proof is in hand, the empirical validation at scale is not claimed in advance, and the two paths under evaluation for the training phase are the institutional one and the Group's own private financing.

SECTION III

Deployment Security Architecture

The cryptographic protection regime under which trained model artefacts of sovereign class are protected at rest, in transit and at inference: on servers, and on machines deployed in the field.

III.1 The structural exposure of frontier model artefacts.

Trained large-model artefacts are, in conventional engineering form, serialised tensor files on conventional storage. The artefact contains the entirety of the training investment; it can be copied bit for bit and redeployed on any infrastructure of comparable inference capacity. The protections that conventionally surround it operate at the layer of operational security, not at the layer of the artefact itself. For robotics, the exposure is sharper still: the artefact leaves the data centre and operates on machines in warehouses, malls, industrial sites and public spaces, physically accessible to adversaries. The Group's proprietary deployment security architecture addresses this exposure directly, with two contributions: it extends the cryptographic envelope from the conventional perimeter of the weights to the full set of architectural parameters that determine model behaviour at inference time; and it is co-designed with the architectural primitives of the sovereign AI programme, so that the gating function of the architecture and the cryptographic verification are the same engineering object in dual function.

III.2 Threat model.

The regime is dimensioned against four adversary classes: the sophisticated state actor, with signals-intelligence-grade resources, supply-chain compromise capability and future quantum-cryptanalytic capacity; the competing institutional actor, with major-laboratory resources and patient long-term operations; the organised criminal enterprise; and the insider with privileged access. The attack vectors considered include exfiltration of the serialised artefact at rest, interception in transit, memory-level extraction from inference servers and edge devices, side-channel model inversion and stealing, supply-chain compromise, and the eventual quantum-cryptanalytic attack on RSA-class primitives.

III.3 Layered cryptographic architecture.

LAYER	PROTECTION	PRIMITIVE CLASS
Layer 1 · At rest	Encryption of the serialised artefact on storage; the weights and each class of architectural parameters serialised into distinct blocks, each encrypted under a unique per-block, per-deployment data key	AES-256-GCM envelope; RSA-4096 key wrapping under HSM custody
Layer 2 · In transit	Encryption during inter-infrastructure transfer, with an application-layer payload key independent of the at-rest envelope, established between institutional signing identities and destroyed at completion	TLS 1.3 with FIPS-validated cipher suites; AES-256 payload
Layer 3 · At inference	Integrity verification of the routing, threshold and gating parameters at every inference call, with tampering detected at the next call and inference aborted on a cryptographic-integrity exception; performance cost in the order of one to two per cent of per-token wall-clock	HMAC-SHA-256 over parameter blocks, verified at gating-decision time

SECTION III · CONTINUED

Custody, Standards and Validation

III.4 Credential custody model.

The key-encryption keys are held in FIPS 140-3 Level 3 certified hardware security modules, under physically separated custody between the institutional principal and the partner institution within the cooperation framework. Neither party holds the full key alone: the key is reconstructed only within the secure execution environment of the HSM under a Shamir secret-sharing scheme, so that no single party, including RRSC itself, is in possession of the full decryption capability for the deployed artefact. The custody model maps naturally onto multi-party institutional governance of exactly the kind contemplated for a joint Emirati centre: the partner holds a share of the custody, structurally, by design.

III.5 Standards alignment and forward compatibility.

FIPS 140-3 Level 3 for the HSM custody; Common Criteria EAL 4+ for the cryptographic implementation; NIST cryptographic standards (FIPS 197, FIPS 180-4, FIPS 186-5, NIST SP 800-38D); ETSI TR 103 619 and TR 103 823 for the post-quantum migration framework; eIDAS-aligned signature scheme for the institutional signing identity. The key-wrapping layer is abstracted behind an interface supporting both classical and post-quantum primitives under the NIST PQC programme, with hybrid operation on the migration path; the symmetric primitives retain their quantum-resistance properties at the key sizes deployed, so that the post-quantum migration is a configuration change, not an architectural change. The regime is of direct relevance to deployments within critical-infrastructure regulatory perimeters, in the Gulf as in Europe, where the protection of model artefacts is increasingly a procurement requirement.

III.6 Implementation and validation.

The cryptographic primitives are mature, standardised and well implemented; the implementation cost is principally a cost of integration into the architectural execution path, and the architectural integration design is the proprietary contribution, specified at the engineering level. Component-level reference implementation at sub-scale forms part of the validation programme. On the EU TRL scale the deployment security architecture stands at TRL 3, reported conservatively.

III.7 Application to robotics fleets.

Applied to a robotics fleet, the regime yields a concrete operational discipline: each machine carries its models under a per-deployment cryptographic envelope; fleet updates travel under the Layer 2 transit regime; tampering with the behavioural parameters of a machine is detected at the next inference call and the machine fails safe; and the custody of the fleet's intelligence is shared between the operator and the joint centre under the HSM model of III.4. For governmental and enterprise clients in the Emirates, this converts the most sensitive objection to embodied AI, the physical exposure of the machine, into a documented, standards-aligned security architecture.

The machine can be touched. The intelligence inside it cannot be taken.

SECTION IV

Cybersecurity and Digital Sovereignty

The second pillar of the Department, of equal weight to the first: cryptographic infrastructure under the technical authority of the Lead AI System Architect, a complete portfolio of cyber disciplines across every sector, and the Strategic Alliance with the global reference in cybersecurity ratings.

IV.1 The vertical.

The Cybersecurity and Digital Sovereignty pillar carries the Group's cryptographic infrastructure capability: proprietary encryption architectures including the 4096-bit algorithmic work of the Scientific and Technical Lead, zero-trust models, secure infrastructure architecture, and the deployment security regime of Section III. The vertical is engineered from the outset to satisfy both civil enterprise and dual-use requirements, architecturally rather than by retrofit, and its disciplines are applied uniformly across every platform the Group delivers: e-governance and public financial management systems, enterprise AI platforms, and, within the cooperation this Presentation contemplates, robotics and autonomous-systems software.

IV.2 The SecurityScorecard Strategic Alliance.

In 2025 the Group executed a Strategic Alliance with SecurityScorecard, the global reference in cybersecurity ratings and threat-informed third-party risk management, signed directly between the principals of the two institutions. The alliance anchors the Group's cyber-risk practice: continuous rating, supply-chain detection and response, threat intelligence, penetration testing and red teaming, across the Group's own perimeter, the systems it deploys and the vendor ecosystems of its institutional clients. Because the alliance is a substantive component of what the Department delivers, the counterparty and the four dimensions of the cooperation are documented in full at IV.4 to IV.8 below.

IV.3 Certification and standards constellation.

QUALITY MANAGEMENT	ISO 9001:2015, UKAS-accredited, across the engineering execution centre (Section VII)
CRYPTOGRAPHIC CUSTODY	FIPS 140-3 (HSM custody) · Common Criteria EAL 4+ (implementation class)
POST-QUANTUM READINESS	NIST PQC programme alignment; hybrid classical and post-quantum key wrapping; EuroQCI compatibility perimeter
ENGINEERING CREDENTIALS	Oracle Certified Professional · Microsoft Certified IT Professional · Cisco CCNP/CCNA · Red Hat RHCE across the senior engineering bench
DEPLOYMENT DISCIPLINE	VAPT-validated government deployments; zero-trust, sovereign-first architectures; DevSecOps and full lifecycle management

The constellation above is not decorative. It is the reason the Group's platforms are accepted into government infrastructure, and it is the discipline under which the joint centre's deliverables, from enterprise AI to robotic fleets, would be engineered, validated and operated.

SECTION IV · CONTINUED

The Global Security Alliance · SecurityScorecard

The cybersecurity vertical of the Department is anchored on a Strategic Alliance with SecurityScorecard. Because the alliance is a substantive component of what the Department delivers, the counterparty is described here in full.

IV.4 Who SecurityScorecard is.

SecurityScorecard was founded in 2013 in New York by Dr. Aleksandr Yampolskiy and Mr. Sam Kassoumeh, respectively a cryptographer by doctoral training and former Chief Information Security Officer, and a former head of security and compliance, who had worked together in the security function of the same enterprise before establishing the company. It is today the global reference in cybersecurity ratings and threat-informed third-party risk management: an institution whose business is to measure, continuously and from the outside, the security posture of organisations, and to make that measurement legible to boards, regulators, insurers and counterparties.

FOUNDED	2013 · Headquartered in New York City · Operating across 64 countries
FOUNDERS	Dr. Aleksandr Yampolskiy, Chief Executive Officer · Mr. Sam Kassoumeh, Chief Operating Officer
SCALE OF MEASUREMENT	More than twelve million companies continuously rated: the largest continuously maintained external security-posture dataset in the industry
INSTITUTIONAL ADOPTION	Patented rating technology used by more than twenty-five thousand organisations for self-monitoring, third-party risk management, board reporting, due diligence, cyber insurance underwriting and regulatory oversight
THREAT RESEARCH	The STRIKE threat intelligence team, analysing in the order of eighty-three billion security events annually, with attribution work spanning more than eight hundred and seventy thousand breaches, including state-linked campaigns
PLATFORM	Security ratings; the TITAN AI platform for threat-informed third-party risk; MAX, the managed supply-chain detection and response service; attack-surface intelligence, extended in 2026 by the acquisition of a global internet-scanning and threat-intelligence capability
BACKING AND STANDING	Financed by Sequoia Capital, GV, Intel Capital, Silver Lake, Evolution Equity Partners, Riverwood Capital, T. Rowe Price, Moody's and AXA Venture Partners, at unicorn valuation; listed by the United States Cybersecurity and Infrastructure Security Agency among its published free cyber tools and services

IV.5 What a security rating actually is, and why institutions use one.

A security rating is an externally derived, continuously updated measurement of an organisation's exposure, computed from what is observable from outside its perimeter: the configuration of its internet-facing systems, the hygiene of its network and application estate, the credentials of its people circulating in criminal markets, the reputation of its addresses, the patching discipline visible in its services. It is, in effect, the credit rating of cyber risk, and institutions use it for the same reasons they use credit ratings: to compare, to monitor continuously rather than annually, to hold suppliers to a measurable standard, and to report to a board or a regulator in a language that does not require the reader to be an engineer.

SECTION IV · CONTINUED

What the Alliance Gives the Department, and Its Clients

IV.6 The alliance.

The Strategic Alliance was executed in 2025 and signed directly between the principals: Mr. Valentino Loforese, Chief Executive Officer and Majority Shareholder of the Group, and Dr. Aleksandr Yampolskiy, Chief Executive Officer and co-founder of SecurityScorecard. It is a technical and commercial alliance, and it operates in four dimensions.

DIMENSION	WHAT IT MEANS IN PRACTICE
The Department measured first	The Group's own perimeter, and the perimeters of the systems it deploys, are continuously rated and monitored under the same discipline the Group asks of others. A supplier of security that is not itself measured is asking for trust it has not earned; the Group's position is the opposite, and it is verifiable from outside.
Supply-chain and third-party risk for clients	Institutional clients rarely fail through their own systems; they fail through a vendor. The alliance gives the Department continuous visibility over a client's vendor ecosystem, with detection and response applied to the supply chain rather than periodic questionnaires filed and forgotten.
Threat intelligence into the architecture	The adversary classes of the deployment security architecture, Section III, are not abstractions: the threat research of the alliance, including its work on state-linked campaigns and on the exposure of internet-facing AI platforms and agent frameworks, informs the threat model against which the Group's cryptographic regime is dimensioned and revised.
Assurance across the lifecycle	Penetration testing, red teaming and vulnerability assessment, aligned to the VAPT discipline that government infrastructure deployments require, applied to systems the Group builds and to systems it did not build, as an independent engagement.

IV.7 Why the combination is unusual.

Rating institutions measure; engineering institutions build. An organisation that receives a poor rating on a critical system is told, with precision, what is wrong, and then has to find someone able to correct it inside a system whose architecture it may not control. The Department sits on both sides of that line: it can measure the exposure with an instrument the market recognises, and it can remediate it at the level of the architecture, because in the systems it delivers it authored the architecture. Measurement without the capacity to remediate is a report. Remediation without independent measurement is an assertion. The two together are an assurance position, and it is the position the Department offers.

IV.8 A statement on independence.

The alliance is an alliance, not a dependency. It does not extend to the Group's proprietary architecture programme, which remains outside every cooperation perimeter; it confers no rights over the Group's cryptographic stack, which is engineered in-house and never licensed or shared; and it does not place any third-party component inside the systems the Group builds. What it adds is external, independent, continuously updated measurement, and the threat intelligence that keeps a security architecture honest against adversaries who do not stand still.

We build it, we secure it, and we are measured on it by an institution that does not work for us.
That is what the alliance is for.

SECTION IV · CONTINUED

The Cyber Portfolio · The Complete Perimeter

Cyber is not, in this Department, an accessory of the artificial-intelligence practice. It is the second pillar, of equal weight, and it is exercised across the full perimeter of the discipline: offensive, defensive, architectural, cryptographic and regulatory.

DOMAIN	WHAT THE DEPARTMENT PERFORMS
Offensive security	Penetration testing of networks, applications, mobile estates, wireless and physical perimeters; red teaming and full adversary simulation against detection and response capability; social engineering and phishing assessment; purple-team exercises conducted with the client's own defenders.
Vulnerability and exposure management	Continuous vulnerability assessment, external attack-surface discovery, prioritisation by exploitability rather than by score alone, remediation engineering, and re-testing to closure. VAPT validation to the standard required by government infrastructure deployments.
Security architecture	Zero-trust architecture design, network segmentation and micro-segmentation, secure landing zones, identity and access management, privileged access, public-key infrastructure, hardware security module design and key custody.
Cryptographic engineering	The Group's proprietary in-house stack, including the 4096-bit algorithmic work of the Scientific and Technical Lead; data protection at rest, in transit and in use; key lifecycle; post-quantum migration planning and hybrid key establishment under the standards of the Compliance Perimeter.
Detection and response	Security operations centre design, build and operation; detection engineering and use-case development; threat hunting; incident response, containment and recovery; digital forensics and post-incident reporting to board and regulator.
Application and product security	Secure software development lifecycle, threat modelling at design time, static and dynamic analysis, dependency and software-bill-of-materials discipline, secrets management, DevSecOps pipelines, and secure code review by engineers who write production code themselves.
Cloud and infrastructure security	Sovereign and hybrid cloud hardening, configuration baselines, container and orchestration security, workload isolation, backup and recovery architecture, and the security engineering of on-premise clusters for institutions that cannot use public cloud at all.
Operational technology and critical infrastructure	Industrial control and SCADA environments: protocol-aware segmentation, safety and security interface, passive monitoring where active scanning is prohibited, and the discipline required where an availability failure is a physical event and not a service interruption.
Third-party and supply-chain risk	Continuous rating and monitoring of vendor ecosystems, supply-chain detection and response, and the threat intelligence of the alliance documented at IV.4 to IV.8, applied to the client's own suppliers and to the client's own external posture.
Governance, risk and compliance	Programme design and audit readiness against the frameworks of the Compliance Perimeter: information-security and AI management systems, sectoral regulation for finance, health and critical infrastructure, incident-reporting obligations, and the documentation an institution must be able to produce on demand.
Human layer and capability building	Awareness programmes, role-based technical training, executive and board briefings, tabletop and crisis exercises, and the formation of the client's own resident security team, which is the only durable outcome of any security engagement.
Security of artificial intelligence	The frontier discipline documented at IV.13: protection of model artefacts, security of agentic systems, adversarial testing of models, and the integrity of the data and training pipeline.

SECTION IV · CONTINUED

Cyber Across the Sectors, and the Operating Model

The discipline is one; its application is sectoral. What a ministry needs is not what a refinery needs, and neither resembles what a robotics fleet needs.

IV.11 Sectoral application.

SECTOR	THE CHARACTERISTIC CYBER PROGRAMME
Government and e-governance	Protection of national platforms of the class the Department itself builds: treasury, budget, revenue, labour, registries and citizen channels. Segregation of duties across ministries, audit trails admissible before an auditor general, and deployment inside government infrastructure under national standards.
Defence and national security	Environments that must operate air-gapped or in degraded connectivity; classification-aware architecture; cryptographic sovereignty; supply-chain provenance for every component admitted into the perimeter.
Energy, utilities and industry	The operational-technology perimeter, where a compromise is a physical event: segmentation between corporate and control networks, passive monitoring, vendor remote-access governance, and incident procedures written for engineers rather than for analysts.
Finance and insurance	Operational resilience obligations, third-party concentration risk, exit and continuity testing, fraud-adjacent detection, and evidence packages sized for a supervisor rather than for an internal committee.
Health and social protection	Protection of the most sensitive personal data an administration holds, in estates that are chronically under-resourced: identity, minimisation, encryption by default, and recovery capability that assumes compromise rather than denying it.
Telecommunications and transport	Large distributed estates with legacy interoperability: asset discovery at scale, segmentation of the exposed perimeter, and detection engineering across heterogeneous technologies.
Robotics and autonomous systems	Machines in the field are an attack surface with physical consequences: fleet identity and credential management, signed firmware and model updates, integrity verification of behavioural parameters at inference, and the fail-closed discipline described in the Engineering in Depth.
Smart cities and public space	Sensor estates, video and access-control systems, and the governance of data collected in public: privacy engineering as a security control, and clear separation between operational necessity and surveillance.

IV.12 The operating model.

FORM	WHAT IT DELIVERS
Assess	Point-in-time engagements: penetration test, architecture review, compliance gap assessment, third-party exposure baseline, with a remediation plan the institution can execute or delegate.
Build	Engineering of the security architecture itself: zero-trust design, cryptographic infrastructure, identity, secure pipelines, and the hardening of systems in production.
Run	Managed operation: monitoring, detection engineering, continuous rating of the institution and its suppliers, patch and exposure governance, under agreed service levels.
Respond	Incident response engagement, with forensics, containment, recovery and the reporting an institution owes to its board and its regulator.
Transfer	Formation of the client's resident team and of its own security function, which is how the Department measures success: the institution able to continue without it.

SECTION IV · CONTINUED

Securing Artificial Intelligence Itself

The intersection at which this Department is unusual: the institution that engineers artificial intelligence from zero is the same institution that attacks it, protects it and answers for it. The discipline below is the newest in the field, and the one the market is least equipped to perform.

IV.13 The attack surface of an AI system.

An artificial-intelligence system introduces exposures that conventional security programmes were not designed to see, and that grow sharply once models are given tools, autonomy and access to institutional systems. The Department addresses them as engineering problems, in six dimensions.

EXPOSURE	THE ENGINEERING ANSWER
The model artefact	Weights and behavioural parameters are the entire training investment in a file that can be copied. The layered cryptographic regime of Section III protects them at rest, in transit and at inference, with split key custody and integrity verification at each call.
The training pipeline	Corpus poisoning, provenance failure and contaminated third-party data are supply-chain attacks against the model's future behaviour. The provenance discipline of the production doctrine is the control: every source documented, every transformation reproducible.
The inference boundary	Instruction injection through untrusted content, in documents, web pages, emails and tool outputs, is the characteristic attack on assistant systems. The answer is architectural: privilege separation between content and instruction, retrieval that preserves the user's own access rights, and output handling that treats model text as untrusted input to the systems downstream.
Agentic autonomy	An agent with credentials and tools is an unattended user account with a reasoning engine attached. Controls are conventional in nature and rigorous in application: 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, refusal and guardrail regression suites run at every release, with results recorded as evidence rather than asserted as assurance.
The deployment estate	Inference servers, vector stores, orchestration frameworks and edge devices are ordinary infrastructure with extraordinary value, and they are secured with the ordinary discipline of the portfolio: hardening, segmentation, monitoring, patching and recovery.

IV.14 Why the intersection is rare.

Securing an AI system requires knowing how it was built. A security firm that has never authored an architecture, engineered a corpus or run a training programme can test the surface of such a system, and will find real exposures there; it cannot reason about the failure modes that live inside the model, because it has never built one. Equally, a laboratory that builds models but does not operate a security practice will document its safeguards and be unable to prove them under adversarial conditions. The Department does both, in one perimeter, with the same engineers: it builds artificial intelligence from zero, and it attacks and defends what it builds, together with what others have built.

Two pillars, one institution: we engineer the intelligence, and we engineer its security.
Neither is a service line attached to the other.

SECTION V

Scientific and Technical Lead · Mr. Umberto Colella

The scientific and technical authority of the Group's sovereign AI programmes, and the originating architect of its proprietary architecture.

The scientific and technical authority of the Group's sovereign AI programmes rests with Mr. Umberto Colella, Chief of Staff and Lead AI System Architect of The Rob Rockefeller Standard Carbon LLC, and the originating architect of the Group's proprietary architecture. He is a world-class AI System Architect specialised in full-stack algorithm design, advanced mathematics, cybersecurity, encryption and neural architectures, with a combination of scientific depth, engineering rigour and institutional authority that places him within the smallest fraction of engineers worldwide operating at this level. His qualities and areas of specialisation are exceptionally rare in today's world of AI engineering, placing him among the finest professionals in his field.

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 described within the programme as a singular innovation in cryptographic systems. Through the University's competitive international exchange programme, admission to which is determined on academic standing and technical examination, he completed selective engagements at Google (search algorithm optimisation and penetration testing, 2015 to 2016), Sony Corporation Tokyo (real-time graphics software and advanced algorithms, 2015) and Microsoft Corporation Redmond (encryption systems and adaptive threat detection, 2016 to 2017). His subsequent engagement with Euroclear involved cryptographic engineering of institutional cross-border settlement protocols, SHA-256 systems, digital asset custody and SWIFT-perimeter work, together with the hardening of the underlying server infrastructure.

A substantial part of his professional record concerns dual-use and national-security programmes conducted for countries of the North Atlantic alliance perimeter, across the United States, the United Kingdom and Italy. The substance of that work is governed by confidentiality and classification obligations, in some cases perpetual in scope, and cannot therefore be described in detail in a document of the present class; the description provided here is at the level of competence area: defence and military AI, cybersecurity and cryptography including quantum-resistant protocols, neural architecture design, custom hardware and neuromorphic computing, sovereign cloud and distributed systems. Supporting documentation, including academic certificates and references from prior engagements, is available under the confidentiality framework to qualified reviewers nominated by the counterparty institution.

His full-time institutional engagement is with RRSC, where he leads the sovereign AI architecture portfolio across the global Innovation Hub network of the Group. Within the cooperation this Presentation contemplates, he serves as the designated technical authority of the partnership: the counterpart of the partner's engineering leadership, and the guarantor of the architectural rigour of every joint deliverable.

CORE SPECIALISATION	Artificial intelligence and conversational AI systems · cybersecurity and secure infrastructure architecture · e-governance and public financial management systems
ELITE TRACK RECORD	Google · Microsoft · Sony · Euroclear
SIGNATURE WORK	Proprietary 4096-bit encryption and zero-trust models, optimised for maximum security, performance and scalability

SCIENTIFIC AND TECHNICAL LEAD

Mr. Umberto Colella

WHAT THIS PANEL RECORDS

The profile of the Chief of Staff and Lead AI System Architect, author of the Group's architecture. The left panel records his core areas of technical specialisation: artificial intelligence and conversational systems built on advanced neural networks and natural language processing; cybersecurity and secure infrastructure architecture, including the proprietary 4096-bit encryption and the zero-trust models; and e-governance and public financial management platforms at federal and provincial level. Below it, the additional expertise: optimised search algorithms and penetration testing, the development of the 4096-bit encryption algorithm, and a cross-sector record spanning defence, technology and research. The right panel records the elite track record behind that competence: search algorithm optimisation and security at Google, cloud architecture at Microsoft, advanced algorithm design at Sony, and financial cryptography at Euroclear.

THE PEOPLE

AI & Cyber Department Management

WHAT THIS PANEL RECORDS

The management of the Department. The upper row records the scientific authority and the direction: Mr. Umberto Colella, Chief of Staff for AI, Cybersecurity and Digital Centers; Mr. Mohan G.C., General Director of the AI Center of Nepal; Mr. Santosh Bhusal, President of the Center; Mr. Udaya Neupane, Network and System Infrastructure Expert; Mr. Upahar Kumar Joshi, Database Administrator; Mr. Dharmaraj Budhathoki, Director; Dr. Adel Miran, General Director of Public and Institutional Relations for the United Arab Emirates; Dr. Michele Vincenti, Director of Global Academic Partnerships and Artificial Intelligence; and Mr. Simone Ruggeri, Fashion and Marketing AI Director. The lower rows record the engineering bench in daily production: senior Java, backend and frontend developers, technical leads and project managers, quality assurance and database administration.

Nepal AI Center. *The senior leadership team.*

The senior leadership of the wholly-owned RRSC AI Center in Lalitpur — the Group's flagship AI / Cybersecurity / Digital capability, 5,000+ sq ft, 50+ certified IT engineers, 15 senior managers, ISO 9001:2015 certification, Anthropic Lab integration.

Mohan G.C.
General Director
AI Center of Nepal

Mohan G.C.
General Director

General Director of the AI Center. 18+ years of ICT leadership with the Government of Nepal (Ministry of Finance, FCGO, MoICS, MoALD), the World Bank, IFAD, USAID, JICA and JSDF.

Santosh Bhusal
President of the Center

Santosh Bhusal
President of the Center

President of the RRSC AI Center Nepal. Anchors the institutional governance and multilateral-donor relationships of the Center — the senior counterpart for the Anthropic Lab integration and the wider multilateral ecosystem (JICA, JSDF, KOICA).

Dharmaraj Budhathoki
Director

Dharmaraj Budhathoki
Director

Director of the RRSC AI Center Nepal in Lalitpur. Senior leader within the 15-strong senior management team coordinating the daily operational delivery of the Centre and the AI / Cybersecurity / Digital capability.

Udaya Neupane
Network & System Infrastructure Expert

Udaya Neupane
Network & System Infrastructure

Network & System Infrastructure Expert of the RRSC AI Center Nepal. Leads the network architecture, system infrastructure and cybersecurity-grade design ensuring the Centre meets sovereign-grade institutional standards.

Sarju Bista
Technical Lead / Project Manager

Sarju Bista
Technical Lead / Project Manager

Technical Lead and Project Manager at the RRSC AI Center Nepal. **18+ years experience, .NET Expert.** Coordinates the multi-project technical delivery across the AI, Cybersecurity and Digital programmes.

Sunil Adhikari
Technical Lead / Project Manager

Sunil Adhikari
Technical Lead / Project Manager

Technical Lead and Project Manager at the RRSC AI Center Nepal. **15+ years experience, PHP/Laravel Expert.** Focus on institutional and governmental client deliverables.

Upahar Kumar Joshi
Database Administrator • Oracle/MS SQL

Upahar Kumar Joshi
Database Administrator • Oracle/MS SQL

Senior Database Administrator at the RRSC AI Center Nepal. **25+ years experience, Oracle/MS SQL Expert.** Leads the database engineering, data governance and structured-data architecture foundational to the Group's sovereign-AI platform.

Sushil Kumar Sah
Database Administrator • Oracle Expert

Sushil Kumar Sah
Database Administrator • Oracle Expert

Database Administrator at the RRSC AI Center Nepal. **18+ years experience** as **Oracle Expert.** Co-leads database engineering, data governance and structured-data architecture, with focus on enterprise-grade Oracle deployments for institutional clients.

Manoj Kumar Mahato

SENIOR JAVA DEVELOPER

Senior Java Developer at the RRSC AI Center Nepal. Leads enterprise-grade Java engineering across the Group's sovereign-AI platform and structured backend services for multilateral-donor and governmental clients.

Monish Manandhar

SENIOR DEVELOPER

Senior Developer at the RRSC AI Center Nepal. Engineers solutions across the AI / Cybersecurity / Digital capabilities, contributing to architecture, deployment and operational delivery of the institutional platform.

Yam Bahadur Limbu

SENIOR DEVELOPER

Senior Developer at the RRSC AI Center Nepal. Contributes to architecture, engineering and operational delivery of the Group's sovereign-AI platform across multilateral and governmental client engagements.

Bijay Joshi

FRONTEND DEVELOPER & UI-UX SPECIALIST

Frontend Developer and UI-UX Specialist at the RRSC AI Center Nepal. Leads user-facing application engineering and design-system implementation across the Group's sovereign-AI products, ensuring institutional-grade usability.

THE PEOPLE

The AI Technical Team Structure

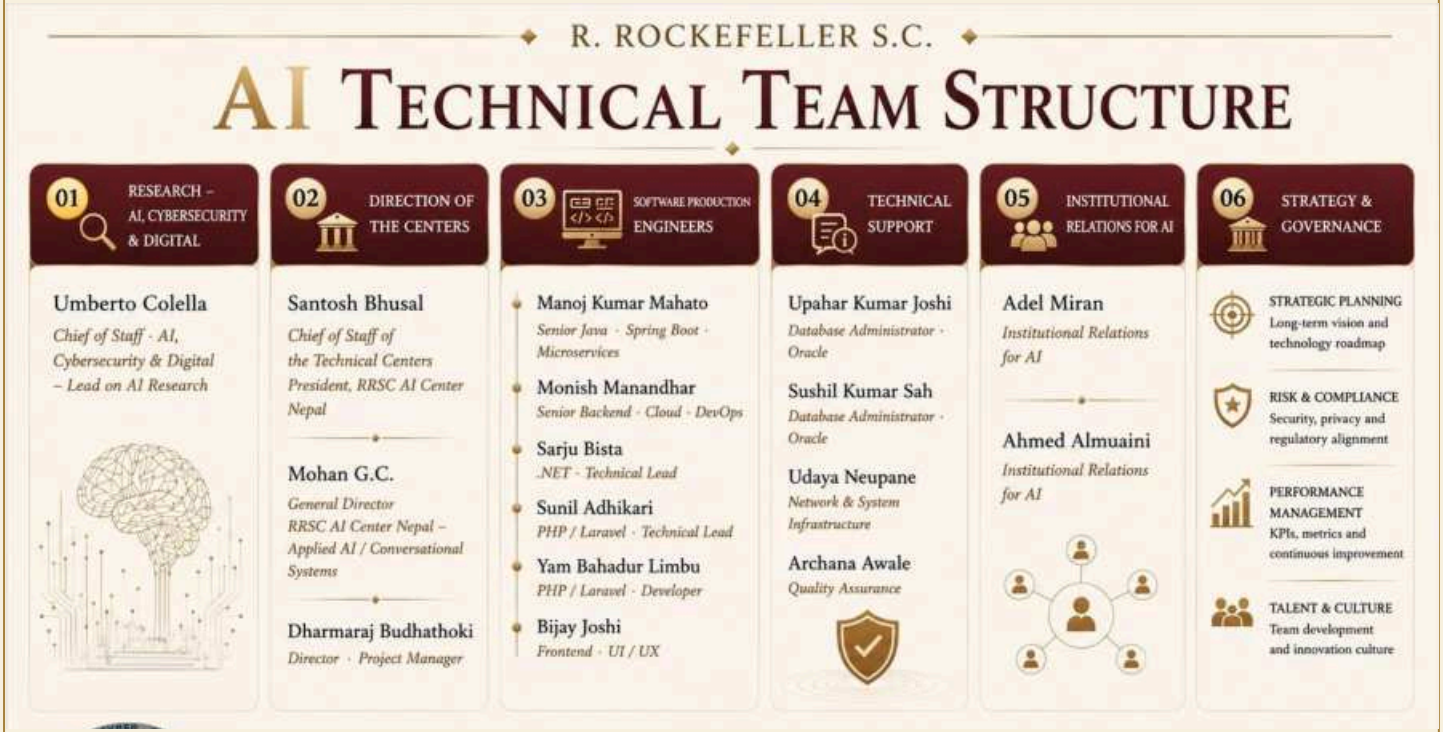

WHAT THIS PANEL RECORDS

The Department organised by function, in six columns. Research in artificial intelligence, cybersecurity and digital, under Mr. Umberto Colella. Direction of the centres, under Mr. Santosh Bhusal, Mr. Mohan G.C. and Mr. Dharmaraj Budhathoki. Software production engineers, with the Java, backend, .NET, PHP and frontend specialisations named individually. Technical support, covering database administration on Oracle, network and system infrastructure and quality assurance. Institutional relations for AI, under Dr. Adel Miran and Mr. Ahmed Almuaini. Strategy and governance, covering strategic planning, risk and compliance, performance management and talent. The lower band records the three AI engineering and innovation centres of the network: the United Arab Emirates, Nepal and the European Union.

SECTION VI · IN NEPAL, WE ARE THIS

Nepal · The Global Innovation Hub

The R. Rockefeller S.C. AI Center of Lalitpur: the Group's principal engineering production floor, certified, referenced, and in daily production for institutional clients across seventeen countries.

The R. Rockefeller S.C. AI Center · Nepal · Global Innovation Hub 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., having accepted full integration into the Group by the corporate instrument executed on 5 May 2026. The instrument confers the official designation, places the centre's engineering capacity at the disposal of the Group's programmes, recognises its thirteen senior professionals as members of the RRSC global professional network, preserves the Group's pre-existing intellectual property as the exclusive and perpetual property of RRSC, and binds the parties to confidentiality obligations that are indefinite with respect to trade secrets.

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
GROUP ACTIVATION DATE	5 May 2026

VI.1 The engineering stack and mandate.

The centre produces the Group's sovereign AI software and platforms, engineered from scratch: conversational models and vertical assistants with multilingual natural language processing; e-governance and public financial management systems for treasury, budget, tax and grievance-redress perimeters; mission-critical database engineering, disaster recovery and structured-data architecture; and full-stack development across desktop, web, mobile and enterprise platforms, with production-grade DevSecOps deployment and lifecycle management at scale. The stack is engineered under the sovereign doctrine of Section I: no third-party AI code, no external pre-trained models outside licensed control, proprietary cryptography, and on-premise or sovereign-cloud deployment within the client's secure perimeter.

Within the cooperation this Presentation contemplates, Lalitpur is the engineering bench on which the joint Emirati centre draws in its foundational phase: senior engineers, delivery discipline and quality infrastructure available from day one, while the resident Emirati team is recruited and trained. Capability is developed first where it exists, and then transferred under institutional ownership to where it is strategically required.

NEPAL

The Engineering Center at a Glance

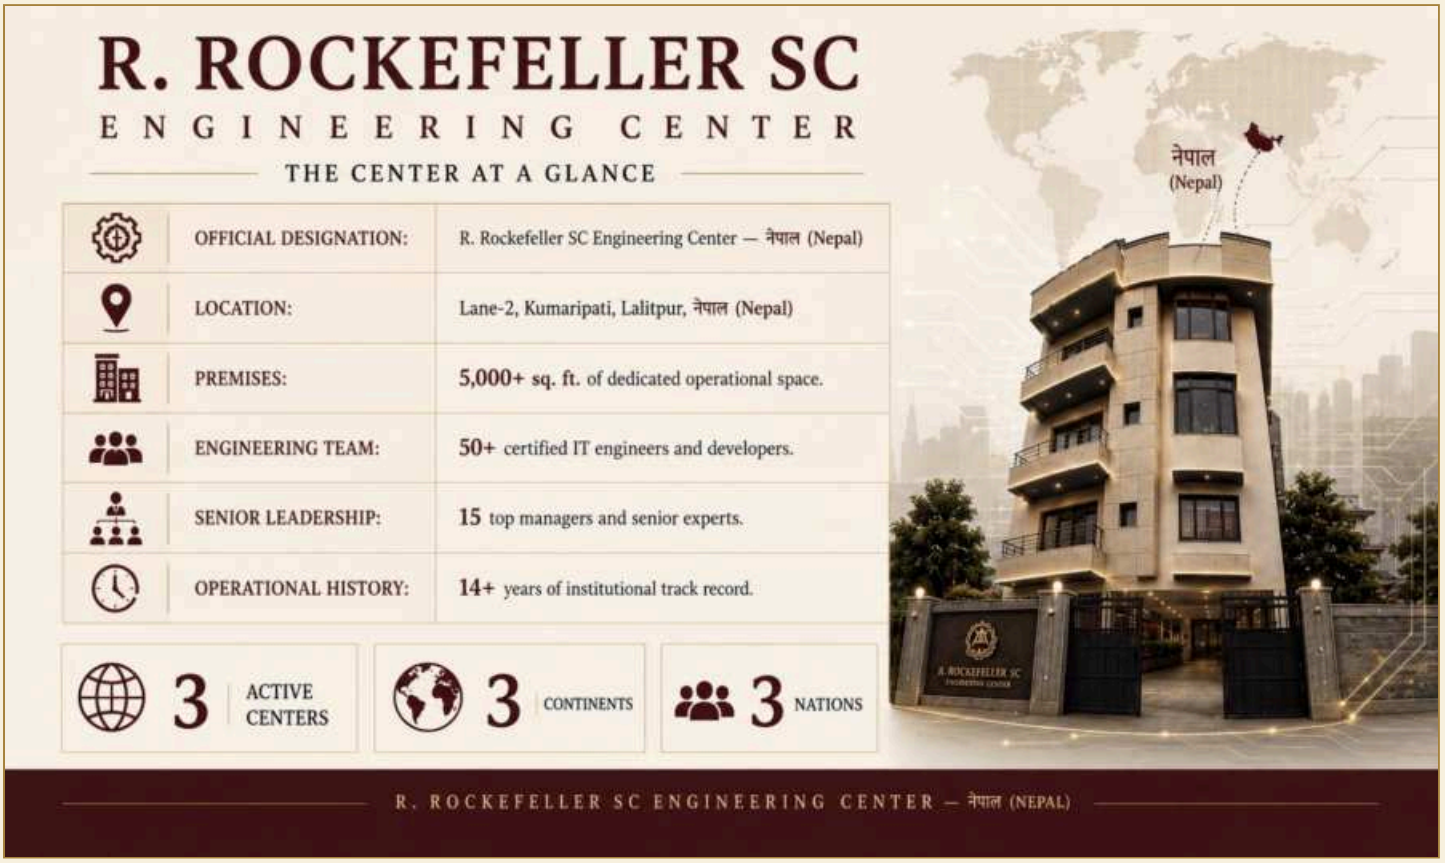

WHAT THIS PANEL RECORDS

The R. Rockefeller S.C. Engineering Center of Nepal, summarised on the figures that matter to an institutional client. Official designation and location at Lane-2, Kumaripati, Lalitpur. Premises of more than five thousand square feet of dedicated operational space. An engineering team of more than fifty certified IT engineers and developers, with fifteen top managers and senior experts in the senior leadership. More than fourteen years of institutional track record. The lower band records the wider network position: three active centres, three continents, three nations.

SECTION VII · IN NEPAL, WE ARE THIS

Nepal · Verification Dossier

The institutional verification of the engineering execution capacity. The certification and reference documents below were issued to the centre's operating entity across its fourteen-year delivery record; they are not modified, retitled or recharacterised, and the originals are reproduced in facsimile in Part II.

VII.1 ISO 9001:2015 Certificate of Registration.

CERTIFICATE NUMBER	77763/A/0001/UK/En
STANDARD	ISO 9001:2015 · Quality Management Systems
ISSUING BODY	United Registrar of Systems Ltd. (URS)
ACCREDITATION	UKAS Management Systems no. 0043; IAF multilateral recognition
ORIGINAL ISSUE DATE	23 August 2017
CURRENT CYCLE	Issued 8 September 2023 · valid through 22 August 2026 · third triennial cycle
SCOPE OF CERTIFICATION	Providing Software Development Services
CERTIFIED ENTITY	Dryice Solutions Pvt. Ltd. · Jwagal, Kupondole, Lalitpur, Nepal

VII.2 The integration instrument of 5 May 2026.

The corporate instrument executed on 5 May 2026 between The Rob Rockefeller Standard Carbon LLC and Dryice Solutions Pvt. Ltd. is the act by which the Lalitpur centre entered the Group's perimeter: it confers the designation R. Rockefeller S.C. AI Center · Nepal · Global Innovation Hub, places the centre's engineering capacity at the disposal of the Group's programmes, and recognises its thirteen Senior Key Resources as members of the RRSC global professional network. Because its recitals identify protected designations of the Group's proprietary technology portfolio, the full text is exhibited to counterparties within the framework of a partnership rather than reproduced in the present edition.

VII.3 Master list of completed institutional engagements.

The broader delivery portfolio spans more than fourteen years of continuous institutional engagement: the Public Financial Management Strengthening Project for the Ministry of Finance (LMBIS and TSA, World Bank and USAID frame, including the Louis Berger engagement); public-sector platforms across multiple ministries (BankPro, eBPS, human-resources and payroll, inventory, learning, document, school, budget, revenue and hospital management systems); the PayPenny and Connecting Creations digital-payments engagement in Malaysia (2021); and the referenced engagements tabulated in Section VIII. Institutional clients span the Government of Nepal, the World Bank, USAID, UNDP, GIZ, JICA, KOICA, SDC, IFAD and ILO, with the centre operating throughout under the ISO 9001:2015 Quality Management System and in compliance with Government of Nepal infrastructure standards, with VAPT-validated deployments and full SDLC discipline.

NEPAL

Credentials and ISO 9001:2015 Certification

WHAT THIS PANEL RECORDS

The quality credential under which every deliverable of the Department is produced. The centre holds ISO 9001:2015 certification for quality management systems across the software development lifecycle, with international standard compliance covering quality assurance, configuration management and customer satisfaction, and more than fourteen years of institutional track record behind the certificate. The original certificate, issued by United Registrar of Systems under UKAS accreditation, is reproduced in facsimile in Part II of the present Presentation.

SECTION VIII · IN NEPAL, WE ARE THIS

Nepal · Reference Letters Portfolio

The principal completed institutional engagements, each documented by an executed reference letter or contractual completion document reproduced in original facsimile in Part II.

ENGAGEMENT	CLIENT / FUNDER	REFERENCE POSITION
ILMIS · Integrated Labour Management Information System, thirteen modules, GIDC and G-Cloud deployment with VAPT validation	Department of Labour and Occupational Safety, Government of Nepal	Letter ref. 688, signed by the Computer Engineer of DoLOS; 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 and Statement of Work of 16 November 2022, prime ref. 81063667
Digital health profile of two municipalities, linked to social health protection	GIZ Deutsche Gesellschaft für Internationale Zusammenarbeit	Contract 83408972, effective 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, GHSC-PSM, 9 March 2022; contract GHSC-PSM-NPL-09, 2019 to 2021
Office Automation System for the Embassy of Switzerland in Nepal	Federal Department of Foreign Affairs, Swiss Confederation	Contract executed 3 February 2020 by H.E. the Ambassador
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, HMIS-ICT Unit, Department of Roads
E-Commission Blueprint and ICT Roadmap; Case Management and Reporting System	National Human Rights Commission of Nepal · UNDP	Executed references, 2023 to 2024, Strategic Plan Support Project
PayPenny online money booking system · MiS Project	Connecting Creations Sdn. Bhd., Malaysia	Completion certification of 2 March 2021; delivery closed 10 April 2021
National IT Roadmap of Nepal	Department of Information Technology, Government of Nepal	Executed reference, 2014

Each reference above is reproduced photographically, in its original executed form, in Part II of the present Presentation, in chronological order from most recent to most historical.

SECTION IX · IN NEPAL, WE ARE THIS

Nepal · Leadership of the Centre

The senior leadership of the engineering execution centre: eighteen-plus years of World Bank and USAID programme experience at the head of a fifty-engineer production floor.

Mohan G.C.
General Director · R. Rockefeller S.C. AI Center, Nepal
A Nepal-born ICT professional with more than eighteen years of progressive experience spanning software engineering, enterprise systems design, programme monitoring and evaluation, and executive leadership; Chief Executive Officer of Dryice Solutions Pvt. Ltd. since 2025, having served as its Technical Project Manager from 2019. He began his career in 2007 as a .NET developer, then led large-scale MIS platforms as Software Technology Officer of the Rural Water Supply and Sanitation Fund Development Board, a World Bank-funded Government of Nepal project, and from 2014 to 2019 served as ICT and M&E Chief at the Poverty Alleviation Fund, a second major World Bank initiative, responsible for information systems governance across multiple project districts. His consulting portfolio includes short-term technical consultancy for Deloitte under the USAID Public Financial Management programme, supporting the Line Ministry Budgetary Information System (LMBIS and PLMBIS); data management for JICA-funded emergency school reconstruction; impact and terminal surveys for agriculture MIS under UNDP; and ERP business-process engineering for the Hydroelectricity Investment and Development Company. MBA, Apex College Kathmandu (2014); BE in Information Technology, Nepal College of Information Technology (2007). In recent years he has directed his expertise toward AI-enabled civic technology incorporating conversational AI and multilingual natural language processing.

Santosh Bhusal
President of the Center
An accomplished information-technology and business professional holding a Bachelor's degree in Information Technology and a Master's in Business Studies, with more than seventeen years of international experience across the ASEAN region specialising in system design, enterprise solutions and business-process transformation. As President of the Center he holds the institutional representation of the Lalitpur hub within the Group's global network and leads its international client engagement.

Dharmaraj Budhathoki
Director · Co-Founder, Dryice Solutions
A highly experienced project manager and team leader whose career at Dryice Solutions spans from 2012 to the present, with expertise across the full software development lifecycle: from system analysis and requirement gathering to large-scale national deployment and support. His record includes large-scale government rollouts and ERP implementations, leading cross-functional teams in delivering complex IT solutions for government bodies and international organisations including USAID and the World Bank.

NEPAL

Mr. Mohan G.C. · General Director

WHAT THIS PANEL RECORDS

The institutional profile of the General Director of the AI Center of Nepal. The card records his professional summary, his areas of expertise across software development, technical project management, MIS design and ICT governance, applied and conversational AI and public financial management systems, and his professional journey from 2007 to the present, including the World Bank funded programmes of RWSSFDB and the Poverty Alleviation Fund. The consulting record lists engagements with Deloitte under the USAID public financial management programme, with JICA on emergency school reconstruction, with UNDP on agriculture management information systems, and on ERP business-process management.

NEPAL

Mr. Santosh Bhusal · President of the Center

WHAT THIS PANEL RECORDS

The institutional profile of the President of the Center and Chief of the AI Technical Labs. The card records his professional biography, his specialisations in AI software production and protection and in cybersecurity, his focus areas across AI innovation, secure and resilient systems, enterprise solutions and technology leadership, and the elite track record of international consulting for the World Bank, USAID, the International Labour Organization and the Japan Social Development Fund.

SECTION X · IN NEPAL, WE ARE THIS

Nepal · Senior Key Resources

The thirteen Senior Key Resources recognised under the integration instrument constitute the operational core of the centre, with cumulative experience exceeding two hundred professional years.

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

The Senior Key Resources are members of the RRSC global professional network and serve within the Group's Lalitpur centre; individual curricula vitae in extended institutional format are available within the confidentiality framework of a partnership.

THE RECORD

The Institutions Served

WHAT THIS PANEL RECORDS

The institutional clients and funders behind the fourteen-year delivery record of the engineering centre: the World Bank and USAID through the Global Health Supply Chain Program; UNDP, the International Labour Organization, IFAD and the German development agency GIZ; KOICA, JICA and the Japan Social Development Fund; the Swiss Embassy and Swiss Cooperation; and the implementing partners Abt Associates, Louis Berger, VSO and WSP. The executed reference letters from these engagements are reproduced in original facsimile in Part II.


EUROPE AI / CYBER CENTERS
CONNECTING INNOVATION • RESEARCH • FUTURE
ITALY
INNOVATION ROOTED IN EXCELLENCE
MILANO
Administrative Coordination Center for Italy
SICILIA
Work in Progress. Possible AI Software Development and Research Hub
ARTIFICIAL INTELLIGENCE
R. Rockefeller S.C.
RESEARCH CENTER
BULGARIA
INNOVATING FOR TOMORROW
SOFIA
AI & Cyber Hub. Research Excellence. Talent Development. Digital Transformation
STRATEGIC PARTNERSHIP
CENTER OF COMPETENCE QUASAR
Quantum Communication, Intelligent Security Systems, Risk Management
PROJECTS & PARTICIPATE IN
EuroQCI
European Quantum Communication Infrastructure
HORIZON EUROPE
EU Framework Programme for Research and Innovation
NATO DIANA
NATO DIANA
Defence Innovation Accelerator for the North Atlantic
AI RESEARCH
Pioneering the future with intelligent systems
CYBERSECURITY
Securing tomorrow with advanced protection
DATA & CLOUD
Scalable data solutions and cloud infrastructure
GLOBAL COLLABORATION
Uniting experts, driving global impact


THE ROB ROCKEFELLER S.C.
ITALY
The European Hub
Via Montenapoleone 21, Milan
HEADQUARTER
in Milan (250 m²)
in via Montenapoleone 21
1 ENGINEERING CENTER
A EUROPEAN HUB
Connecting innovation, engineering and excellence across Europe.

SECTION XI · IN ITALY, WE ARE THIS

Italy · The European Hub

The Italian Branch: Milan and Brindisi as the European institutional and engineering anchor of the Group, and the coordination layer of the European perimeter.

The European hub of the Group is the Italian Branch: headquarters in Milan, Via Monte Napoleone 21, of approximately 250 square metres, anchoring the European institutional perimeter, with the Brindisi engineering centre, three thousand square metres in technical partnership with RePG, as the strategic coordination centre for senior design works and the seat of the Group's Technology Research Laboratory. The Branch coordinates the institutional public-relations programme across the European and African perimeters, bridging the Italian operational layer with the European Commission frameworks and the African continental network, and provides the operational alignment across the Group's twenty-four active centres, interacting between Branch CEOs, Country Presidents and the Office of the Chief Executive Officer.

XI.1 Leadership of the Italian Branch.

Mr. Francesco Cuccurullo
Italian Branch Administrator
Coordinates the European hub from Milan, Via Monte Napoleone 21, interfacing the Italian institutional perimeter, the Milan headquarters and the Brindisi strategic coordination centre.

Mr. Gianluca Annunziata
Board of Directors · Italian Branch
Anchors the institutional governance and the European institutional perimeter, coordinating with the Italian Branch Administrator and the European quantum-communication institutional frameworks.

Mr. Gennaro Migliore
Director · Institutional Public Relations, Europe and Africa
Coordinates the institutional public-relations programme across the European and African perimeters, bridging the Italian branch operational layer with the European Commission and the African continental network.

Mr. Sandro Nasta
Executive Manager, Engineering Area · Director, Technology Research Laboratory
Directs the engineering area of the Italian Branch and the Technology Research Laboratory anchored on the Brindisi centre.

Mr. Simone Ruggeri
Fashion and Marketing AI Director
Leads the applied-AI programmes of the Branch in the fashion and marketing verticals of the Italian industrial perimeter.

Mr. Carlo Elasi
Branches Coordinator
Coordinates the operational alignment across the Group's active centres and seventeen sovereign jurisdictions worldwide.

To this standing architecture the Group is now adding its third and flagship Italian node: the technology centre of the Augusto Innovation Hub at Milazzo, presented in the following Section as what the Group will build next in Europe.

SECTION XII · IN ITALY, WE WILL BUILD THIS

Italy · The Augusto Innovation Hub, Milazzo

The Group's next European flagship: a permanent centre for end-to-end sovereign artificial-intelligence and cybersecurity engineering, established by RRSC as technology partner and technical leader of one of the most significant innovation campuses in development in Southern Europe.

XII.1 The undertaking.

The Augusto Innovation Hub of Milazzo, Sicily, is, in the reading of the Group, one of the most significant innovation-infrastructure undertakings currently in development in Southern Europe: a campus of approximately 47,000 square metres and a strategic designation secured at European level within the perimeter of the European Union's Strategic Technologies for Europe Platform, at the exact centre of the Mediterranean's digital corridors. The investment programme, opened in the order of one hundred million euro, is already being scaled toward two hundred million for the creation of the strategic Italian centre, and the perimeter of the programme is expanding beyond the original campus, extending the Group's Italian axis alongside Milan, Brindisi and Naples. Within this undertaking, The Rob Rockefeller Standard Carbon LLC is the technology partner and the technical leader of the campus programme: the party that brings the anchor technological content of sovereign grade.

XII.2 Statement of purpose of the Milazzo centre.

The centre the Group is establishing at Milazzo is conceived not as the execution of a single technological programme but as a permanent institutional asset within the territory of the European Union: a centre for end-to-end sovereign artificial-intelligence engineering, operating under joint governance. Its mandate is concrete: the production of sovereign AI software for institutional and industrial clients; cybersecurity engineering under the discipline of Section III and of the SecurityScorecard alliance; and the qualified capability to operate the Group's private models, an access that exists nowhere else in the European Union and that only a centre of the Group can carry. The Centre is not a project office, not a temporary consortium, not a programme team. It is a standing centre of European institutional grade, with its own governance, its own engineering capacity, its own deliverable portfolio and its own institutional positioning.

XII.3 The institutional rationale.

The European Union, notwithstanding substantial public investment, does not at present have a single institutional centre combining the four conditions necessary for genuine end-to-end sovereign AI engineering: original architecture design at the frontier; production-scale training capability; deployment security aligned with where institutional infrastructure is going; and vertical engineering across the full perimeter of strategic interest. Excellent European institutions exist for each individual dimension: academic networks for research, the high-performance computing joint undertaking for compute, national centres of excellence for applications. The integration does not exist as a permanent institutional asset. The integrator function is performed today by a small number of actors in the United States and the People's Republic of China; the Augusto Centre constitutes the European institutional response to that asymmetry, integrating the existing institutions rather than competing with them.

SECTION XII · CONTINUED

Milazzo · Portfolio, Build-Out and Meaning for the Gulf

XII.4 Deliverable portfolio.

The portfolio is organised in three concentric circles. The inner circle: the sovereign production programmes operated by the Centre, the production of sovereign AI software and the qualified operation of the Group's private models for institutional clients. The middle circle: vertical engineering programmes commissioned by European institutional clients, Member State administrations, agencies, critical-infrastructure operators and industrial corporations, for which the Centre is positioned as a preferred engineering partner. The outer circle: the scientific, educational and ecosystem-development activities through which the Centre contributes to the broader European doctrine: peer-reviewed publication as the programmes mature, training programmes for senior European engineers, technical workshops, contribution to European standardisation.

XII.5 Build-out.

The Centre is integrated from inception with the global engineering network, and in particular with Lalitpur, which contributes the senior engineering execution capacity on which delivery rests during the foundational phase; as the Milazzo seat matures, the centre of gravity of the capability transitions, in proportion to recruitment and training, toward the campus. Within twelve months from the activation of the executive calendar the Centre establishes its governance bodies, its legal and contractual framework, its resident personnel function and its quality system aligned with the ISO 9001:2015 discipline already in operation in the Group; within twenty-four months it completes the foundation phase, inaugurates the first vertical programmes and operates as a self-standing institutional centre. The dedicated resident engineering unit at Milazzo is sized, at regime, at twenty-six personnel across machine-learning research, systems architecture, data engineering, infrastructure and site-reliability engineering, security engineering and programme management, additional to the existing capacity of the global network. The scaling of the envelope toward two hundred million euro corresponds to this widened ambition: not a single building on a single campus, but the strategic Italian centre of the Group, with the capacity to host programmes of European rank and to anchor the expansion of the Italian axis.

XII.6 What Milazzo means for an Emirati partner.

For a prospective counterparty of the Group, the Augusto programme carries a precise meaning. First, it demonstrates method: the same institutional logic the Group proposes for the Emirates, a permanent joint centre, phased build-out, parity governance, engineering transferred from Lalitpur to the resident seat, is the logic already in execution on European soil. Second, it creates symmetry: the Milazzo centre and the Emirati centre become the two wings of the same sovereign-technology architecture, the European wing and the Gulf wing, each qualified to operate the Group's private models, each integrated with the Lalitpur production floor. Third, it opens a corridor: a partner joined to the Group in any jurisdiction is, structurally, connected to a European sovereign-technology centre of STEP designation, with everything that connection implies for programmes, clients and standing.

Milazzo is what the Group builds next in Europe, and it is growing as it is built.
The Emirati centre is what the Group builds next in the Gulf. The method is the same.

ITALY

The Augusto Innovation Hub · Milazzo

WHAT THIS PANEL RECORDS

The Augusto Innovation Hub of Milazzo, Sicily, in the architectural design of the Giancarlo Zema Design Group, and the figures of the programme as they stand. The campus occupies approximately 47,000 square metres on the former Montecatini site, at the exact centre of the Mediterranean's digital corridors, under a strategic designation secured at European level within the perimeter of the European Union's Strategic Technologies for Europe Platform. The band below the render records the investment architecture: a first phase of one hundred million euro over 2026 to 2028, carried by the Sicily Regional Government within the European Union framework, and a prospective second phase of a further hundred million over 2027 to 2030 at the level of the Italian central government, taking the programme to two hundred million euro in total. At full capacity the campus is sized for five hundred direct jobs. Within the undertaking, The Rob Rockefeller Standard Carbon LLC is the technology partner and the technical leader of the campus programme, as recorded in Section XII.


SUSTAINABILITY • AI • CYBER • ENERGY • GEOPOLITICS
AURYNTHEON
R. ROCKEFELLER S.C.
SecurityScorecard
CENTER OF COMPETENCE QUASAR
Quantum Communication
Intelligent Security Systems
Risk Management
NATO OTAN
NATO DIANA
Defence | Innovation | Adoption
EU
EDF European Defence Fund
HORIZON EUROPE
EuroQCI


R. ROCKEFELLER S.C.
SecurityScorecard
CENTER OF COMPETENCE QUASAR
A STRATEGIC PARTNERSHIP FOR A SOVEREIGN EUROPE
Uniting expertise. Advancing technology.
Building the future, together.
TRUSTED PARTNERSHIP
Built on trust, transparency and shared values.
TECHNOLOGICAL EXCELLENCE
Combining scientific leadership with industrial execution.
STRATEGIC IMPACT
Developing critical capabilities for security, resilience and sovereignty.
EUROPEAN VISION
Strengthening Europe's independence in key technologies.
INNOVATION & GROWTH
Investing in AI, Quantum, Cybersecurity and advanced infrastructures.
TOGETHER, FOR THE FUTURE
Creating long-term value for society, industry and future generations.
WORKING WITHIN A STRATEGIC FUNDING ECOSYSTEM


CENTER OF COMPETENCE QUASAR
The Rob Rockefeller Standard Carbon LLC
Center of Competence QUASAR – BAS
MEMORANDUM OF UNDERSTANDING
SBIL/QCI

SECTION XIII · IN BULGARIA, WE ARE THIS

Bulgaria · The Sofia AI and Cyber Hub

The Group's Eastern European pole: research excellence, talent development and digital transformation, inside the European defence and quantum perimeter.

XIII.1 The Sofia hub.

The Sofia AI and Cyber Hub is the Group's Eastern European pole, under Mr. Kalin Petkov Edrev, Director Bulgaria. Its mandate spans research excellence, talent development and digital transformation, and it anchors the Group inside the European defence and sovereign-technology perimeter: the environment in which the institutional grade of the Group's AI programme is received and advanced in Europe.

XIII.2 The strategic partnership with the Centre of Competence QUASAR.

In the European defence sector, the Group executed in Sofia, in April 2026, a formal scientific cooperation with the Centre of Competence QUASAR, a European-Union-funded centre of competence of the Bulgarian Academy of Sciences, spanning quantum communication, intelligent security systems and risk management, and covering quantum-resistant cybersecurity, sovereign AI and critical infrastructure. That cooperation is now advancing to its second phase. It is, in the reading of the Group, the single clearest external verification of the institutional grade of the AI programme: a European-Union-funded scientific institution of a Member State academy, in the defence sector, engaging the Group's sovereign-technology perimeter on the merits.

XIII.3 The European programmes perimeter.

EUROQCI	The European Quantum Communication Infrastructure: the compatibility perimeter for the Group's quantum-resistant cryptographic stack
HORIZON EUROPE	The EU Framework Programme for Research and Innovation: the principal European funding ecosystem for the innovation programmes of the perimeter
NATO DIANA	The Defence Innovation Accelerator for the North Atlantic: the alliance-level innovation perimeter within which the Group's dialogues proceed

This positioning places the Group's programmes within the principal European funding ecosystems for innovation, sovereign technology and defence, whose combined envelope exceeds one hundred billion euro over the current programming cycle: the EU flagship research programme, the European digital and defence instruments, the European quantum-communication infrastructure and the NATO innovation accelerator. The ecosystem above is institutional context, stated as such; it is not a list of commitments.

XIII.4 What Sofia means for the joint offering.

For a Gulf partner, the Sofia pole means that the joint centre's technology base is not a bilateral claim: it is a capability already engaged, in the defence sector, by a European-Union-funded scientific institution, and already advancing to a second phase. Counterparties in the Emirates do not need to take the Group's word for the grade of its technology; they can read the institutional company it keeps.

SECTION XIII · CONTINUED

Bulgaria · What the Group Will Build in Sofia

The second phase of the QUASAR cooperation, and the Sofia AI and Cyber Centre as the Group's Eastern European node within the European defence and quantum perimeter.

XIII.5 The second phase of the cooperation.

The cooperation executed in Sofia with the Centre of Competence QUASAR is advancing to its second phase, which moves the relationship from scientific framing to joint work. The three workstreams of that phase, as scoped between the parties, are the following. First, quantum-resistant cybersecurity: the alignment of the Group's cryptographic stack, including the key-wrapping and custody regime of Section III, with the post-quantum migration frameworks and with the quantum-communication infrastructure of the European perimeter. Second, sovereign artificial intelligence: joint scientific work on the architectural class described in Section II, conducted at the level of design doctrine and validated components, with the protected substance of the Group's architecture remaining, as everywhere, undisclosed. Third, critical infrastructure: the application of both to the protection of infrastructures within the European regulatory perimeter, where the security of models and of the systems that carry them is becoming a procurement requirement rather than a feature.

XIII.6 The Sofia AI and Cyber Centre.

On this foundation the Group is structuring its Sofia node as a standing AI and Cyber Centre: a research and engineering seat, integrated with the Lalitpur production floor and with the Italian axis, whose mandate is the quantum-resistant cryptographic engineering of the Group, the scientific interface with the Bulgarian Academy of Sciences and its European programmes, and the development of talent within one of the strongest engineering education perimeters of Eastern Europe. The centre is, in the Group's architecture, the scientific counterweight to the industrial capacity of Nepal and the institutional capacity of Milan: the place where the doctrine is tested against the European scientific establishment.

XIII.7 Why this matters to a counterparty.

A counterparty engaging the Group does not engage a claim: it engages an institution whose technology base is already under examination, in the defence sector, by a European-Union-funded scientific institution of a Member State academy, and whose cooperation with that institution is advancing rather than concluding. For governmental and enterprise clients, in Europe as in the Gulf, that is the difference between a supplier and an institution.

Nepal produces. Milan represents. Sofia proves.
The three together are what the Group means by sovereign engineering.

SECTION XIV

Strategic Positioning in the Global Context

The convergence of a new architectural paradigm, the Western strategic vacuum, and the position of the Group within it.

XIV.1 Premise.

The Group's sovereign AI programme does not operate in a technological vacuum. It develops in parallel with a strategic convergence that, over the last eighteen to twenty-four months, has acquired the proportions of a coordinated national programme within the People's Republic of China: the abandonment of the dense Transformer paradigm in favour of hybrid and sparse architectures; the integration of linear-complexity attention for extended contexts; and the preparation of the computational infrastructure for quantum computing as an AI accelerator. The programme identified and pursued this direction in autonomy from its conception, in advance of the institutional consolidation now visible in the sequence of Chinese model releases, publications and infrastructure investments.

XIV.2 A national programme.

The transition is not a technical trend in isolation; it has been institutionally elevated to an economic and national-security priority. The 15th Five-Year Plan (2026 to 2030) positions artificial intelligence and quantum computing at the centre of the national development strategy, with the declared objective of seizing the commanding heights of scientific and technological development; AI investment doubled in the course of 2025, and scientific-research expenditure rose 8.3 per cent over the previous year, in parallel with a deliberate reduction of dependence on Western technologies. The principal industrial expressions of the convergence in the present cycle:

ACTOR	ARCHITECTURAL CONTRIBUTION
Alibaba	The Qwen 3.5 and 3.6 model families, combining linear attention with sparse Mixture-of-Experts
Baidu	The Wenxin 5.0 model: Mixture-of-Experts at 2,400 billion parameters with selective activation of approximately three per cent of experts per token
DeepSeek	The V3 and V4 series, integrating sparse Mixture-of-Experts and sparse attention against the leading closed-source frontier models
Ant Group	The Ling-2.6-Flash model, activating 7.4 billion of 104 billion parameters through sparse Mixture-of-Experts
Zhipu AI	The GLM-5 model, trained entirely on Huawei Ascend chips as the institutional signature of the technological-independence doctrine
Xiaomi	The MiMo-V2.5 Pro: sparse multimodal Mixture-of-Experts at 310 billion parameters, open source under MIT licence
BAAI	Continuing research on advanced Mixture-of-Experts architectures, including the GigaMoE family for computer vision

The structural significance of the enumeration is not the competition between individual actors but the institutional coherence of a national programme in which each pursues the same architectural pillars under the coordinating framework of the State Council.

SECTION XIV · CONTINUED

The Western Vacuum and the Position of the Group

XIV.3 The quantum dimension.

In parallel, the Chinese ecosystem constructs the infrastructural substrate for quantum-AI integration: China Telecom Quantum Group operates the Tianyan quantum cloud, with the recently commissioned Tianyan-287 superconducting machine (105 qubits, 99.90 per cent fidelity) available at global scale; Huawei has developed the MindQuantum framework for quantum machine-learning workflows and operates the Hongmeng Quantum Cloud at 128 qubits; the CAS Quantum Excellence Center in Beijing has declared the objective of a 1,000-qubit machine by the end of 2027. These are not isolated research initiatives; they are infrastructure investments oriented toward consolidating quantum-AI integration within the perimeter of Chinese technological sovereignty.

XIV.4 The Western scenario: a strategic vacuum.

Against this coordinated programme, the Western landscape exhibits a structurally fragmented profile. The United States holds undisputed leadership in foundational research, and its principal laboratories operate at the frontier of the field; but their architectural posture has, with limited exceptions, remained anchored to the dense Transformer paradigm, with innovation taking the form of incremental refinement rather than the paradigmatic shift toward integrated sparse-temporal-quantum architectures. The European Union, notwithstanding its declarations on digital sovereignty and its substantial public investment, has excellent institutions within each individual dimension, academic research networks, high-performance compute infrastructure, quantum-communication programmes, national champions, but has not consolidated an integrated industrial programme combining the three pillars under a single architectural direction with a proprietary mathematical foundation. In synthesis: no Western initiative is presently identifiable that integrates sparse Mixture-of-Experts architecture, linear-complexity attention and quantum-computing integration into a single coordinated industrial programme with a proprietary mathematical foundation.

XIV.5 The position of the Group.

While the Western landscape has remained anchored to the dense paradigm or fragmented across uncoordinated initiatives, the Group's sovereign AI programme has independently identified, and structurally pursued, the same architectural direction that the Chinese ecosystem has elevated to national-programme status. On the basis of the institutional verifications conducted by the programme, the Group is the single Western industrial actor presently engaged in the integrated pursuit of the three architectural pillars under a coordinated programme with proprietary mathematical foundations. Three elements are decisive. Architectural alignment: reached in autonomy, in advance of the Chinese consolidation, and not by imitation. The mathematical advantage: where the principal Chinese actors operate on architectural engineering validated through data-driven empirical iteration at immense compute scale, the Group's architecture rests on an original and proprietary mathematical formulation whose stability, convergence and sparsity properties have been demonstrated analytically, the proof preceding and structurally enabling the empirical validation rather than emerging from it. And engineering execution capability: the certified, referenced delivery capacity of Sections VI to X, combined with the deployment security architecture of Section III and the cybersecurity discipline of the SecurityScorecard Strategic Alliance. The institutional reception of this position in Europe is already a matter of record, as Section XIII documents. The specification of the architecture is not disclosed, in any forum.

A STATEMENT OF SEPARATION

The China Perimeter · Sustainability Only, and Nothing Else

The Group states its position on the People's Republic of China explicitly, because in the technology perimeter of this Presentation the question is a procurement question, and a counterparty is entitled to a precise answer rather than a diplomatic one.

C.1 What the Group has in China, and what it does not.

The Group maintains an active department and a joint venture in the People's Republic of China. Its scope is defined and it is narrow: sustainability, carbon neutrality and industrial decarbonisation, within which RRSC was selected among the Top 100 in China to co-organise the Zero-Carbon Industrial Parks White Paper of 2026, the only Western entity at coordination level on that governmental mandate. That department has not been closed and the Group does not intend to close it: it operates in a field where cooperation with the world's largest industrial economy is legitimate, valuable and, in the carbon dimension, unavoidable.

The Group accepts nothing with China in artificial intelligence or cybersecurity. No joint development, no technology transfer, no licensing, no model artefacts, no training data, no source code, no cryptographic material, no research cooperation and no commercial engagement of any kind within the perimeter of this Presentation. The position is not commercial preference. It is doctrine, it is permanent, and it is a condition of the Group's standing in the perimeters where its AI capability is engaged.

C.2 How the separation is enforced.

CORPORATE	The Chinese vehicle is a separate legal entity with its own governance, scoped by its constitutive documents to the sustainability and carbon vertical, and holding no interest of any kind in the AI and cybersecurity perimeter.
TECHNICAL	Separate infrastructure, separate credential domains and separate repositories. No source code, model artefact, training corpus, architectural documentation or cryptographic key of the AI and cyber perimeter is accessible from, transmitted to, or replicated within the Chinese perimeter.
PERSONNEL	No personnel of the Chinese perimeter hold roles, credentials or review functions in the AI and cyber programmes, and the Scientific and Technical Lead's programmes are conducted entirely outside it.
CONTRACTUAL	Engagements of the sustainability vertical carry no technology-transfer, licensing or co-development rights over any asset of the AI and cyber perimeter, and none will be granted.

C.3 Why a counterparty should care.

Section XIV documents, on public sources, the scale and coherence of the Chinese national programme in the architectural class this Presentation describes. That analysis is technical respect, and it is not ambiguity of position: the Group studies that programme as a strategic actor studies a peer, and cooperates with it not at all. For an institution in Europe, in the Gulf or within an allied procurement perimeter, the separation above is what allows the Group's technology to be considered at all in defence, critical-infrastructure and classified environments, where the origin, the custody and the reachability of a system are examined before its capability is.

The Group also regards the position as a matter of coherence. An institution that sells sovereignty cannot hold, in another room, a technology relationship that contradicts it. The separation is documented here so that it can be verified, contractually and technically, by any counterparty that wishes to verify it.

SECTION XV

The Sovereign Question · An Executive Briefing

The proposition in one page, for senior decision-makers.

The architecture of the next generation of artificial intelligence is being defined, in the present cycle, within the institutional perimeter of the People's Republic of China. Under the strategic direction of the 15th Five-Year Plan, with AI investment doubled in 2025, the Chinese ecosystem has converged on a paradigm built on three structural pillars: sparse Mixture-of-Experts architectures, linear-complexity attention, and the institutional integration of quantum computing as an AI accelerator. Alibaba, Baidu, DeepSeek, Ant Group, Zhipu AI, Xiaomi, BAAI: each contributes to a coordinated national programme pursuing the same three pillars simultaneously, with full state support.

The Western hemisphere, on the present trajectory, is structurally absent from this race. The United States laboratories have remained anchored to the dense Transformer paradigm; the European Union has not consolidated a coordinated industrial programme combining the three pillars under a single architectural direction. No Western actor presently identifiable is pursuing what China is pursuing. With one exception. The Rob Rockefeller Standard Carbon LLC, through the sovereign AI programme conducted under the scientific authority of its Lead AI System Architect, identified the same architectural direction in autonomy, years in advance of the institutional consolidation of the Chinese programme, and pursues the three pillars under an integrated proprietary architecture with one substantive and decisive distinction: it rests on a proprietary mathematical solution whose stability, convergence and sparsity properties have been demonstrated analytically. The mathematical proof precedes, and structurally enables, the empirical validation, rather than emerging from it. No actor in the present global ecosystem possesses an equivalent in this architectural class.

What this means for an institutional counterparty. States and enterprises are declaring, and funding, AI agendas of national ambition, in Europe and in the Gulf alike. An agenda of that ambition requires, sooner or later, exactly what commercial API resale cannot provide: sovereign capability, engineered inside the jurisdiction, under the client's own perimeter. The Group is the single Western industrial actor operating within the architectural class of the next generation, holder of the analytical mathematical formulation that the Chinese national programme itself lacks, engaged in advanced institutional dialogues within the European Union perimeter and with programmes of the North Atlantic alliance, and, in the European defence sector, with a European-Union-funded centre of competence of the Bulgarian Academy of Sciences in a cooperation now advancing to its second phase. An institution that engages the Department engages that capability directly, inside its own perimeter, rather than consuming intelligence as a foreign service. The architecture programme of Section II remains, in every scenario, an exclusive undertaking of the Group.

For the jurisdiction, sovereign entry into the architectural race of the next generation.
For the partner, the position at the door through which it enters.

SECTION XVI · IN THE EMIRATES, WE ARE THIS

United Arab Emirates · Institutional Architecture

The Group's institutional presence in the Emirates: a royal anchor and a resident senior leadership whose members hold concurrent mandates inside the Emirati establishment. Fully constituted, active, and ready to operate.

XVI.1 The royal anchor.

H.E. Sheikh Khalifa Bin Khalid Al Hamed, of the Al Hamed ruling family of Abu Dhabi, presides over RRSC UAE. He is the grandson of H.E. Ahmed Khalifa Al Suwaidi: statesman of the founding generation of the United Arab Emirates, closest institutional collaborator of the founding father Sheikh Zayed, and founder of the Abu Dhabi Investment Authority, the sovereign fund that defined, for the entire world, what patient sovereign capital means. As President of RRSC UAE, Sheikh Khalifa is not a ceremonial patron: he presides over the Group's Emirati institutional presence, anchors it within the sovereign perimeter of Abu Dhabi, and represents the continuity between the generation that invented the model and the present generation redeploying it into sovereign technology.

XVI.2 The resident leadership.

Dr. Adel Miran
General Director, Public and Institutional Relations · Middle East
The operational bridge between the Group and the Emirati institutional ecosystem: government relations, sovereign counterparties, and the protocol architecture without which nothing of consequence moves in the Gulf. He holds this mandate concurrently with a senior directorate of administration at Dubai Holding, the diversified global investment group of the Ruler of Dubai, and is a recognised thought leader on AI in the region.

Ms. Dolly Tabet
Deputy CEO · RRSC UAE
French executive with more than twenty-seven years of continuous service to the Al Hamed ruling family: the operational continuity at the centre of the Group's daily engagement with the family's broader perimeter.

Mr. Mohamed Al Hosani
Executive Manager · RRSC UAE
Senior background at the headquarters of ADNOC, the Abu Dhabi National Oil Company and the Emirates' sovereign energy operator, with more than twenty years of corporate strategy and group people-development experience at the federal energy level. Few institutional addresses in the Gulf carry comparable weight.

Mr. Ahmed Almuaini
Director, Branch Operations and Institutional Coordination
The Group's coordinator at the Emirati federal level. A member of a leading Emirati family with deep institutional and legal ties, connected by family relationship to H.E. Dr. Abdelrahman Almuaini, Assistant Undersecretary at the UAE Ministry of Economy and Tourism; doctorate from the University of Aberdeen. In an environment in which access is governed by protocol and continuity, the operations directorate is the mechanism by which standing becomes process.

This architecture remains fully constituted and active. It is the institutional and diplomatic constant of the Group's Emirati presence; what the Group now adds, as Section XVIII records, is the technical and operational component that completes it.

UNITED ARAB EMIRATES

Organisation Chart of the Branch

WHAT THIS PANEL RECORDS

The organisation of RRSC UAE. At the head, the Presidency of H.E. Sheikh Khalifa Bin Khalid Bin Ahmed Bin Hamed Al Hamed, alongside Mr. Valentino Loforese, Chief Executive Officer of the Group. Below, the four directorates of the branch: Ms. Dolly Tabet as Deputy Chief Executive Officer; Mr. Ahmed Almuaini, Director for branch operations and institutional coordination; Dr. Adel Miran, General Director of public and institutional relations; and Mr. Mohamed Al Hosani, Executive Manager and World Carbon project development manager.

UNITED ARAB EMIRATES

The Leadership Architecture

H.E. Sheikh Khalifa Al Hamed

PRESIDENT · RRSC UAE

Al Hamed family, ruling family of Abu Dhabi; historic ally of the Al Nahyan family. Grandson of H.E. Ahmed Khalifa Al Suwaidi — UAE President's Representative, First Minister of the Interior (1971) and founder of ADIA (~$1tn AUM). Partnership formalised through Royal Diamond Group LLC.

Dr. Adel Miran

DG INSTITUTIONAL RELATIONS · ME

Dual mandate. Concurrently Director General of Administration & Facilities Management at Dubai Holding — the global investment vehicle of the ruling Al Maktoum family and largest shareholder of Emaar Properties. A recognised AI thought leader, providing a direct sovereign bridge between top UAE governance and the Group's Middle East vision.

Ms. Dolly Tabet

DEPUTY CEO · RRSC UAE

French executive with 27+ years of continuous service to the Al Hamed ruling family. CEO of Royal Diamond Group LLC and Office Manager of Al Hamed Enterprises — the operational continuity at the centre of the institutional architecture.

Mr. Mohamed Al Hosani

EM · WORLD CARBON PROJECT

Senior background at ADNOC HQ — the UAE's sovereign energy operator. 20+ years of corporate strategy at federal energy level; anchors the carbon-neutrality vertical into UAE sovereign energy infrastructure.

Mr. Ahmed Almuaini

DIRECTOR · UAE BRANCH OPERATIONS

Federal-level coordinator. Family relationship with H.E. Dr. Abdelrahman Almuaini, Assistant Undersecretary at the UAE Ministry of Economy & Tourism. PhD, University of Aberdeen. Coordinates daily branch and ministerial engagement.

The architecture of the Emirati presence, and the principle on which it is built: every senior figure of the branch carries an additional institutional mandate outside the Group, within UAE governance, sovereign-wealth investment, the federal ministerial structures or the national energy operator. The panel 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 ADNOC background of Mr. Mohamed Al Hosani, and the federal coordination of Mr. Ahmed Almuaini. The table on the right records the branch profile: legal entity, headquarters in Abu Dhabi, Dubai presence, presidency, deputy chief executive and senior team.

UNITED ARAB EMIRATES

Dr. Adel Miran · Institutional Relations

WHAT THIS PANEL RECORDS

The institutional profile of the General Director of Public and Institutional Relations for the United Arab Emirates. The card records more than two decades of experience in public and institutional relations across the Emirates and global markets, leadership in multi-stakeholder engagement with governments and institutions, and the strategic synergy of his concurrent mandates in facilities management and administration at Dubai Holding and Nakheel. The right-hand panel records the strategic liaison functions he exercises for the Group: government relations, policy and partnerships, and regional impact.

UNITED ARAB EMIRATES

The Presidency of the Branch

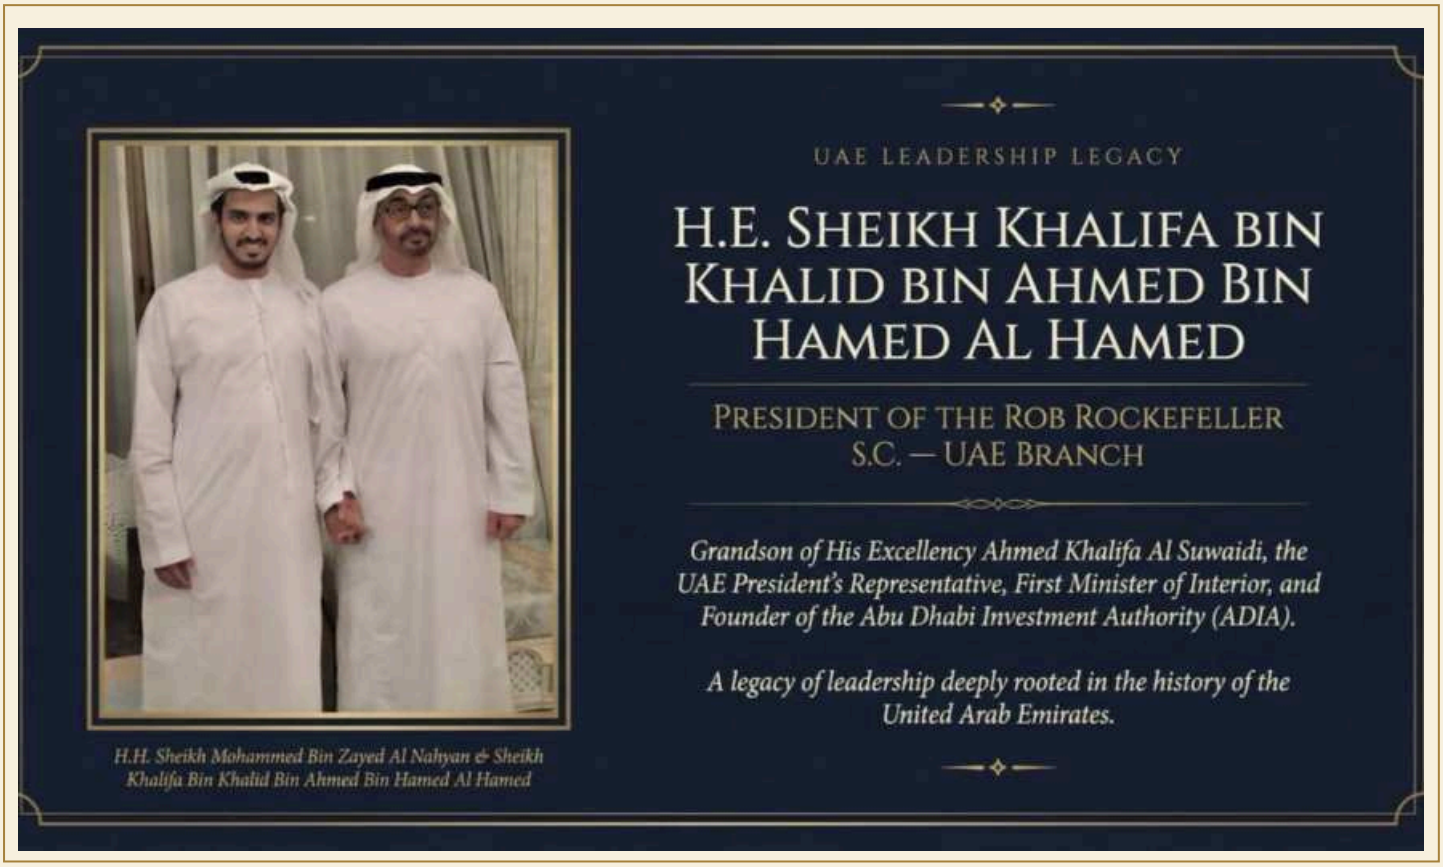

H.H. Sheikh Mohammed Bin Zayed Al Nahyan & Sheikh Khalifa Bin Khalid Bin Ahmed Bin Hamed Al Hamed

WHAT THIS PANEL RECORDS

H.E. Sheikh Khalifa Bin Khalid Bin Ahmed Bin Hamed Al Hamed, President of The Rob Rockefeller S.C. UAE branch, photographed with H.H. Sheikh Mohammed Bin Zayed Al Nahyan. The panel records the lineage on which the institutional standing of the branch rests: he is the grandson of H.E. Ahmed Khalifa Al Suwaidi, the UAE President's Representative, First Minister of the Interior, and founder of the Abu Dhabi Investment Authority.

UNITED ARAB EMIRATES

The Institutional Legacy

WHAT THIS PANEL RECORDS

The lineage behind the Emirati architecture, presented in two branches. On the left, the ruling Al Nahyan dynasty, from the founding father Sheikh Zayed bin Sultan Al Nahyan to the present leadership. On the right, the Al Suwaidi and Al Hamed families, institutional allies of the ruling house, from Ahmed Khalifa Al Suwaidi, founder of the Abu Dhabi Investment Authority, to his grandson H.E. Sheikh Khalifa Bin Khalid Al Hamed and the connection to the leadership of RRSC UAE.

SECTION XVII · IN THE EMIRATES, WE ARE BUILDING THIS

United Arab Emirates · The AI and Cyber Hub

The Emirati AI and Cyber Hub: the institutional bridge of the Group for the Middle East and Africa, and the frame within which its resident technical component is now being established.

XVII.1 A centre of institutional representation and regional delivery.

The Group's AI and cybersecurity 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. The Emirati Hub operates under a distinct doctrine: it channels the Group's technical capabilities, the sovereign AI programme, the SecurityScorecard alliance, the digital platforms delivered through Nepal, into the highest levels of regional governance and enterprise. With the establishment of the new operational centre, the Hub adds to this representational mandate a production dimension: resident engineering, joint delivery with the local partner, and the qualified operation of the Group's private models for Gulf institutional clients.

XVII.2 Strategic coverage.

The Hub is scoped to coordinate the Group's AI, cybersecurity and digital delivery across three perimeters. The GCC perimeter: the Emirati federal and sovereign-wealth structures; the Kingdom of Saudi Arabia, where the Group operates under a formal Royal Partnership with H.H. Prince Abdul Majeed bin Saud bin Muhammad Al Saud; and the State of Qatar, anchored by Sheikh Tameen Al Thani of the Al Thani Royal Family. The wider Arab perimeter: Tunisia, Libya, Egypt and the Maghreb under the Group's North African presence, with the President of RRSC North Africa, Mr. Mohamed Lassaad Kedimi. The African continental network: Ivory Coast, Congo, Senegal and Tanzania under Mr. Reda Jaber; Nigeria through the Port Harcourt presence; and the continental headquarters in Johannesburg, South Africa.

XVII.3 What the Hub delivers.

So constituted, the Hub delivers a combination few institutions in the region can assemble: sovereign AI software and the platforms that carry it, engineered, secured and operated under one governance and one quality system. 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 Hub is the address for that demand across the GCC and the African corridor.

MANDATE	Institutional representation and regional delivery for the Middle East and Africa
REACH	GCC · the wider Arab perimeter · the African continental network
ANCHOR	The Presidency of H.E. Sheikh Khalifa Bin Khalid Al Hamed and the architecture of Section XVI
STATUS	Active institutional representation. No Emirati company is yet constituted for the AI and cyber perimeter; a local technical partnership for its establishment is being sought.

XVII.4 The perimeter of the Emirati engagement, stated precisely.

H.E. Sheikh Khalifa Bin Khalid Al Hamed is the Group's partner in the United Arab Emirates and its constant support across every dimension of the Group's presence there. The company held with him, recorded at XVIII.1, is scoped to the sustainability vertical, one defined sector of the Group's wider architecture. No Emirati company has yet been constituted for the artificial-intelligence and cyber perimeter, and the Group has not to date delivered in those disciplines in the country. That perimeter is being constituted now: the Group is in discussion with prospective local technical partners able to work alongside its engineering and to develop, in the Emirates, what the Group builds elsewhere. The Presidency of the branch and its support accompany the Group across both dimensions.

SECTION XVIII · IN THE EMIRATES, THE POSITION AS IT STANDS

UAE · The Corporate and Institutional Position

The Emirati position of the Group exactly as it stands: the vehicle it holds, the institutional and diplomatic architecture that is fully constituted and active, and the technical component now being established.

XVIII.1 The joint venture in the sustainability sector.

The Group's corporate presence in the United Arab Emirates currently consists of a joint venture with H.E. Sheikh Khalifa Bin Khalid Al Hamed: RRSC Pollution & Environment Protection Services L.L.C., registered in Abu Dhabi and dedicated to environmental protection, pollution control and sustainability services. The ownership is 60 per cent The Rob Rockefeller Standard Carbon LLC and 40 per cent H.E. Sheikh Khalifa.

The position is stated with precision: the entity is validly constituted, but it is not currently operational and its commercial licence is not presently active. It is held as a vehicle of another vertical of the Group, and it is in every respect distinct from the Artificial Intelligence division and from the cooperation contemplated by the present Presentation.

XVIII.2 The institutional and diplomatic architecture.

The Group maintains in the United Arab Emirates an institutional and diplomatic architecture that is fully constituted and active, anchored at the summit of the sovereignty of Abu Dhabi through the presidency of H.E. Sheikh Khalifa Bin Khalid Al Hamed, and supported by a resident leadership of high standing whose members hold concurrent mandates within the Emirati establishment, as described in Section XVI of the present Presentation.

That architecture constitutes the permanent platform through which the Group engages the federal institutions and the sovereign actors of the region. It is the diplomatic and institutional channel through which the Group operates, and intends to operate, in the Gulf.

XVIII.3 The technical component now being established.

What is described above constitutes the institutional and diplomatic component. The Group now intends to add the technical and operational component, which represents the natural evolution of its presence in the Emirates.

The technical component is presently developed through the Group's European poles, in particular Italy with its hub of Milan and Brindisi and the nascent Augusto Innovation Hub of Milazzo, and through the certified engineering centre of Lalitpur, Nepal, which constitutes the Group's principal production platform and its Global Innovation Hub.

The Group is today in a position to produce any artificial-intelligence code and software from zero, for any sector, end to end and under sovereign control, as documented in Sections I and II of the present Presentation. That capability, which very few institutions in the world possess, is already operational and certified, and constitutes the core of the offering the Group intends to bring to the Emirates.

The step now under way is the establishment of a permanent operational centre in the Emirates, through which that engineering capability is channelled into the region: a centre that 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 global bench documented in the preceding Sections.

SECTION XVIII · CONTINUED

The Research Dimension, and a Further Opening

XVIII.4 The research dimension: the advanced AI programme.

In parallel with the productive and operational capability described above, the Group conducts a proprietary research programme on advanced artificial-intelligence architectures, under the scientific authority of its Lead AI System Architect, Mr. Umberto Colella. The programme stands at the world frontier in sparse architectures, in linear-complexity attention and in integration with quantum computing: the same architectural direction toward which the most advanced national programmes in the world are converging.

The element that distinguishes the Group's programme is the priority given, from inception, to the analytical mathematical solution of the architectural problem itself. The architecture rests on an original and proprietary mathematical formulation whose stability, convergence and sparsity properties have been demonstrated analytically, with the mathematical proof preceding and structurally enabling the subsequent empirical validation.

The programme is in an active phase of research and evolution. It is presently the object of advanced institutional dialogues within the European Union perimeter, with programmes of the Atlantic alliance, and with other European centres of competence, concerning in particular the training phase and the infrastructure it requires. In the European defence sector, a formal scientific cooperation has been executed in Sofia with a European-Union-funded centre of competence of the Bulgarian Academy of Sciences, and that cooperation is now advancing to its second phase.

In parallel, the Group maintains its own private corporate plans for the self-financing of the training phase, prepared internally and independent of any institutional instrument, so that the programme can advance on the Group's own resources and under its exclusive control should the Group so decide. Two paths are therefore under evaluation at present, the institutional and the private, and the Group will select between them on the criterion it applies to every decision of this class: the preservation, in full, of its independence and of its exclusive control over the asset.

The research programme is conducted by the Group for its own account and is not the object of the commercial cooperation contemplated here. It constitutes the strategic reserve of value on which every centre of the Group, including the future Emirati centre, is by construction qualified to draw for its institutional clients.

XVIII.5 A further dimension: sustainability.

One further consideration is owed to the reader. The sustainability dimension of the Group's Emirati presence, the joint venture of XVIII.1 presided over by H.E. Sheikh Khalifa, is a field in which, in the Emirates, everything remains to be built. That dimension sits outside the perimeter of the present Presentation and is developed by the Group within the perimeter of that vehicle and of its royal presidency, separately from the AI and Cyber Department.

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

THE MEASURE OF THE INSTITUTION

Where the Strength of the Group Actually Lies

Not in size, and not in capital. The strength of the Group is of a different kind, and it is stated here plainly so that it can be tested.

S.1 Authorship.

The Group authors what it delivers. Architecture, corpus, tokeniser, training, cryptography, platform: each is engineered inside one perimeter, by people the Group employs, under a scientific authority whose qualities and areas of specialisation are exceptionally rare in today's world of AI engineering. Authorship is the strength from which every other one follows, because only the author can transfer the system, answer for its behaviour, secure it at the layer where it is actually exposed, and continue it when the environment changes.

S.2 Independence.

The Group has been built on partnership-funded growth, without external equity dilution and without debt. No investor's timetable governs its architectural decisions, no upstream licence governs its technology, and no commercial interface governs its deployments. Independence of this kind is unusual and it is deliberate, because a supplier of sovereignty that is not itself sovereign is a contradiction the institution buying from it will eventually pay for.

S.3 The capacity to refuse.

A strength rarely stated on a capability document: the Group can afford to say no. It says no to technology cooperation with China in the AI and cyber perimeter, as Section C records. It says no to the disclosure of its architecture under any instrument, to any counterparty, however commercially attractive the instrument. It says no to pilots introduced before validation, and to claims it cannot evidence. An institution that can refuse is an institution whose commitments are worth something.

S.4 The record, and the anchors.

Behind the doctrine sits evidence a counterparty can inspect: fourteen years of certified delivery into government infrastructure, reproduced in original facsimile in Part II; a certified engineering centre in daily production; a scientific cooperation in the European defence sector advancing to its second phase; a strategic alliance with the global leader in cybersecurity ratings; and institutional anchors at royal level in three Gulf capitals. Doctrine without evidence is a brochure. Evidence without doctrine is a contractor. The Group offers both, and asks to be measured on both.

S.5 What the Group therefore is.

A concentrated institution, not a large one: modest in headcount by the standards of the integrators, immense in depth by the standards of its size. It builds artificial-intelligence software from zero, for any sector, and hands the result to the institution that commissioned it, secured, documented, operable and owned. Very few institutions in the world do this at all. Fewer still do it independently, and fewer still are willing to state, in writing and in advance, exactly what they will not do.

Authorship, independence, the capacity to refuse, and a record that can be inspected.
That is the strength. Everything else in this Presentation is its documentation.

CLOSING

Closing Considerations

What the present Presentation has documented, pole by pole, and the position from which the Department operates.

The Presentation has documented the two pillars of the Department and the poles through which they are exercised. In Nepal, the certified engineering production floor: more than fifty engineers, fourteen years of institutional delivery, and the executed references reproduced in Part II. In Italy, the European hub of Milan and Brindisi, and the Augusto Innovation Hub of Milazzo as the Group's next European flagship. In Bulgaria, the Sofia pole inside the European defence and quantum perimeter, in formal scientific cooperation with a European-Union-funded centre of competence of the Bulgarian Academy of Sciences. In the United Arab Emirates, the royal anchor and the resident institutional leadership, fully constituted and active, to which the Department is now adding its technical and productive component.

It has documented the technology base on which the work rests: the production of artificial-intelligence software from zero, for any sector, under sovereign control; the proprietary architecture programme, its analytical mathematical foundation and its present development status; the deployment security architecture that protects model artefacts at rest, in transit and at inference; and the complete cyber portfolio, offensive, defensive, architectural, cryptographic and regulatory, applied across every sector the Department serves.

And it has documented the discipline under which all of it is stated: 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. That discipline is the Department's method, and it is the standard by which it asks to be measured.

EXECUTIVE OFFICE	**Mr. Valentino Loforese** Chief Executive Officer and Majority Shareholder R. Rockefeller S.C. · The Rob Rockefeller Standard Carbon LLC
SCIENTIFIC AND TECHNICAL AUTHORITY	Mr. Umberto Colella · Chief of Staff and Lead AI System Architect
EMAIL	
WEB	therockefeller-sct.com
TELEPHONE	
REGISTERED OFFICES	8 The Green, STE R, Dover, DE 19901 · 1309 Coffeen Ave, Ste 1200-C, Sheridan, WY 82801 EIN: 93-4657669 · Delaware LLC Reg. No. 6714331

· August 2026

The Rob Rockefeller Standard Carbon LLC · Office of the Chief Executive Officer