

The Harvest Now, Decrypt Later Model

Understanding the Emerging Quantum Risk and What Organizations Must Do Now



The Model Behind the Risk

Harvest Now, Decrypt Later

Traditional security models are built on a straightforward assumption: data must be protected now it is accessed, and encryption prevents unauthorized use. Within this model, the primary question organizations ask is whether their data can be broken today.

The “harvest now, decrypt later” model challenges that assumption at its core. It introduces a separation between when data is collected and when it is ultimately exploited. In this model, encrypted data is intercepted and stored today, retained over time, and only decrypted when sufficient computational capability becomes available.



This shift changes the fundamental question of security. It is no longer whether encryption can be broken today, but whether data will remain protected across its full lifecycle.

While timelines for large-scale quantum capabilities remain uncertain, the long lifecycle of sensitive data requires action before those capabilities are realized.

Encryption, as it is commonly implemented, protects the content of data but does not prevent access to the encrypted form of that data. As information moves across networks, it traverses multiple systems, routing layers, and infrastructure components. At each stage, encrypted payloads can be copied, captured, and stored without needing to be decrypted. This is not a theoretical edge case. It is a natural byproduct of how modern digital systems operate.

In practical terms, this means that encrypted network traffic, virtual private network communications, and even encrypted archives can be collected and retained at scale. The presence of encryption delays access to the underlying information, but it does not prevent acquisition. This distinction is subtle, but it is critical. It enables a model in which data can be harvested long before it can be understood.

The viability of this approach is closely tied to the concept of data longevity. Not all data carries the same long-term value. Some information, such as real-time operational telemetry, loses relevance within minutes or hours. Other data, including financial transactions, may retain value for weeks or months. However, certain categories of data, such as identity credentials, system architecture, or national security information, can retain value for years or even decades. It is this long-life data that becomes the primary target for collection under this model.

Advances in storage economics have further reinforced the feasibility of this approach. The cost of storing large volumes of data has decreased dramatically, while distributed storage systems have made it possible to retain and organize vast datasets over extended periods. What was once cost-prohibitive is now routine. As a result, retaining encrypted data indefinitely is no longer a constraint. It is an expected capability.

This leads to an important strategic insight. Adversaries do not need to possess advanced decryption capabilities at the time of collection. They only need the ability to access data and store it. The cost of collecting and retaining encrypted information is relatively low, while the potential future value is high. This creates a low-risk, high-upside model in which data is accumulated now with the expectation that it will become accessible later.

The implications of this model extend beyond technical considerations. It introduces the concept of delayed breach, where the consequences of data exposure may not be realized until years after the initial compromise. It also creates the potential for retroactive exposure, where data that was once considered secure becomes vulnerable as computational capabilities evolve. For organizations handling sensitive or long-life data, this represents a shift from point-in-time protection to lifecycle protection.

Understanding this model requires a change in perspective. A breach should no longer be defined solely by when data is accessed or exfiltrated. Instead, it should be understood in terms of when



that data becomes readable and usable. In this context, the moment of greatest risk is not when data is stolen, but when it can finally be decrypted.

How Encryption Actually Works in Real Systems

Architecture, Key Management, and Where It Breaks

Encryption is often described as a control that can be applied to data. In practice, it is not a single control or a single implementation. It is a distributed system that spans applications, infrastructure, networks, and identity frameworks. Understanding how encryption works in real environments requires looking beyond algorithms and examining how data actually moves through systems.

When data is transmitted between systems, it is typically protected using protocols such as Transport Layer Security. From a high-level perspective, this process appears straightforward. A client connects to a server, verifies its identity through a certificate, and establishes a secure session. Behind this exchange, however, multiple steps occur that introduce both protection and exposure.

During session establishment, asymmetric cryptography is used to exchange information securely and derive a shared secret. Once that secret is established, symmetric encryption is used for the remainder of the session to improve performance. This combination of asymmetric and symmetric methods allows systems to communicate securely at scale.

What is often overlooked is how many times data is decrypted and re-encrypted along the way. In a modern enterprise environment, data rarely travels directly from a user to a destination system. It may pass through load balancers, proxy layers, inspection points, and logging systems. At each of these points, encrypted data is often decrypted for processing and then encrypted again before continuing its path.

This means that even when encryption is used consistently, data exists in an unencrypted form at multiple points within the system. These transitions are necessary for functionality, but they also create exposure that is not always visible or tracked. The security of the system depends not only on encryption in transit, but on how these intermediate systems handle data.

Data at rest presents a different but equally important set of challenges. Organizations often assume that encrypting storage systems, databases, or disks provides sufficient protection. While these controls are important, they are frequently implemented in ways that reduce their effectiveness.

One of the most common issues is the proximity of encryption keys to the data they protect. In many environments, keys are stored within the same system or accessible through the same control plane. If an attacker gains access to the system, they may gain access to both the encrypted data and the keys required to decrypt it. In this scenario, encryption provides limited protection.



Another consideration is how applications use data. Even when data is encrypted at rest, it must be decrypted when accessed by an application. During this time, it exists in memory in an unencrypted form. Logs, debugging processes, and system integrations can also inadvertently expose sensitive data. These operational realities mean that encryption at rest is only one part of the overall protection model.

Backup systems represent one of the most significant and often overlooked areas of exposure. Backups are designed to preserve data, which means they frequently contain complete datasets, including historical records and sensitive information. They are retained for long periods of time and are often subject to less monitoring than production systems.

In many cases, backup environments rely on long-lived keys or static encryption configurations. This combination of high-value data, long retention periods, and limited oversight creates a concentrated risk. From the perspective of a long-term collection strategy, backup systems are among the most attractive targets.

Beyond data storage and transmission, encryption also underpins identity and trust across modern systems. Public key infrastructure, certificates, and authentication mechanisms rely on cryptographic methods to establish and verify trust. When a system validates a certificate or authenticates a user, it is relying on the integrity of these cryptographic foundations.

If these mechanisms are compromised, the impact extends beyond data exposure. It affects the ability to trust systems, verify identities, and secure communications. This is why the discussion of cryptography must include not only data protection but also trust infrastructure.

At the center of all these systems is key management. Encryption is only as strong as the processes used to generate, store, distribute, and rotate keys. Different organizations adopt different models, ranging from provider-managed keys to customer-managed systems and externally controlled key infrastructure.

Each model introduces tradeoffs between control and complexity. Provider-managed keys reduce operational burden but require trust in the provider's environment. Customer-managed and external key systems increase control but introduce integration and operational challenges. Regardless of the model, weak key management practices can undermine otherwise strong encryption.

In real-world environments, encryption failures are rarely caused by weak algorithms. They are more often the result of how encryption is implemented and managed. Hardcoded keys, outdated libraries, inconsistent practices across systems, and lack of lifecycle management all contribute to gaps in protection.

Another important factor is the concept of trust boundaries. Data does not remain within a single controlled environment. It moves between systems, across cloud and on-prem environments, and through partner integrations. Each transition introduces a new boundary where assumptions about control and security may no longer hold.



An organization may have strong encryption within its own environment, but weaker controls at the boundaries where data is shared or transferred. These boundary conditions are often where exposure occurs.

Taken together, these factors highlight an important reality. Encryption is not a single control that can be enabled and assumed to be effective. It is a distributed system with dependencies across multiple layers of architecture and operations. Its effectiveness depends on how consistently it is implemented, how well it is managed, and how clearly it is understood across the organization.

The most significant risks do not arise because encryption is inherently weak. They arise because encryption is complex, distributed, and often only partially visible.

How Data Is Collected and Stored in Practice

Collection Paths, Aggregation, and Long-Term Retention

Data collection is often imagined as a targeted action, carried out against a specific system at a specific moment in time. In modern environments, this is rarely the case. Collection is better understood as a continuous process that occurs across networks, systems, and platforms simultaneously. It is not defined by a single point of access, but by a series of opportunities that exist wherever data is created, transmitted, processed, or stored.

One of the most scalable forms of collection occurs at the network level. Data moving between systems must traverse infrastructure that includes enterprise boundaries, cloud interconnects, and broader routing environments. Even when that data is encrypted, the encrypted traffic itself can be captured and retained. This does not require decryption. It requires only access to the data as it moves through the network.

This distinction is important. Decrypting data may be difficult, but collecting it is not. As long as data is transmitted, it can be observed, copied, and stored. In this sense, encryption protects content but does not prevent the creation of a stored record of that content in its encrypted form.

Beyond the network, collection frequently occurs through direct access to systems. Endpoints, servers, and application environments represent points where data exists before it is encrypted or after it is decrypted. Access at this level does not require breaking encryption at all. It bypasses it entirely.

For example, when a user accesses a secure application, the data is decrypted within the application environment so it can be processed and displayed. If that environment is compromised, the data is exposed in its usable form. Similarly, encryption keys are often present in memory during active sessions, creating another potential point of access. These realities highlight a limitation of focusing solely on encryption as a control. Protection depends on the entire system, not just the algorithm.



Modern architectures introduce additional collection paths through applications and APIs. Data rarely moves in a simple, linear path. Instead, it flows through service layers, microservices, and integration points that process and transform it along the way. These systems are often highly trusted and optimized for performance, which can lead to less scrutiny compared to external-facing controls.

In practice, this means that sensitive data may pass through multiple internal services, be logged for debugging or monitoring, and be replicated across environments. Each of these steps creates an opportunity for collection. The more complex the system, the more opportunities exist.

The expansion of cloud computing and third-party services further extends the collection surface. Organizations routinely share data with external platforms, rely on managed services, and integrate with partner systems. This creates a distributed environment in which data is no longer contained within a single organizational boundary.

As a result, exposure is often indirect. It may not originate within an organization's own systems, but within those of a vendor or partner. This introduces a dependency on external security practices and increases the difficulty of maintaining consistent control.

As data is collected from these various sources, its value increases when it is aggregated. Individual pieces of data may have limited meaning on their own. When combined, they can reveal patterns, relationships, and context that are not visible in isolation.

Aggregation typically occurs in systems such as data lakes, analytics platforms, centralized logging environments, and backup repositories. These systems are designed to consolidate information, which makes them highly valuable from a collection perspective. They often contain historical data, multiple data types, and records from across the organization.

The concentration of data in these systems creates a different kind of risk. Rather than targeting individual systems, an adversary can focus on environments that provide access to large volumes of information at once. This shifts the focus from isolated access to high-density targets.

Underlying all of this is the role of storage. The ability to retain collected data over long periods of time is what enables the "harvest now, decrypt later" model. Advances in storage technology have made it possible to store large volumes of data at relatively low cost. Distributed storage systems allow data to be replicated, indexed, and retrieved efficiently.

This changes the economics of collection. It is no longer necessary to decide in advance which data will be valuable. Data can be collected broadly and stored until its value can be determined later. The decision to exploit the data can be deferred.

Even without decryption, collected data can still provide value through metadata. Information about communication patterns, timing, system interactions, and relationships can reveal insights into behavior and structure. In some cases, this metadata is sufficient to support analysis without ever accessing the underlying content.



Taken together, these factors illustrate that data collection is not opportunistic or isolated. It is systematic and scalable. It leverages the natural behavior of modern systems, the distribution of data across environments, and the declining cost of storage.

For organizations, this creates a fundamental shift in how exposure should be understood. Data does not need to be immediately usable to be valuable. It only needs to be accessible and retain its relevance over time.

The implication is that exposure begins earlier than most organizations assume. It begins at the point where data can be accessed and retained, not at the point where it is decrypted. This reinforces the need to think about security not only in terms of immediate threats, but in terms of long-term data lifecycle and future accessibility.

What Data Is Most at Risk

Time Value, Mission Impact, and Practical Prioritization

Organizations often approach data protection with the assumption that all sensitive data should be treated equally. While this is a reasonable starting point, it is not practical in complex environments where data volumes are large, systems are distributed, and resources are limited. In practice, risk is not uniform across all data. It varies based on both the sensitivity of the information and the length of time that information remains valuable.

Understanding this distinction is critical in the context of long-term cryptographic risk. The “harvest now, decrypt later” model does not target all data equally. It prioritizes data that will retain value well into the future. This introduces the concept of long-life data, which becomes the primary driver of risk in this model.

Long-life data refers to information that remains relevant, useful, or exploitable over extended periods of time. Unlike transient data that loses value quickly, long-life data continues to provide insight or advantage long after it is created. This includes information that reveals identity, system design, operational intent, or strategic positioning.

For example, real-time operational telemetry may be highly valuable in the moment, but it often loses relevance within minutes or hours. Financial transactions may retain value for a longer period, particularly for auditing or fraud detection, but their usefulness typically declines over time. In contrast, identity data, system architecture, and certain forms of research or intelligence can remain valuable for years or even decades.

This difference in time-value fundamentally shapes how data is targeted. Data that remains useful over long periods becomes a candidate for collection, even if it cannot be immediately decrypted. The assumption is not that the data will be useful today, but that it will be useful when access becomes possible.



Several categories of data consistently fall into this high-risk, long-life category. National security and defense-related information is one example. Operational plans, communications, and system capabilities can reveal strategic intent and vulnerabilities long after they are created. Exposure of this data does not simply reveal historical information. It can inform future actions and decisions.

Critical infrastructure data presents a similar risk profile. Information related to energy systems, transportation networks, industrial control systems, and utilities often reflects how essential systems operate. This type of data can be used to understand dependencies, identify points of failure, and plan disruptions. Because infrastructure systems evolve slowly, the value of this data can persist for extended periods.

Intellectual property and research data also represent long-life value. Proprietary designs, algorithms, and research findings can provide competitive or strategic advantage long after they are developed. In some cases, the impact of exposure may not be immediately visible, but may influence outcomes years later.

Identity and credential data are particularly sensitive because of their durability. Unlike other types of data, identity information does not expire in the same way. Credentials can be reused, identities can be impersonated, and relationships can be mapped over time. Even partial exposure can enable further access or compromise.

Financial and transactional data, while often considered short-term, can also carry longer-term value when aggregated. Patterns of behavior, relationships between entities, and historical trends can all be derived from transaction data. This becomes especially relevant when combined with other data sources.

This introduces another important dimension of risk: aggregation. Data does not exist in isolation. Its value often increases when it is combined with other data. Individual records may appear insignificant, but when aggregated, they can reveal patterns, relationships, and context that are far more meaningful.

For example, combining identity data with transaction history can reveal behavior. Linking communication records with operational data can expose coordination and intent. Correlating system logs with user activity can uncover patterns that are not visible within any single dataset. In this way, the value of aggregated data is often greater than the sum of its parts.

High-risk data is also not confined to a single location. It is distributed across production systems, backup environments, analytics platforms, and partner systems. In many cases, the same data exists in multiple forms across different environments. A dataset that is tightly controlled in a production system may be less protected in a backup or testing environment.

Backup and archival systems are particularly important in this context. These systems are designed to preserve data over time, which means they often contain complete historical records. They are typically retained for long periods and may not be subject to the same level of



monitoring as active systems. This combination of high-value data and reduced visibility makes them a focal point for long-term risk.

Despite the importance of identifying high-risk data, many organizations struggle with prioritization. Data classification efforts may be incomplete or inconsistent. Ownership of data may be fragmented across teams. Visibility into where data resides and how it is used may be limited. As a result, organizations often apply broad controls rather than targeted protections.

A more practical approach is to evaluate data based on a small set of guiding questions. How long must this data remain secure? What is the impact if it is exposed in the future? Where does this data reside, and how does it move across systems? These questions help shift the focus from general sensitivity to long-term risk.

Not all data requires the same level of attention. Prioritization allows organizations to focus resources where they will have the greatest impact. Data that is both high in sensitivity and long in lifespan should be addressed first, particularly if it is widely distributed across systems.

It is also important to recognize that some of the most critical data is not immediately obvious. System configuration data, encryption keys, certificates, and internal communications can all provide insight into how systems operate. These elements may not be classified as sensitive in traditional models, but they can enable broader access or reveal structural vulnerabilities.

Ultimately, the challenge is not simply identifying sensitive data. It is understanding which data must remain protected over time and ensuring that it is managed accordingly. This requires a shift from focusing on immediate protection to considering the full lifecycle of data.

The most important data is not necessarily the data that is most active or most visible. It is the data that retains its value over time and can be used to create insight, access, or advantage in the future.

Why Current Cryptography Will Not Hold

Foundations, Assumptions, and System-Level Implications

Modern cryptographic systems are often described as secure, and in the context of today's computing capabilities, that description is accurate. However, it is important to understand what "secure" means in this context. Cryptography does not provide absolute protection. It provides protection based on the assumption that certain problems are computationally infeasible to solve within a practical timeframe.

This distinction matters. Security is not determined by whether a problem can be solved, but by how long it would take and what resources would be required. If solving the underlying problem is impractical, the system is considered secure.



Most modern systems rely on two complementary forms of cryptography. Symmetric cryptography uses a single key for both encryption and decryption and is optimized for speed and efficiency. It is commonly used to protect data at rest and to encrypt large volumes of data in transit. Asymmetric cryptography, by contrast, uses a pair of keys, one public and one private. This approach enables secure key exchange, digital signatures, and identity verification.

While both forms of cryptography are important, they serve different roles within a system. Symmetric cryptography provides performance. Asymmetric cryptography provides trust. It is asymmetric cryptography that enables systems to establish secure communication without pre-shared secrets and to verify the identity of other systems.

This distinction is critical because the long-term risk associated with emerging computational models affects these two approaches differently.

The security of widely used asymmetric systems, such as RSA and elliptic curve cryptography, is based on specific mathematical assumptions. RSA relies on the difficulty of factoring large numbers into their prime components. Elliptic curve cryptography relies on the difficulty of solving discrete logarithm problems. For classical computers, both problems become exponentially more difficult as the size of the inputs increases. This is what makes them practical for securing modern systems.

However, these assumptions are tied to the capabilities of classical computation. They do not represent inherent limits of mathematics. They represent limits of the tools currently used to solve these problems.

Quantum computing introduces a different model of computation that changes how certain problems can be approached. Instead of evaluating possibilities sequentially, quantum systems can explore multiple states simultaneously and use interference to amplify correct solutions. While this does not make all problems easy, it does affect specific classes of problems that current cryptographic systems depend on.

In particular, algorithms such as Shor's algorithm are expected to significantly reduce the effort required to solve factoring and discrete logarithm problems. If these problems become tractable at scale, the foundation of widely used asymmetric cryptography is weakened.

The implications of this extend beyond individual encryption mechanisms. Asymmetric cryptography is deeply embedded in the systems that establish trust. It is used to secure communication channels, validate identities, and ensure the integrity of data. If the underlying assumptions are no longer valid, the impact is not limited to confidentiality. It affects authentication, authorization, and system trust.

For example, if the integrity of digital signatures cannot be assured, it becomes possible to impersonate systems or users. If certificate validation can be bypassed, the mechanisms used to establish secure communication can be undermined. These are not isolated failures. They represent a breakdown in the systems that allow distributed environments to function securely.



Symmetric cryptography is affected differently. Quantum approaches are expected to reduce the effective strength of symmetric keys but not eliminate their usefulness. In practical terms, this means that larger key sizes can compensate for the impact. Increasing key lengths can restore the desired level of security for symmetric systems.

This difference leads to a common misconception. It is often assumed that increasing key sizes will solve the broader problem. While this approach can be effective for symmetric cryptography, it does not address the underlying vulnerability in asymmetric systems. The issue is not the size of the key, but the mathematical structure on which the system is built.

Because asymmetric cryptography is so deeply integrated into modern systems, replacing it is not a simple task. It is not a matter of updating a single component or applying a patch. It requires changes across protocols, identity systems, applications, and infrastructure. Systems that were designed with specific cryptographic assumptions must be re-evaluated and, in some cases, redesigned.

This creates a time-related challenge. Even if new cryptographic approaches are available, transitioning to them across a large and complex environment takes years. Systems have dependencies, integrations, and operational constraints that limit how quickly changes can be made. At the same time, data that is generated and stored today may need to remain secure for decades.

This overlap between long data lifecycles and long transition timelines is where risk accumulates. Organizations are effectively operating within a window in which current systems remain effective, but future requirements are already becoming relevant.

The strategic reality is that current cryptographic systems are not failing today. They are functioning as designed. The issue is that they were not designed for a world in which new computational models can change the difficulty of the problems they rely on.

Understanding this distinction is essential. It shifts the conversation from reacting to failure to preparing for change. It reinforces the need to evaluate systems not only based on their current effectiveness, but on their ability to remain effective over time.

Where Organizations Are Most Exposed

Real Environments, System Patterns, and Concentrated Risk

When organizations think about cryptographic risk, they often look for a specific system, application, or dataset where exposure might exist. This framing suggests that risk is localized and identifiable. In practice, exposure is rarely concentrated in a single place. It is distributed across systems, embedded in architecture, and accumulated over time.

Modern environments are built from layers of infrastructure, applications, and integrations that evolve independently. As a result, data is not stored or processed in one location. It exists



simultaneously across production systems, backups, analytics platforms, and partner environments. Each of these locations introduces a different form of exposure, and taken together, they create a network of conditions in which risk persists.

One of the most significant concentrations of risk exists in long-life data stores, particularly in backup and archival systems. These environments are designed to preserve data, often for years, and they frequently contain complete copies of operational datasets. Because they are not part of day-to-day operations, they are typically subject to less monitoring and fewer access controls than production systems.

In many organizations, backups are assumed to be secure because they are encrypted. However, this assumption often overlooks how those systems are implemented. Backup environments may rely on long-lived keys, static encryption configurations, or shared access models that are not regularly reviewed. They also tend to aggregate large volumes of sensitive data into a single location. This combination of high-value content, long retention periods, and limited visibility makes them particularly attractive in a long-term collection model.

Legacy systems represent another major area of exposure. These systems often remain in operation because they support critical functions, even as the surrounding environment evolves. Many were designed with older cryptographic libraries, hardcoded assumptions, or limited support for updates. As a result, they are difficult to modify and may not be capable of adopting new cryptographic approaches without significant redesign.

In practice, legacy systems create a form of persistent risk. They continue to operate as part of the environment, but they do not evolve at the same pace as modern systems. Security improvements are often deferred because of the perceived risk of disruption. Over time, this leads to a gap between how systems are protected and how they need to be protected.

Exposure is also shaped by the increasing reliance on third-party systems and supply chain integration. Organizations routinely share data with cloud providers, SaaS platforms, and external partners. These relationships extend the environment beyond direct organizational control and introduce dependencies on external security practices.

In this context, exposure is often indirect. It may not originate within an organization's own infrastructure, but within the systems of a vendor or partner. Even when strong controls exist internally, inconsistencies across the broader ecosystem can create gaps. Organizations may assume that encryption and key management are handled appropriately by their providers, but in many cases, they have limited visibility into how those controls are implemented.

Another area that is frequently overlooked is machine-to-machine communication. Modern systems rely heavily on internal APIs, service-to-service interactions, and automated workflows. These communication paths handle large volumes of data and are often treated as trusted channels within the environment.

Because they are internal, they may not receive the same level of scrutiny as external interfaces. Encryption may be applied inconsistently, key management practices may vary, and monitoring



may be limited. Over time, these internal pathways become some of the most active and least visible areas of data movement.

Identity systems introduce a different form of exposure with broader implications. Cryptographic mechanisms are used to establish and maintain trust across systems, including certificates, authentication tokens, and key-based identity models. If these systems are compromised or weakened, the impact extends beyond data confidentiality.

A failure in identity systems can enable impersonation, unauthorized access, and the breakdown of trust relationships across environments. In this sense, identity-related cryptography is not just protecting data. It is enabling the system to function securely. Weakness in this area creates systemic risk rather than isolated exposure.

Endpoints and edge systems add another layer of complexity. User devices, IoT systems, and operational technology often operate outside centralized control. They may have limited processing capabilities, inconsistent update cycles, and weaker protection for keys and data. These systems frequently handle sensitive information at the point where it is created or consumed.

Encryption provides limited protection in these scenarios if the endpoint itself is compromised. Data can be accessed before it is encrypted or after it is decrypted, and keys may be exposed through memory or local storage. This reinforces the idea that encryption is only one component of a broader system of controls.

The shift toward hybrid and multi-cloud environments further expands the exposure surface. Data moves between on-premises systems, multiple cloud providers, and partner environments. Each transition introduces a new trust boundary, often with different assumptions about security and control.

In practice, this creates inconsistency. Encryption models may vary across environments, key management practices may differ, and visibility may be fragmented. Even if each individual environment is reasonably secure, the transitions between them can introduce gaps.

A common thread across all these areas is the assumption that encrypted data is inherently safe. While encryption is an essential control, it does not eliminate risk, particularly when data is stored for long periods or distributed across multiple systems. Encrypted data can still be collected, retained, and later decrypted if underlying assumptions change.

Another important factor is the presence of unknown or poorly understood systems. Over time, organizations accumulate applications, integrations, and data stores that are not fully documented or centrally managed. These may include internally developed tools, legacy integrations, or shadow systems that operate outside standard governance.

These unknowns represent some of the most significant sources of risk. They are difficult to assess, difficult to monitor, and often overlooked in security planning. As a result, exposure may exist in areas that are not visible to the teams responsible for managing risk.



Taken together, these patterns highlight an important reality. Exposure is not created by a single failure or weakness. It emerges from the way systems are designed, integrated, and operated over time. It is reinforced by complexity, distribution, and limited visibility.

Understanding where exposure exists requires moving beyond individual systems and looking at how data flows across the environment. It requires recognizing that risk is not confined to obvious targets but is often concentrated in areas that are less visible, less monitored, and more persistent.

Why Visibility and Control Are So Difficult

Structural, Technical, and Organizational Barriers

At this point, the challenge is no longer understanding the risk. It is understanding why, despite widespread use of encryption and significant investment in security, most organizations still struggle to clearly answer a basic question: where are we exposed?

The difficulty is often attributed to gaps in tooling or process. While these factors play a role, they are not the primary cause. The challenge is structural. It is rooted in how modern systems are designed, how organizations operate, and how responsibilities are distributed.

Encryption is not implemented in a single place. It exists across applications, infrastructure, databases, devices, and identity systems. Each layer introduces its own methods, dependencies, and assumptions. Over time, these layers evolve independently, often under different teams and priorities. What emerges is not a single, unified system, but an ecosystem of interconnected components.

In this environment, there is rarely a central inventory of cryptographic usage. Different systems may use different algorithms, different libraries, and different key management approaches. Some systems are modern and well maintained. Others are legacy systems that have remained unchanged for years. Still others may be experimental or internally developed tools that were never fully integrated into standard governance processes.

Consider a typical enterprise environment. External communication may be secured using current versions of TLS, while internal systems rely on older encryption methods that have not been updated. Backup systems may use a separate encryption model entirely, often with longer-lived keys. Each of these implementations may function correctly in isolation, but no single team has a complete view across all of them.

This fragmentation is reinforced by how responsibilities are distributed within organizations. Security teams define policies and standards, but they do not implement every system. Application teams build and deploy services, often making decisions about encryption within the context of performance and functionality. Infrastructure teams manage platforms and environments, focusing on availability and scalability. Data teams prioritize access and usability.



Each group operates with a partial view of the system and a specific set of objectives. Coordination across these groups is possible, but it is not always consistent. As a result, encryption practices can vary across the organization, even when high-level policies are in place.

Tooling adds another layer of complexity. Organizations invest in monitoring and security tools with the expectation that they will provide visibility into how systems operate. In practice, most tools provide insight into specific aspects of the environment rather than a comprehensive view.

Some tools focus on network traffic, others on certificates, and others on data storage. Few provide a unified perspective that connects cryptographic usage, key management, and data movement across systems. This leads to a situation where multiple tools are in use, each offering a partial view, but none providing a complete picture.

In addition to managed systems, organizations must contend with what are often referred to as shadow systems. These include internally developed applications, experimental environments, and integrations that operate outside formal governance structures. Developers may implement encryption using standard libraries, but without centralized oversight, key management and configuration practices may vary.

From the perspective of the development team, encryption may be implemented correctly. From the perspective of the organization, it may be unmanaged. This distinction is important because it highlights that correct implementation at a local level does not guarantee effective control at a system level.

Scale further complicates visibility. Large organizations may operate thousands of applications, process millions of data transactions, and maintain environments across multiple cloud providers and on-premises systems. These systems are not static. They change continuously as new features are deployed, integrations are added, and infrastructure is updated.

Even when visibility is achieved at a specific point in time, it can quickly become outdated. Maintaining an accurate understanding of the environment requires continuous effort, not a one-time assessment.

Despite these challenges, many organizations maintain a degree of confidence based on the presence of encryption. The assumption is that if data is encrypted, it is protected. While this is true in a narrow sense, it does not account for how encryption is implemented, how keys are managed, or how data moves through systems.

Encryption usage can create a sense of assurance without providing full control. Data may be encrypted in storage and in transit, but still exposed through logs, intermediate systems, or poorly managed keys. This gap between usage and understanding is one of the central challenges organizations face.

Another significant barrier is the lack of crypto-agility in existing systems. Many applications and platforms were designed with specific cryptographic methods embedded directly into their



logic. These assumptions were reasonable at the time of development, but they limit the ability to adapt.

Changing cryptographic methods in these systems often requires code changes, system updates, and revalidation of functionality. This creates friction. Even when the need for change is recognized, the effort required to implement it can lead to delays or incremental approaches.

Organizational constraints further reinforce these challenges. Teams operate within limits on time, budget, and resources. Security improvements must compete with operational priorities, feature development, and system stability. In many cases, the decision is not whether to address risk, but when and how to do so without disrupting critical operations.

This leads to a pattern where risk is acknowledged but deferred. Improvements are made incrementally, and systemic issues persist. This is not a failure of awareness. It reflects how complex systems and organizations operate.

Taken together, these factors explain why visibility and control remain difficult. The challenge is not caused by a single weakness that can be corrected. It is the result of distributed systems, fragmented ownership, partial tooling, and operational constraints.

The most important implication is that strong encryption alone does not eliminate risk. When encryption exists within a system that lacks visibility and control, it becomes difficult to assess exposure or respond effectively to change.

Understanding this dynamic is a critical step. It shifts the focus from isolated technical fixes to broader questions about architecture, governance, and coordination. It also provides a more realistic foundation for addressing the challenges that follow.

Post-Quantum Cryptography and the Transition Reality

New Foundations, Practical Tradeoffs, and Implementation Challenges

As the limitations of current cryptographic approaches become better understood, attention naturally turns to what comes next. Post-quantum cryptography has emerged as the primary path forward, offering new methods designed to withstand the types of attacks that future computational models may enable.

It is important, however, to approach this transition with clarity. Post-quantum cryptography is not a simple replacement for existing systems, nor is it a single solution that can be applied uniformly across all environments. It represents a shift in the mathematical foundations of cryptography, and with that shift comes a new set of considerations.

Traditional asymmetric cryptography relies on problems such as factoring large numbers or solving discrete logarithms. These problems are difficult for classical computers, which is why they have been effective for securing modern systems. Post-quantum approaches are built on



different types of problems that are not known to be vulnerable to the same computational techniques.

Among the most widely discussed approaches are lattice-based methods, which rely on geometric problems in high-dimensional space, as well as hash-based and code-based techniques. Each of these approaches has different characteristics, strengths, and limitations. There is no single method that addresses all use cases equally well.

This diversity is important because it means that post-quantum cryptography is not a single upgrade path. Organizations will need to evaluate which approaches are appropriate for different parts of their environment, considering performance, compatibility, and operational constraints.

Efforts to standardize these new approaches are well underway, with leadership from the National Institute of Standards and Technology. Initial post-quantum algorithms have been selected, and guidance has been provided to support early adoption. This represents a significant step forward, as it provides a common reference point for vendors and organizations.

At the same time, standardization does not eliminate complexity. The existence of approved algorithms does not mean that all systems can immediately adopt them. Many existing protocols, applications, and infrastructure components were designed with specific cryptographic assumptions that do not align directly with post-quantum approaches.

One of the most visible differences is the size of keys and cryptographic artifacts. Post-quantum methods often require significantly larger keys than those used in current systems. This has practical implications for storage, bandwidth, and system performance. In environments where resources are constrained, such as embedded systems or edge devices, these differences can be difficult to accommodate.

Performance characteristics also change. Some post-quantum operations require more computation or introduce additional latency, particularly during key exchange or signature verification. While these impacts can often be managed, they must be understood and tested within the context of real systems.

Perhaps more important than these individual differences is the impact on how systems are designed and integrated. Cryptographic methods are not isolated components. They are embedded within communication protocols, identity systems, and application logic. Changing the underlying cryptography often requires changes to these surrounding systems.

For example, secure communication protocols such as TLS were originally designed around RSA and elliptic curve cryptography. Incorporating post-quantum methods into these protocols requires updates to how keys are exchanged, how certificates are structured, and how compatibility is maintained across different systems.

Because of these challenges, many organizations are exploring hybrid approaches as a transitional strategy. In a hybrid model, classical cryptographic methods are used alongside post-



quantum methods. The goal is to ensure that data remains protected unless both approaches are compromised.

This approach provides a level of continuity while new methods are adopted. It allows organizations to begin integrating post-quantum capabilities without fully abandoning existing systems. However, it also introduces additional complexity, as systems must support multiple cryptographic methods simultaneously.

The impact of this transition is not limited to communication protocols. Identity and trust systems are also affected. Certificates, public key infrastructure, and authentication mechanisms all rely on cryptographic assumptions that must be revisited. Updating these systems requires careful coordination to ensure that trust relationships are maintained throughout the transition.

Embedded and long-lifecycle systems present a particularly difficult challenge. Devices that are deployed in the field, such as industrial control systems or IoT devices, may have limited processing capabilities and long operational lifespans. In many cases, they cannot be easily updated to support new cryptographic methods.

This creates a situation where some systems may remain dependent on older cryptographic approaches for extended periods of time. Managing this risk requires a combination of compensating controls, segmentation, and careful planning.

Key management also becomes more complex in a post-quantum environment. Larger keys, new algorithms, and hybrid models all place additional demands on how keys are generated, stored, and rotated. Organizations must ensure that their key management practices evolve alongside their cryptographic methods.

The scale of the transition adds another layer of difficulty. Large organizations operate thousands of systems with interdependencies that are not always fully understood. Updating cryptographic methods in one system may require changes in others, as well as coordination with external vendors and partners.

This interconnectedness means that migration cannot be approached as a single event. It must be planned and executed over time, with careful consideration of dependencies and sequencing. In many cases, progress will be incremental rather than linear.

There are also common misconceptions that can slow progress. One is the belief that it is best to wait until post-quantum cryptography is fully mature before taking action. While it is important to avoid premature adoption of unstable technologies, waiting for complete certainty can compress timelines and increase risk.

Another misconception is that vendors will handle the transition entirely. While vendors play a critical role, organizations are responsible for integrating, configuring, and operating their systems. Readiness depends on both the capabilities provided by vendors and the actions taken internally.



The most important shift introduced by post-quantum cryptography is not a specific algorithm, but a change in how organizations approach cryptographic systems. Rather than treating cryptography as a fixed component, it must be viewed as an evolving capability that can adapt over time.

This perspective connects directly to the concept of crypto-agility, which enables systems to change without requiring complete redesign. Without this capability, even well-chosen cryptographic methods can become difficult to maintain as requirements evolve.

The transition to post-quantum cryptography is therefore both necessary and complex. It requires technical understanding, operational planning, and organizational alignment. It is not a one-time upgrade, but a multi-year effort that must be approached deliberately.

Crypto-Agility and Transition Strategy

From Awareness to Execution in Real Environments

Understanding the limitations of current cryptographic systems and the emerging role of post-quantum approaches is only the beginning. The more difficult challenge is determining how to transition in a way that is both practical and sustainable. For most organizations, this is not a matter of selecting a new algorithm. It is a matter of enabling change across systems that were not designed to change easily.

This is where the concept of crypto-agility becomes central. At a high level, crypto-agility refers to the ability to update cryptographic methods without requiring fundamental redesign of systems. In practice, it represents a broader capability. It requires that cryptography be treated as a component that can evolve, rather than a fixed assumption embedded in code and infrastructure.

Most existing systems do not meet this standard. Cryptographic methods are often tightly coupled to application logic, implemented directly through specific libraries, and integrated into workflows that assume a particular structure. These decisions were reasonable at the time they were made, but they create friction when change is required.

For example, an application may rely on a specific encryption library, assume a particular key format, and interact with other systems that expect the same structure. Changing the underlying cryptography in this context is not isolated. It requires updates to the application, coordination with dependent systems, and validation to ensure that functionality is preserved. This complexity is one of the primary reasons transitions are slow.

Moving toward crypto-agility requires a shift in how systems are designed. Rather than embedding cryptographic decisions directly within applications, organizations must introduce abstraction layers that separate cryptographic functions from business logic. This allows cryptographic methods to be updated independently, reducing the impact on the surrounding system.



In practical terms, this often involves introducing a service layer that handles encryption, decryption, and key management. Applications interact with this layer through defined interfaces, rather than implementing cryptography directly. When changes are required, they can be made within the service layer, minimizing disruption to the broader system.

Key management plays a central role in this model. It becomes the control plane through which cryptographic operations are coordinated. Effective key management requires visibility into where keys are used, control over how they are accessed, and processes for rotation and revocation. Without this foundation, even well-designed cryptographic systems become difficult to manage.

Transitioning to this model begins with understanding the current environment. Organizations need to identify where cryptography is used, how keys are managed, and which systems depend on specific algorithms. This is often referred to as inventory and mapping, but in practice it is more complex than the term suggests.

Documentation may be incomplete, dependencies may be hidden, and legacy systems may not be fully understood. Attempting to create a perfect inventory before taking action can delay progress. A more effective approach is to begin with high-risk systems and expand visibility over time. This allows organizations to make progress while continuing to refine their understanding.

Prioritization is equally important. A full-scale, simultaneous transition across all systems is not feasible. Organizations must determine where to focus first based on factors such as data longevity, system criticality, and external exposure. Identity systems, external communication channels, and high-value data stores are often logical starting points because of their broad impact.

As organizations begin to plan for transition, they must also account for dependencies on vendors and external platforms. Many systems rely on third-party components that define how cryptography is implemented. Cloud providers, SaaS platforms, and hardware vendors each have their own timelines and capabilities.

This creates a coordination challenge. Even if an organization is ready to adopt new cryptographic methods, it may be constrained by the readiness of its ecosystem. Engaging with vendors early, understanding their roadmaps, and aligning expectations becomes a critical part of the process.

Testing and validation represent another significant effort. Cryptographic changes affect system behavior in ways that are not always immediately visible. Performance characteristics may change, compatibility issues may arise, and interactions between systems may need to be re-evaluated. Ensuring that systems continue to function as expected requires careful testing across different environments.

Because of these factors, transition is best approached as a phased process rather than a single event. Organizations typically move through stages that include awareness, initial assessment,



planning, pilot implementations, and gradual expansion. Progress is not always linear. Some areas may advance quickly, while others require more time due to complexity or dependency.

Throughout this process, alignment across teams is essential. Security, infrastructure, application development, and data teams all play a role in implementing and operating cryptographic systems. Each group brings a different perspective and set of priorities. Coordinating these efforts requires clear communication and shared understanding of objectives.

Organizational constraints must also be acknowledged. Teams operate within limits on time, budget, and resources. Changes must be balanced against the need to maintain system availability and support ongoing operations. This often leads to incremental progress rather than large-scale transformation.

Delaying action, however, has consequences. The longer organizations wait to begin the transition, the more systems become dependent on existing approaches, and the more data accumulates under assumptions that may not hold in the future. This increases both the scale of the eventual transition and the potential impact of exposure.

The goal of this process is not to reach a fixed end state. It is to develop the capability to adapt. Cryptographic requirements will continue to evolve, and systems must be able to evolve with them. Crypto-agility provides the foundation for this adaptability, enabling organizations to respond to change without starting from scratch.

In this sense, the transition to post-quantum cryptography is part of a broader shift. It requires organizations to move from static security models to dynamic ones, where change is expected and managed as part of normal operations.

What Organizations Should Do Now

From Understanding to Action

The challenge outlined throughout this paper is not theoretical, and it is not confined to a future point in time. The conditions that create long-term cryptographic risk are already present in today's systems, data, and decisions. As a result, the appropriate response is not to wait for a specific technological milestone, but to begin addressing the factors that are already within organizational control.

One of the most important shifts is how this problem is framed. It is not necessary for every individual within an organization to understand the details of quantum computing or advanced cryptographic methods. What is required is a shared understanding that data has a lifecycle, that risk can accumulate over time, and that decisions made today influence future exposure.

This understanding must extend beyond leadership to include the teams that design, build, and operate systems. Engineers, architects, security professionals, and data teams all play a role in



how cryptography is implemented and managed. When awareness is distributed across these groups, it becomes possible to align decisions and reduce fragmentation.

A practical starting point is to identify what data truly matters over time. Not all information requires the same level of protection, and attempting to treat it all equally can dilute focus. Organizations should begin by asking which data must remain secure for extended periods, where that data resides, and how it moves across systems. This exercise does not require perfect accuracy. It requires enough clarity to begin prioritizing.

Building visibility is the next step. Many organizations operate with limited understanding of where cryptography is used, how keys are managed, and how data flows through their environments. Improving this visibility does not require a complete inventory on day one. It can begin with high-value systems and expand over time.

As visibility improves, organizations can move into structured assessment. This involves evaluating where long-life data exists, how it is protected, and where dependencies on current cryptographic methods may create future risk. It also includes examining key management practices, which often determine the effectiveness of encryption in practice.

Planning for transition should follow naturally from this assessment. This does not mean committing to immediate, large-scale changes. It means establishing direction. Organizations should understand which systems are most critical, which dependencies may influence timelines, and how emerging standards can be incorporated over time.

Engagement with external partners is an essential part of this process. Many systems rely on vendors, cloud providers, and third-party platforms that define how cryptography is implemented. Understanding their roadmaps, capabilities, and constraints helps organizations align their own plans and avoid surprises.

It is also important to recognize that progress will not occur in a single step. Transition will be phased, and different parts of the organization will move at different speeds. Some systems can be updated relatively quickly, while others may require longer timelines due to complexity or operational constraints. The goal is not uniformity. It is steady, prioritized progress in the areas that matter most.

Throughout this process, organizations should avoid two common pitfalls. The first is waiting for complete certainty before taking action. While it is important to make informed decisions, waiting for all variables to be resolved can delay progress and increase risk. The second is attempting to solve the problem through a single initiative or transformation. The scale and complexity of the challenge require a sustained effort rather than a one-time project.

Progress should be measured in terms of direction rather than completion. Indicators such as improved visibility, clearer prioritization, initial assessments, and pilot implementations all represent meaningful movement. Over time, these efforts build the foundation for broader adoption and more significant change.



At a strategic level, this challenge can be viewed through three lenses. It is a risk management issue, as it involves protecting data and systems over extended periods. It is an operational issue, as it requires systems and processes that can adapt to change. And it is a strategic issue, as organizations that move early are better positioned to manage both risk and opportunity.

The most important conclusion is that the response does not require immediate, large-scale disruption. It requires starting. Organizations that begin to build awareness, improve visibility, and plan for change will be in a stronger position than those that defer action in the hope of future clarity.

The conditions that create future exposure are already in place. Addressing them does not require certainty about the future. It requires a clear understanding of the present and a willingness to act on it.



SecureFi Institute Research Deep Dive Paper 011

Harvest Now, Decrypt Later

Understanding the Emerging Quantum Risk and What Organizations Must Do Now

SecureFi Institute

SecureFi Institute focuses on the convergence of cyber, artificial intelligence, high-performance computing, and quantum technologies and their impact on secure infrastructure.

Through research, executive briefings, and training, the Institute helps organizations move from awareness to readiness.

For additional insights and research, visit securefi.com

Research. Awareness. Preparedness.

Figures and Analytical Models

All figures, diagrams, and analytical models presented in this research brief were developed by SecureFi Institute as part of its research on emerging computing architectures and cybersecurity implications.

Research Disclaimer

This research brief is provided for informational and educational purposes and reflects analysis from SecureFi Institute on emerging computing technologies and cybersecurity trends. The views expressed are intended to support awareness and discussion of technology and infrastructure challenges and do not represent official policy positions.

@ 2026 SecureFi Institute. All rights reserved.

