SECURITY POSTURE BRIEF
DOC · SEC-BR-02 · REV 2026.07

OUR COMMITMENT TO CLOUD, AI & DATA SAFETY

Securing Your Data

How we design for protection, privacy & resilience

A principle-level view of the security approach DataExos builds into every layer of its cloud, AI, and data systems — the commitments that guide how we architect, operate, and improve.

SCOPE	POSTURE	ISSUED
Cloud · AI · Data	Compliance-Aware	July 2026

01 A security-first way of building

Security at DataExos is not a feature bolted on at the end — it is the frame we build inside. The commitments below describe how we design and operate our systems, not audited outcomes or certifications we claim to hold.

We start from the data. Every system we design begins with a simple question: what would it take to protect this information as if it were our own? That lens shapes architecture decisions before a single line of production code is written — where data lives, who can reach it, how it moves, and how we would know if something went wrong.

Because we work across **cloud, AI, and data** together, we treat these as one continuous surface rather than three separate problems. A model is only as trustworthy as the pipeline that feeds it; a pipeline is only as safe as the cloud it runs on. Our approach is designed to hold all three to the same standard.

P1
Defense in depth
No single control is trusted to stand alone. We layer protections so that a failure at one boundary is contained by the next.

P2
Least privilege by default
Access is granted narrowly and deliberately. The absence of a reason to grant access is treated as a reason to withhold it.

P3
Assume-breach mindset
We design as though a boundary will eventually be tested — so detection, containment, and recovery are first-class, not afterthoughts.

P4
Privacy as a design input
Data minimization and purpose-limitation are considered at design time, not retrofitted once a system is in production.

P5
Human oversight
Automation accelerates our work; people remain accountable for it. Consequential actions are designed to stay under human review.

P6
Transparency by habit
We document how our systems work and where their limits are — so trust rests on evidence, not on assurances alone.

HUMAN-IN-THE-LOOP

Our operating philosophy is captured in a phrase we hold ourselves to: **“Not AI. AI for Humans.”** Intelligent systems are built to inform and augment human judgment — never to quietly replace the accountability that protects the people whose data we hold.

02 Encryption & crypto-agility

Encryption is the baseline we design toward, not the exception. Our approach protects data through its full lifecycle — moving, at rest, and in use — while keeping the underlying cryptography replaceable as standards evolve.

ENC-01

Encryption in transit

Data moving between clients, services, and storage is designed to travel over strong, modern transport encryption, with plaintext channels treated as a design defect rather than a convenience.

ENC-02

Encryption at rest

Stored data — databases, object storage, backups — is architected to be encrypted at rest, so that access to underlying media alone does not expose readable information.

ENC-03

Key management

Keys are designed to be centrally managed, access-scoped, and rotated on a defined lifecycle, kept separate from the data they protect and away from application code.

ENC-04

Crypto-agility

Algorithms and key lengths are treated as configurable, so that cryptography can be upgraded — including toward post-quantum-ready schemes — without re-architecting the systems around it.

CONTROL	OBJECTIVE	HOW WE DESIGN FOR IT
ENC-01	**Protect data in motion**	Modern transport encryption on service-to-service and client-facing paths; secure defaults over opt-in.
ENC-02	**Protect data at rest**	Storage-layer encryption for primary data and backups; media loss does not equal data loss.
ENC-03	**Guard the keys**	Managed key store, scoped access, planned rotation, and separation of keys from ciphertext.
ENC-04	**Stay upgradeable**	Abstracted crypto interfaces so primitives can be replaced ahead of standards changes.

WHY CRYPTO-AGILITY MATTERS

The strongest algorithm today is a liability tomorrow if it cannot be replaced. Designing for change — rather than for a single fixed cipher — is how a security posture stays durable as the cryptographic landscape shifts.

03 Access control & least privilege

The safest access is the access that was never granted. We design identity and authorization so that every person, service, and system reaches only what its role genuinely requires — and no more.

ACC-01

Strong identity

Access is anchored to verified identity, with multi-factor authentication designed in for privileged and administrative paths rather than left optional.

ACC-02

Role-based authorization

Permissions map to roles and responsibilities, so entitlements are legible, reviewable, and revocable — not scattered across ad-hoc grants.

ACC-03

Secrets management

Credentials, tokens, and keys live in managed secret stores with scoped retrieval and audit — never hard-coded into source or configuration.

ACC-04

Segmentation & isolation

Networks, environments, and tenants are separated by design, so a compromise in one area does not become free movement across the whole.

› **Just-enough, just-in-time**

Elevated access is scoped and, where practical, time-bound — granted for a task and withdrawn when the task is done.

› **Reviewable by design**

Entitlements are structured to be periodically reviewed, so access reflects current need rather than accumulated history.

LEAST PRIVILEGE IN PRACTICE

Least privilege is not a one-time configuration; it is a discipline. We design authorization to make the secure choice the default choice — so access narrows naturally over a system's life instead of quietly widening.

04 Threat detection & monitoring

You cannot protect what you cannot see. Our approach treats visibility as a control in its own right — designing systems that record what they do, surface what looks wrong, and respond in a rehearsed way.

MON-01

Comprehensive logging

Systems are designed to emit meaningful, tamper-evident logs of security-relevant events, so activity can be reconstructed and questioned after the fact.

MON-02

Continuous monitoring

Telemetry is centralized and watched rather than left to accumulate unread, so signal is separated from noise while events are still actionable.

MON-03

Anomaly & threat detection

Baselines of normal behavior let unusual access patterns and anomalies stand out — the earlier a deviation surfaces, the smaller its blast radius.

MON-04

Incident response

Response is designed as a defined path — detect, contain, eradicate, recover, learn — so a real event is handled by a plan, not improvised under pressure.

CAPABILITY	WHAT IT ANSWERS	DESIGN INTENT
Audit logging	What happened, and when?	Durable, security-relevant event records built to support after-the-fact review.
Monitoring & alerting	Is something wrong right now?	Centralized telemetry with alerting tuned to reduce noise and preserve signal.
Anomaly detection	Does this look normal?	Behavioral baselines that make unusual patterns visible early.
Incident response	What do we do about it?	A rehearsed detect-contain-recover-learn cycle with clear ownership.

MONITORING WITH JUDGMENT

Detection generates signals; people decide what they mean. Consistent with **“Not AI. AI for Humans.”** automated monitoring is designed to bring the right events to human attention quickly — not to make irreversible decisions on its own.

05 Resilience, backup & recovery

Security is not only about keeping data safe from others — it is about keeping it available to you. We design for the days when things fail, so that an incident becomes an inconvenience rather than a loss.

RES-01

Resilient architecture

Redundancy and graceful degradation are designed into critical paths, so a single failure does not cascade into an outage.

RES-02

Protected backups

Backups are architected to be regular, encrypted, and isolated from primary systems — so they survive the very events they exist for.

RES-03

Recovery objectives

Recovery-time and recovery-point goals are defined per system, so restoration is measured against an intended target rather than guessed at.

1

Anticipate

Identify what could fail and what each failure would cost.

2

Protect

Back up, replicate, and isolate so recovery is always possible.

3

Restore

Recover to a known-good state against defined objectives.

4

Verify

Test recovery so plans are proven, not merely written down.

RECOVERABILITY IS A SECURITY PROPERTY

A backup that has never been restored is a hypothesis. Our approach treats tested recovery — not just the existence of backups — as the measure of resilience, so continuity plans hold up when they are actually needed.

06 Governance & standards alignment

Our controls are designed around widely recognized security and privacy frameworks. We speak of *alignment* deliberately: we architect toward these standards as targets — we do not claim to hold their certifications or attestations. The frameworks below name the bar we design against.

DOMAIN	CONTROL FOCUS	ALIGNED TO
GRC-01	**Information security management**	ISO 27001
GRC-02	**Service trust & controls**	SOC 2 SECURITY · AVAILABILITY · CONFIDENTIALITY
GRC-03	**Personal data & privacy**	GDPR
GRC-04	**Protected health information**	HIPAA
GRC-05	**Cardholder data**	PCI DSS

Governance, not just controls

Security decisions sit inside a governance model with defined ownership and review — technical controls are the expression of that model, not a substitute for it.

Human accountability

Under our human-in-the-loop model, people remain answerable for how automated and AI systems handle data — governance keeps that accountability explicit.

WHERE WE ARE — STATED PLAINLY

These five frameworks — ISO 27001, SOC 2, GDPR, HIPAA, and PCI DSS — are **alignment targets, not held certifications.** We design toward these standards and will pursue formal assessment where an engagement genuinely requires it.

07 Security is a practice, not a state

Threats change, systems evolve, and yesterday's safe assumption becomes tomorrow's gap. We treat security as an ongoing discipline — reviewed, tested, and improved on a cycle rather than declared finished.

C1 — ASSESS

Look for the gaps

Regular self-review and testing to find weaknesses before they are found for us.

C2 — IMPROVE

Close them deliberately

Findings become tracked changes, so improvement is evidenced rather than assumed.

C3 — VERIFY

Prove the fix

Changes are validated against real conditions, not signed off on intent alone.

C4 — ADAPT

Raise the baseline

As standards and threats move, our commitments move with them — the bar only goes up.

YOUR DATA. OUR COMMITMENT.

Security you can ask hard questions about.

We would rather earn trust through transparency than assert it. To discuss how we protect cloud, AI, and data systems, start a conversation.

WEB
www.dataexos.com

PHONE
1-800-460-2261