Confidencial

Next Gen Cryptographic Protection:

Protecting your data is essential. Here's how you can build the path forward today.

Data and its protection are at a critical juncture as technology and regulatory forces reshape the choices ahead. Opportunity awaits, but so does risk. The platform decisions you make now will define your resilience.

AI is bringing new urgency to a timeless struggle

If language itself is a technology, then cybersecurity began the day humans first concealed a secret. From ancient battlefields to modern enterprises, progress has always depended on moving, building, and trading securely and silently.

The urgency of the AI moment means we're sharing more sensitive data, more broadly across humans, machines, and models. That multiplies both opportunity and exposure exponentially, pushing traditional perimeter and identity-based defenses to their limits, especially when it comes to information-level security.

Attackers and defenders both know that sensitive data is valuable data. So who will win the race to control it, and how will data protection platforms make a difference?

AI is creating and complicating data protection challenges.

As AI raises the stakes on information security, best practices need to evolve, and data protection platforms need to adapt. But people, processes, and technology all have obstacles to overcome.

AI is turning data flows into rivers of risk and opportunity.

Security and compliance teams both understand what attackers already know: the most valuable data sources and the most sensitive data sources are often the same. Given this, they're ideal AI inputs **if the risk can be managed.**

3 Reasons to Worry About Your Encryption

1. AI systems depend on vast, diverse datasets to function effectively.
2. Regulators are tightening expectations, requiring encryption, auditability, and proactive governance.
3. Quantum computing remains on the horizon, but its future impact on current encryption demands action today.

- Organizations fear repeating past cloud migration mistakes, trying to secure a new frontier with legacy tools.
- New AI/MLOps pipelines are the frontline for data risk; security must be embedded directly into these workflows.

Hybrid workplaces (and cloud environments) are here to stay.

Whether you're WFH or RTO, the new normal feels permanently hybrid. With users connecting from more locations and an increased reliance on rapidly multiplying edges, AI's insatiable appetite for data means building and delivering security in new ways.

- AI is getting pushed far and wide; as both inferencing and training workloads get accelerated.
- This complicates security obligations, since the who/where/how of data storage dramatically impacts compliance and regulatory burdens.

Regulators and cyber insurers are intensifying scrutiny.

Two things are happening at once: rules are becoming more prescriptive, and regulators are expecting organizations to become more proactive, especially around data protection. This adds up to much tighter operating constraints, especially for heavily regulated industries.

- New governance updates emphasize continuous risk assessment and control testing. Frameworks like NIST CSF 2.0 and ISO 27001:2022 now require demonstrable encryption and access control.
- Regulators are now pushing down specific control requirements around everything from encryption to multi-factor authentication.

Quantum's shadow has defenders feeling gloomy.

Much like AI, quantum computing has been a long time coming. Just as AI researchers needed the GPU's parallel processing to make LLMs finally work, quantum theorists have been waiting for today's quantum computers, and the results so far are impressive, if not unsettling.

- Experts warn that traditional algorithms (RSA, ECC) will become vulnerable as quantum computing advances. Planning for post-quantum cryptography (PQC) should start now.
- Luckily, researchers have been busy building ways to defend data in a post-quantum era. Many PQC standards are deployable today, without waiting for a cryptographically relevant quantum computer (CRQC).

Even as end-to-end encryption faces mounting regulatory pressure.

Finally, as governments and vendors clash over compliance expectations, huge pieces of security infrastructure are now at risk. End-to-end encryption is no longer a given. As the Apple and UK standoff showed, governments can force tech companies to water down encryption, or at least fight hard to do so.

- Disruptions will leave organizations playing catch-up.
- Decision-makers must design architectures that keep data protected, no matter where it moves or who accesses it.

3 Questions to Ask About Your Platform

1. *Do we have data visibility gaps?*
2. *How fragile is our stack?*
3. *Are we AI-ready?*

Data protection stacks are getting taller, but still coming up short.

Even as challenges multiply in volume and variety, security and compliance teams (along with everybody else) are looking at their current data protection platforms and wondering if they're up to the challenge.

- Can they enable secure, responsible AI?
- Do they protect data persistently, at rest, in transit, and in use?
- Will they scale to meet emerging threats?

The answers to these questions must shape what comes next. Not just for businesses, but platform builders and integrators too, as AI will magnify every weakness in today's fragmented protection stacks.

DSPM (Data Security Posture Management): Great at seeing but not doing.

DSPM tools locate sensitive data across environments and can configure access at the folder level, but this control is still too coarse-grained. Once data moves, is downloaded, or shared, it's no longer protected.

Knowing where your crown jewels are is critical, but without enforcement, monitoring, or response, discovery alone can leave them exposed. DSPM highlights risk; real protection requires persistent controls that travel with the data.

DLP (Data Loss Protection):Often too much too soon, and too little too late.

DLP aims to block data leaks but disrupts workflows through restrictive policies that slow legitimate access, create document delays, and generate false positives that overwhelm security and IT.

DLP can also struggle with hybrid cloud and AI-driven data flows, leaving critical gaps in both protection and productivity.

Finally, users and security face constant interruptions and workarounds.

Traditional All-or-Nothing Encryption: Forces too many choices, too early.

Traditional encryption secures all data in a file equally: if you have access, you see everything; if not, nothing is accessible. This offers no nuance between public and protected fields.

It lacks precision and persistent policy controls, making it poorly suited for multi-cloud, multi-agent environments where granular, adaptive, and enduring protection and sharing are required.

Three truths to guide your journey:

→ Visibility must go wide and deep.

→ Granular is always greater.

→ Selective is effective.

As threats, trends, and tools change, the journey remains fundamental.

Even as external and internal forces combine to create all kinds of disruption, some things remain foundational, including your data protection journey. But fundamental doesn't mean inflexible. This 'journey' must be revisited with these new world challenges in mind.

1 DISCOVERY CREATES VISIBILITY

Discovery now extends across structured, unstructured, and AI-generated data. Inputs and outputs alike must be scanned for sensitive content. And the pressure for real visibility is on, from regulators, underwriters, and business partners. Blind spots and dark corners are absolutely unacceptable.

But it's not just the breadth of discovery that's changing, it's also depth, especially with AI. Being able to look inside files and understand risk at the object level is critical.

2 CLASSIFICATION SETS CONTEXT

Once discovered, information must be classified, typically by risk level. This will become the basis for future decision making about the information, especially if it's 'a secret'. But classifying isn't about the data, but rather its value to the enterprise. The best plan is often to start at your crown jewels and work your way out from there.

The same granular discovery should translate into more granular classification of files and the data they contain. In reality, a single file can contain both safe and sensitive fields.

3 PROTECTION GIVES CONTROL

Data protection, typically through encryption, operationalizes some of those classification decisions. This also includes the access design required to decide which teams and users (and AI agents) need/want the information. The final result should be control embedded into the data itself.

That granular discovery and classification can then be operationalized in selective encryption. This means individual fields, not the whole file, can be individually managed and protected.

4 AI GOVERNANCE PUTS IT ALL TOGETHER

Finally, at the end of the journey, we end at governance that maps business and compliance needs to technical controls, especially when building for AI. The decisions made here will then be translated into technical controls and encryption choices.

Ultimate, granular classification and encryption enable a more dynamic and precise data protection strategy. This is how you can strike the secure/shareable balance that's critical for the current AI moment.

3 Priorities to Balance

- *Usability (speed and minimal friction)*
- *Persistent protection (always on)*
- *Auditability (provable control)*

Building for new data protection fundamentals.

If the old ways are failing, what kinds of paradigms and best practices can replace them? We've already examined the data protection journey and how granularity and precision make a significant difference. But how else can organizations start building for next-gen cryptographic protection and observability?

Precision drives orchestration. Organizations must protect sensitive data while keeping non-sensitive data AI-ready.

We already saw the power of greater precision throughout the data protection journey. The ability to classify and protect at the object level lets you be much more deliberate about access. This fine-grained approach is also ideal for fast, dynamic, AI/ML pipelines.

Thinking and defending with this granularity doesn't end there. A precise, granular data inventory also lets you really zoom in on crown jewels, especially in an era of "harvest now, decrypt later". You need to be able to find your most sensitive secrets and defend them with everything you have.

Whether it's emerging quantum threats or aggressive regulators, your ability to zoom in and bring a granular "risk-value" view to your data is becoming increasingly critical. Imagine only being able to manage RBAC at the team level, not the user level. Why accept that level of clumsy management with your encryption?

Simpler is better, but elegant integration ensures new protections complement existing systems, not replace them.

This one is straightforward: simpler is always better for both technology and security teams. This is why organizations happily invest in DSPM and other large platforms, reducing the time required to tinker under the hood to get everything to work. Integration, customization, and middleware can all be slow and costly.

But organizations have robust security and compliance “systems of record” in place, and they often represent serious investments in time and technology. Ultimately, the costs of switching or ripping and replacing are usually too high, so next-gen solutions need to be able to mesh with existing tooling.

Elegant integration must be the goal:

- Organizations can still get the most from DSPM and DLP technology investments.
- They can also build on existing discovery and classification work.

Autonomy and persistence unlock protection that won’t and can’t quit. This keeps organizations in control.

If you can’t trust any one else to truly safeguard your data, what now? Persistent, policy-based encryption ensures protection travels with your data wherever it goes. Access policies are embedded at the metadata level, while keys remain securely managed to protect sensitive content across multi-cloud, hybrid, and collaborative environments.

This self-enforcing, data-centric approach ensures protection persists, even as information is shared externally or moves between all the applications, devices, and users. Access controls remain inseparable from content and adapt automatically, potentially limiting insider threats and exposure risk while boosting ongoing compliance.

This is where AI benefits defenders:

- Automation and **ML-driven classification** can automatically detect and trigger encryption policies for sensitive content.
- Organizations can replace old data operations with more nimble, modern orchestration.

It should all add up to more dynamic data protection defenses.

Granular encryption enables adaptive defenses aligned with Zero Trust principles, verifying every access, human or machine. By classifying and securing specific data objects, organizations ensure that only authorized human or machine users can reach sensitive content. This is well-suited for fast-moving AI and ML pipelines, where access needs and risk levels change in real time.

Integrating access control directly with encryption at a granular level transforms how and when information can be used. Access attempts are continuously evaluated, and decryption only happens if every live attribute (e.g., user identity, device status, contextual signals) meets the current policy. This operationalizes Zero Trust's core guidance of overwhelming verification.

Again, AI helps here, with automation and intelligent policy engines becoming essential to this outcome.

- As data is created or changed, these systems can detect, classify, and encrypt sensitive content instantly, with no human intervention required.
- Conditional encryption policies adapt automatically to changing roles, regulations, and threats.

5 Immediate Steps

1. *Assess coverage breadth*
2. *Review control depth*
3. *Test across the data lifecycle*
4. *Validate solution integration*
5. *Model for real world risks*

An Action Plan for CISOs and IT Leaders

Use this framework to assess visibility, precision, and persistence.

Evaluating your current security and data protection stack starts with an objective look at your current platforms and then identifying where they might fall short.

1. Assess Data Coverage

Start by pinpointing what types of data are covered by your existing tools.

Are your protections limited to perimeter security, or do they extend deeply into applications, unstructured files, emails, and collaboration platforms where critical data lives?

2. Review Control Granularity

How finely can you control access? Can you enforce rules at the level of specific document sections or data fields, or are you stuck with broad, "all-or-nothing" permissions? Today's environments demand precise, policy-driven encryption and micro-segmentation.

3. Test Your Lifecycle Protection

Follow your data's journey. Does protection move with your information as it crosses users, devices, clouds, and apps? True security embeds cryptographic controls inside the data itself instead of relying solely on infrastructure boundaries.

4. Check Integration and Automation

Evaluate whether new protection tools mesh well with your security and productivity solutions. The best ones deploy easily, automate classification and encryption, and require no disruptive "rip and replace" efforts.

5. Model Insider and SaaS Risks

Finally, run scenarios to see if insiders or compromised third-party SaaS services can steal, decrypt, or misuse data. Your current stack should detect, block, or mitigate these insider and external threat vectors effectively.

You'll probably find that all your hard work in data classification and even discussions around data and AI governance will pay off here.

Ask the right questions of vendors and internal teams

Here are some questions to ask to drill down even further into your current data protection needs and wants.

Core Coverage and Capability

- What types of data are covered by your platform (structured, unstructured, documents, emails, collaboration tools)?
- Does your solution provide persistent protection at rest, in transit, and in use, including documents shared externally or stored in SaaS or cloud applications and during AI training, inference, and collaboration?

- Can you specify which data elements need to be protected, down to individual fields, paragraphs, or document objects, or are controls applied only at a file or container level?

Policy and Access Control

- How granular are your access controls? Can different roles, teams, and users have differentiated access to specific parts of the same file?
- Are encryption and access policies adaptive and able to update automatically as roles, risk, or compliance needs change?
- Can policies be managed and enforced both centrally and contextually, such as based on user, device, or location?

Integration and Automation

- How does your platform integrate with existing identity and access management tools like Active Directory, Okta, or SAML and business applications such as Word, Excel, Outlook, Teams, and Slack?
- Is the discovery, classification, and protection of sensitive data fully automated, and can it keep up with high volumes or velocity of data creation?
- Can historical or archived content be discovered and protected automatically including legacy repositories and backups, or is manual intervention required?

Data Lifecycle and Incident Response

- Does protection follow the data, even as it is shared inside and outside the enterprise, across clouds, and between users or applications?
- What audit capabilities exist to track who has accessed protected content and to support compliance or forensics needs?
- How quickly can you investigate and mitigate incidents such as unauthorized access or data leakage?

Usability, Deployment, and Operations

- How easy is the solution for end-users? Does it add significant friction, require complex training, or operate transparently in the background?
- What deployment models are supported (SaaS, on-premises, hybrid), and how flexible is your architecture as infrastructure evolves?
- What is the performance and overhead impact on documents and collaboration workflows?

Zero Trust and Next-Generation Protection

- How do you extend Zero Trust principles to the data layer and not just network or session?
- Are cryptographic keys and policies always under customer or enterprise control, never exposed to vendors?
- How do you support evolving requirements such as post-quantum cryptography, universal SaaS integration, and regulatory mandates?

Asking these questions will help you uncover gaps in your current stack, assess true coverage and persistence, and ensure next-generation data protection and compliance for dynamic, distributed, multi-cloud environments.

1 Partner to Trust

1. *All data, all environments*
2. *Unique selective encryption*
3. *Persistent, portable, AI-ready protection*
4. *Traceable and protected beyond your infrastructure*

Confidencial knows control is always your most critical asset

The world is changing, but you're still in control of data protection and your security and compliance posture as a whole. Confidencial knows that stronger, smarter data protection is critical to this control:

- Regulators hold you responsible for your own risk, plus that of third parties.
- Governments or vendors can make sudden decisions that disrupt everything.
- The quantum computing threat to current classical encryption standards is real, though its timing remains uncertain.
- Successful AI means strong AI governance.
- Greater visibility and control are the only way forward, ensuring protection remains persistent, auditable, and AI-ready. Your data, encrypted in a way that keeps it protected and visible. No matter where it goes, you and your policy stay in control.

State-of-the-art selective encryption protects sensitive data without locking down workflows. Our end-to-end platform lets you discover, classify, and protect across cloud and on-prem stores, all in one place. It's not just about data protection as an idea, but solving the real-world data protection challenges keeping leaders up at night.

- Fine-grained, policy-driven encryption and access at the object level
- Policies travel with your data everywhere
- Integrates with current workflows and document tools
- Audits and tracks access for forensics
- Defends against insider and SaaS cloud misuse
- Integrates directly into existing tools and workflows.

Deploy Your Way: On-Premises, Cloud, or Hybrid

Confidencial's platform is designed to support data sovereignty and compliance mandates and fit your deployment preferences: on-premises, in your private cloud, or as a hybrid solution with no forced SaaS dependency. Organizations retain full control of their data and keys, making it possible to align security with infrastructure needs and compliance requirements.

Two paths to selective, effective encryption

Confidencial accelerates your data protection maturity. Customers can deploy Confidencial in two models:

1. Leverage existing DSPM outputs for instant remediation

Customers with a DSPM in place can snap in selective encryption by simply integrating with our encryption engines. This builds on existing data classification work and policy by adding powerful selective encryption to discovered data.

2. Integrates classification, encryption, policy orchestration, and tracing into a unified data protection fabric.

Deploying the full Confidencial platform gives you a "DSPM+" that enables you to discover, classify, protect and track unstructured data across your entire environment. This gives you full access to the benefits of selective, persistent encryption. Once data is discovered, you can build and integrate policy.

Made for this moment & your momentum

Confidencial empowers organizations to protect sensitive data without compromising usability, compliance, or speed. Our platform offers fine-grained encryption at the object level and persistent policy controls that travel with your data, ensuring robust security across clouds, users, and AI workflows.

We've been building for this moment since our inception. As we move into what's next for data and encryption, we look forward to helping you take that next smart step.

Get your free Data Risk Assessment →
www.confidencial.io/assessment

Compliance Alignment Snapshot

How Selective Encryption and AI-Aware Data Protection Map to Modern Frameworks

Organizations face growing regulatory pressure to ensure persistent encryption, access control, and auditability across AI pipelines, hybrid environments, and unstructured data. Confidencial's selective encryption and policy-based protection model align with leading global standards.

Framework / Regulation	Relevant Requirement / Control	How Confidencial Helps You Comply
NIST Cybersecurity Framework (CSF) 2.0	*PR.DS-1, PR.DS-2:* Data-at-rest and data-in-transit are protected. *PR.AC-4:* Access permissions are managed and reviewed. *GV.SC-02:* AI and emerging tech are governed.	Applies **object-level encryption** and **persistent access controls**; integrates with IAM for continuous policy enforcement; supports governance of AI-driven data flows.
ISO/IEC 27001:2022	*A.8.24 & A.8.25:* Cryptographic controls and key management. *A.5.7:* Threat intelligence and adaptive security.	Implements **fine-grained, policy-based encryption** with enterprise-controlled keys; supports **post-quantum readiness** and adaptive controls that evolve with risk.
NIST AI Risk Management Framework (AI RMF)	*Map & Govern functions:* Identify and manage AI risks, data protection, transparency.	Enables **AI data governance** by discovering, classifying, and protecting sensitive inputs/outputs across LLM, RAG, and fine-tuning pipelines.
ISO/IEC 42001 (AI Management System)	*8.2 & 8.3:* Data governance and security of AI systems.	Provides **auditable encryption** and **usage traceability** across AI models and data pipelines.
GDPR	*Art. 32:* Security of processing (encryption, pseudonymization). *Art. 5(1)(f):* Integrity and confidentiality.	Applies **selective encryption** and **semantics-preserving tokenization** to meet encryption and pseudonymization requirements while retaining usability.
HIPAA (Security Rule)	*§164.312(a)(2)(iv):* Encryption and decryption controls. *§164.312(b):* Audit controls.	Encrypts **PHI at the object level**, embeds policies for role-based access, and **logs all access/decryption events** for compliance.
PCI DSS 4.0	*Req. 3.5, 3.6:* Key management; *Req. 4.2:* Strong encryption for transmissions.	Centralized **key ownership**; **policy-based encryption** that persists across data lifecycle; integrates with DSPM/DLP for coverage validation.
EU AI Act (2024)	*Art. 10:* Data governance and quality; *Art. 15:* Robustness, security, and accuracy.	Ensures **data provenance and protection** across training and inference; enables **secure sharing and auditability** to meet governance obligations.
FedRAMP / DoD SRG (U.S.)	*AC-3, SC-13, SC-28:* Access enforcement, cryptographic protection, protection of information at rest.	Enforces **least-privilege access** and **enduring encryption** within controlled boundaries; supports hybrid deployment for **data residency**.