International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2025: 7.089

Protecting Cardholder Data in Modern Payment Systems: A Threat-Model-Based Comparative Analysis of Tokenization and Encryption for PCI DSS Scope and Breach-Risk Reduction

Chandrashekhar Kharche

Independent Researcher, Payments Technology, Leander, Texas, USA
Email: *cpskharche[at]gmail.com*

Abstract: ***Payment card data is among the most valuable targets in cybercrime, and the systems that process it are heavily regulated under the Payment Card Industry Data Security Standard (PCI DSS). Two design goals dominate its protection: reducing the scope of systems subject to the standard, and reducing the usefulness of card data if it is stolen. A family of techniques (point-to-point encryption, vaulted and vaultless tokenization, format-preserving encryption, and network or EMV payment tokenization) pursues these goals differently. This article compares them under an explicit method: we define a threat model and deployment assumptions, assess each technique separately against security, compliance, and integration criteria, and validate the comparison against three representative payment architectures rather than relying on qualitative positioning alone. We correct a common overstatement: tokenization does not automatically remove systems from PCI DSS scope. The tokenization system itself (vault, detokenization, and keys) remains in the cardholder data environment; a token-only system may be considered out of scope only if the PAN is not computationally recoverable from tokens, the environment is segmented from detokenization and key material, no other cardholder-data channels exist, and that segmentation is validated, with high-value tokens potentially remaining in scope. On the stakes, we report observed survey figures by edition: IBM's Cost of a Data Breach Report put the financial-sector average at about USD 5.56 million in its 2025 edition and USD 6.29 million in its 2026 edition, the latter roughly 26% above that year's USD 4.99 million cross-industry average; we do not attribute a quantified reduction to any single technique. We close with a decision framework and qualify the quantum-era discussion with explicit assumptions.***

Keywords: tokenization, PCI DSS scope, segmentation, payment data security, cardholder data, payments

1. Introduction

A payment card number is a bearer credential: whoever holds the primary account number (PAN) in the clear can often use it. That fact makes stored and transmitted card data a perennial target and has made its protection a regulated obligation. Under PCI DSS, any system that stores, processes, or transmits cardholder data falls within the cardholder data environment (CDE) and inherits a demanding set of controls, assessment, and audit obligations [1]. The economic context is significant: IBM's Cost of a Data Breach Report places the financial sector among the costliest, with an observed average near USD 5.56 million in its 2025 edition and USD 6.29 million in its 2026 edition [6]. These are survey averages for the sector, not figures attributable to the presence or absence of any particular protection technique, and we treat them as context rather than as evidence of a technique's effect.

Two broad strategies reduce the risk and the regulatory burden of handling card data. The first protects the data cryptographically so that intercepted values are not directly usable. The second replaces the PAN with a surrogate so that most systems no longer hold the real number. Both are widely described, but the compliance consequence of the second is frequently overstated, as if replacing the PAN with a token automatically removed a system from PCI DSS scope. It does not. This article compares the principal techniques under an explicit, reproducible method and states the scope consequences precisely. Its contributions are four: a defined threat model and comparison methodology (Section 3); a structured characterization of the techniques (Section 4) with a separate assessment of security, compliance, and integration (Section 5); a precise account of the PCI DSS scope conditions, with a reference architecture (Section 6); and a decision framework validated against representative architectures (Sections 5 and 7).

2. Background: Cardholder Data and PCI DSS Scope

PCI DSS organizes its requirements around wherever the PAN lives. Scope, the set of systems, networks, and people subject to the standard, is determined by data flow: any component that stores, processes, or transmits a live PAN, and any component connected to or that could affect the security of such a component, is in scope [1][11]. Scope is reduced primarily by two means: removing the PAN from a system so it no longer handles cardholder data, and network segmentation that isolates in-scope systems from out-of-scope ones. Segmentation is not a control that PCI DSS mandates, but where it is used to reduce scope it must be implemented and, critically, verified; an unsegmented or inadequately segmented environment leaves connected systems in scope [11]. This distinction, between removing data and merely protecting it, and between claiming segmentation and validating it, is central to the analysis that follows.

3. Comparison Methodology

To avoid the limitations of purely qualitative positioning, we compare the techniques under an explicit method with three elements: a threat model, deployment assumptions, and separately assessed criteria, and we then validate the comparison against representative architectures.

3.1 Threat model

The protected asset is the PAN. The modeled adversary gains read access to data at rest or in transit within the merchant or processor environment: for example, by exfiltrating a database, capturing data on the network, or compromising a non-CDE host. The question each technique must answer is whether the value the adversary obtains can be turned back into a usable PAN. We treat two threats as *out of model* and note them as residual risks: an insider with authorized detokenization or key-access rights, and physical tampering of a capture device. These are addressed by access control (PCI DSS Requirements 7 and 8) and device standards rather than by the choice among the data-protection techniques compared here [1].

3.2 Deployment assumptions

We assume a regulated merchant, processor, or service provider that already operates production systems and seeks to protect card data while reducing assessable scope. We do not assume a greenfield build; integration cost into existing systems is therefore a first-class criterion. We assume competent key and vault management, and we assess each technique at the level of a validated, correctly configured implementation, noting where real-world misconfiguration is the dominant risk.

3.3 Evaluation criteria (assessed separately)

We assess three criteria independently rather than collapsing them into a single ranking. **(i) Security:** if protected data is exfiltrated under the threat model, what is required to recover the PAN (computationally infeasible, a key, a vault lookup, or merely that the value is domain-bound)? **(ii) Compliance (PCI DSS scope):** can the technique reduce assessable scope, under what conditions, and which components remain in scope regardless? **(iii) Integration:** what change to existing systems, formats, terminals, or network onboarding does adoption require? Keeping these separate is deliberate: a technique can be strong on one and weak on another, and conflating them is how overstated scope claims arise.

3.4 Representative architectures for validation

We validate the comparison against three representative architectures rather than asserting a context-free ranking: **A1**, a card-present retail merchant with terminal-based acceptance; **A2**, an e-commerce card-not-present merchant with web checkout and card-on-file; and **A3**, an enterprise or service provider operating a central token vault. Section 5 maps each technique to these architectures. We note the method's limitation plainly: this is a structured, criteria-based analysis grounded in the standards and documented guidance, not an empirical benchmark of deployed systems, and the representative architectures are illustrative archetypes rather than audited case studies.

4. Techniques for Protecting Cardholder Data

4.1 Point-to-point encryption (P2PE)

Point-to-point encryption encrypts card data at the point of capture, typically inside a tamper-resistant payment terminal, so that it travels encrypted through the merchant's systems and is decrypted only within a secure environment at the payment processor [10]. Because the merchant's own systems never hold clear data, a PCI-validated P2PE solution can substantially reduce the merchant's assessable scope, which is why it carries a dedicated PCI SSC standard and validation program. Its applicability is architectural: it protects data in transit from a controlled capture device and is best suited to card-present, terminal-based flows (A1).

4.2 Vaulted tokenization

Tokenization replaces the PAN with a surrogate, a token, generated so that the PAN is not computationally recoverable from the token alone [2]. In the vaulted model, a secure token vault stores the token-to-PAN mapping; downstream systems hold only tokens and call the vault when the real number is genuinely needed. This can keep the PAN out of many systems, but the vault, the detokenization function, and the PAN storage are themselves part of the CDE and remain in scope; the scope benefit accrues to the segmented systems that hold tokens only, and only if the conditions of Section 6 are met [2].

4.3 Vaultless tokenization

Vaultless tokenization avoids a central mapping database, deriving tokens cryptographically so that an authorized party can recover the PAN without a stored lookup table [5]. This removes the vault as a single high-value target and eases scaling, while preserving the goal of keeping PANs out of downstream systems. Its security rests on the strength and management of the underlying keys, so the key-management function occupies the sensitive, in-scope role that the vault holds in the vaulted model.

4.4 Format-preserving encryption (FPE)

Format-preserving encryption encrypts a PAN into a ciphertext of the same length and format, so that it fits existing database fields and message formats without schema or application changes; the NIST-standardized modes FF1 and FF3 define the construction [4]. FPE therefore eases retrofitting protection into legacy systems. Because the protected value is reversible with the key and remains a cipher of the PAN, it is cryptographically protected cardholder data rather than a non-recoverable surrogate; whether a system holding FPE-protected PANs can be treated as out of scope depends on access to the key and on the assessor's judgment, and is generally less clear-cut than for true tokenization [4].

4.5 Network and EMV payment tokenization

Network, or EMV, payment tokenization replaces the PAN with a token issued within the payment network and bound to a specific merchant, device, or channel, usable only within those constraints [3]. It underlies card-on-file and mobile-wallet credentials: even if such a token leaks, its domain restrictions limit misuse. It shifts much of the token lifecycle to the network and issuer, simplifying the merchant's role, and applies primarily to the transactional use of credentials rather than to arbitrary stored data; a merchant that stores only network tokens may reduce scope, again subject to the conditions of Section 6.

5. Comparative Analysis

Figure 1 positions the techniques on two axes that dominate selection: the strength of protection and potential scope reduction they offer, and the simplicity of adopting them in an existing environment. The positioning is a summary of the separate criteria assessed below, not a substitute for them; Table 1 records the three criteria independently, and Table 2 validates the comparison against the representative architectures of Section 3.4.

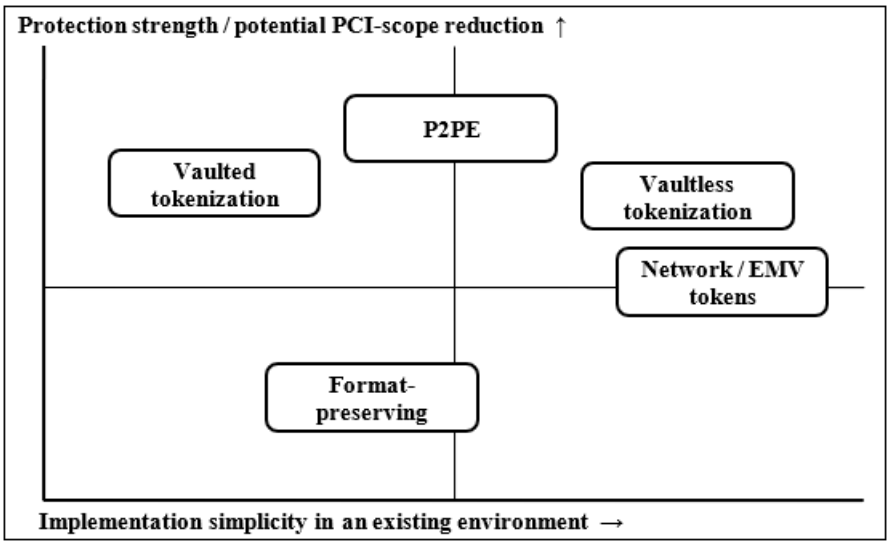


Figure 1: Indicative positioning of cardholder-data protection techniques by protection strength and potential PCI-scope reduction (vertical) against **implementation** simplicity in an existing environment (horizontal). Positions summarize Table 1 and depend on deployment; they are not a context-free ranking.

Table 1: Techniques assessed separately on security, compliance (scope), and integration.

Technique	Security: recovering the PAN requires	Compliance (PCI DSS scope)	Integration change	Reversible to PAN
P2PE	Decryption key, held only at processor	Validated P2PE can reduce merchant scope; HSM/decryption environment in scope	Terminal-centric	At processor
Vaulted tokenization	Access to the token vault	Token-only systems may be out of scope if conditions met; vault/detok/PAN store in scope	Moderate	Via vault
Vaultless tokenization	Access to keys (no vault)	As above; key management in scope	Moderate	Via keys
Format-preserving encryption	The encryption key	Protected PAN, not a surrogate; scope treatment depends on key access and assessor judgment	Low (same format)	Yes, with key
Network / EMV tokens	Not reversible by merchant; domain-bound	Storing only network tokens may reduce scope, subject to conditions	Low for mer-chant	Network-controlled

"May be out of scope" and "may reduce scope" are conditional on the segmentation, access-control, and validation requirements in Section 6; no technique removes systems from scope automatically.

Table 2: Validation against representative architectures: best-fit techniques per context

Architecture	Primary technique	Complementary technique	Scope outcome (if validated)
A1 Card-present retail (terminals)	Validated P2PE at capture	Tokenization for stored references	Merchant systems largely out of scope; decryption env in scope
A2 E-commerce / card-on-file	Network / EMV tokens	FPE for legacy stores	Token-only app tier may be out of scope; checkout/vault in scope
A3 Enterprise token-vault / provider	Vaulted or vaultless tokenization	P2PE for acceptance channels	Segmented token-only consumers out of scope; vault/keys in scope

Outcomes assume the Section 6 conditions are met and independently validated; otherwise connected systems remain in scope

6. PCI DSS Scope: Conditions and Validation

The compliance benefit of tokenization is real but conditional, and stating it precisely is the main correction this paper offers. Per the PCI SSC tokenization guidance, tokenization solutions *do not eliminate* the need to maintain and validate PCI DSS compliance; at best they can simplify validation by reducing the number of in-scope system

components [2]. The tokenization system itself, including detokenization, PAN storage, and key management, is part of the CDE and remains in scope. A token-only system **may be considered outside the CDE and possibly out of scope only if it meets every one of a set of conditions**: the PAN must not be computationally recoverable from a token (even if the token and its host are compromised); the system must be segmented and isolated from any component that can detokenize or that holds the vault or keys; it must have no connection to the tokenization system, vault, or key storage and no credentials that could authenticate to the CDE; it must store, process, and transmit no cardholder data by any other channel; and any legacy cardholder data on it must have been securely removed [2]. Access to detokenization must be governed by strong access control and authentication under PCI DSS Requirements 7 and 8 [1][2].

Two further points temper the claim. First, **validation**: segmentation must be verified, and token usage reviewed, as part of scope validation (for example, during the annual assessment); a claimed reduction that has not been validated does not reduce scope in practice [2][11]. Second, **high-value tokens**: tokens that can themselves initiate transactions may remain in scope even though they cannot reveal a PAN, because they carry exploitable value [2]. Figure 2 shows a reference architecture that makes these boundaries explicit: the capture and protection points and the vault, detokenization, and key functions sit inside the CDE, while a segmented, validated token-only zone is the only part that may fall out of scope, and only when the conditions above hold.

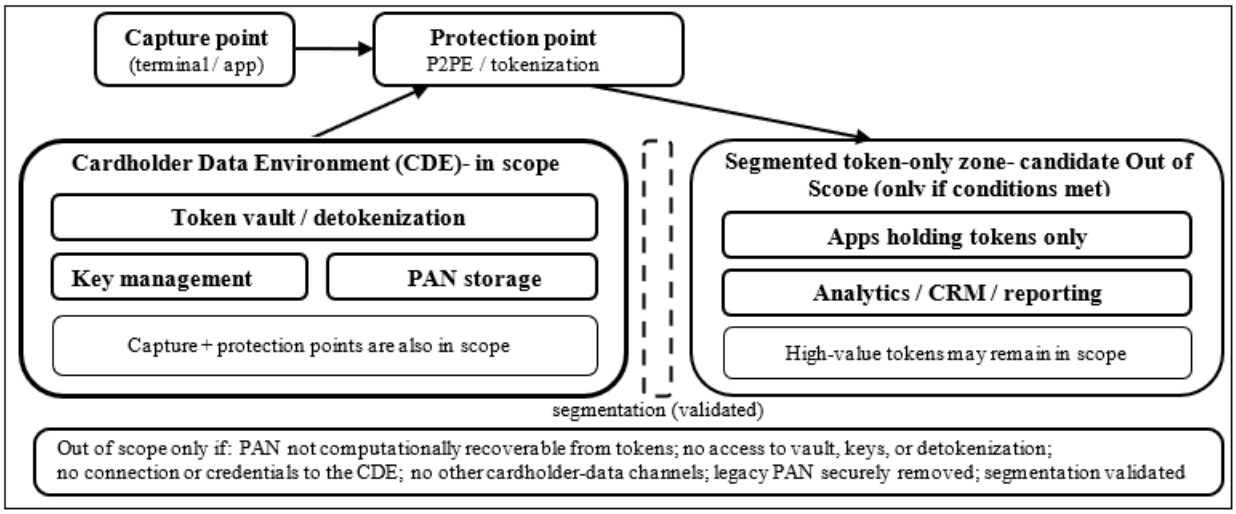


Figure 2: Reference architecture for PCI DSS scope under tokenization. The cardholder data environment (vault, detokenization, key management, and the capture and protection points) remains in scope; only a segmented, validated token-only zone may be out of scope, and only if the listed conditions are met. High-value tokens may remain in scope

7. Breach-Cost Context and the Limits of Attribution

It is tempting to translate the protection techniques directly into breach-cost savings, but the available figures do not support technique-level attribution, and we are careful not to imply it. The headline numbers are observed sector averages from identifiable report editions: IBM's Cost of a Data Breach Report gave a financial-sector average of about USD 5.56 million in its 2025 edition and USD 6.29 million in its 2026 edition, the latter about 26% above that year's USD 4.99 million cross-industry average and the second-highest of any sector [6]. These are *observed* averages across many organizations and incident types, not projections and not controlled comparisons of tokenized versus non-tokenized environments. Industry incident reporting consistently identifies payment-card data as a frequent target [7], which motivates protection, but neither source isolates the marginal effect of a specific technique on breach cost or likelihood.

What can be stated without over-claiming is mechanistic and qualitative. Techniques that render exfiltrated data non-usable (strong tokenization, and encryption whose keys the attacker does not hold) reduce the usefulness of stolen values, and scope reduction, where validly achieved, reduces the number of systems that must be controlled and assessed and therefore the attack surface and compliance effort. Whether these translate into a particular dollar or percentage reduction for a given organization is an empirical question that the cited sources do not answer, and that a credible claim would need to establish with matched data. We therefore present the breach-cost figures as context for why protection matters, not as a measured return on any one technique.

8. A Decision Framework

No single technique is best for every environment; the choice follows from data flow, constraints, and the scope conditions of Section 6. Four heuristics capture the logic, and map onto the architectures of Table 2. **(i) By channel.** Card-present, terminal-based acceptance (A1) favors validated P2PE at capture; card-on-file, wallet, and recurring flows (A2) favor network or EMV tokens, which bind the credential to its use [3][10]. **(ii) By stored-data footprint.** Environments that retain card references across many systems (A3) favor tokenization: vaulted where central control is paramount, vaultless where scale and avoiding a single high-value target matter more [2][5]. **(iii) By legacy constraint.** Where systems expect card-number-shaped data and cannot be re-architected, FPE protects data with minimal change, accepting the less clear-cut scope treatment of a reversible

cipher [4]. **(iv) By objective.** Where the goal is to reduce assessable scope, techniques that keep the PAN out of segmented, validated systems can do so, but only under the Section 6 conditions; encryption that merely protects an in-scope PAN does not by itself remove a system from scope. In practice these techniques are layered: a mature environment often combines P2PE at capture, tokenization for storage, and network tokens for credentials on file.

9. Limitations and Future Directions

This is a structured, criteria-based comparison grounded in the standards and documented guidance, not an empirical benchmark; the positioning in Figure 1 is indicative, and the representative architectures are archetypes rather than audited deployments. Three forward-looking issues deserve explicit, assumption-bound treatment. **Key management** remains the dominant real-world determinant of security for every cryptographic approach [8]. The **migration to network and EMV tokenization** shifts protection and liability toward networks and issuers, with governance consequences that merit study. Finally, the **quantum-era** outlook should be stated conditionally: assuming a cryptographically relevant quantum computer eventually emerges and that current public-key parameters are used, the primary exposure is to the asymmetric cryptography used in key establishment, which NIST's post-quantum standards are intended to address [12]; symmetric and format-preserving ciphers are comparatively less affected but not immune, while the security of true tokenization rests on surrogate randomness and vault or key secrecy rather than on a reversible transform of the PAN, so it is differently exposed. These are assumptions about an uncertain timeline, not predictions, and each reinforces the paper's central point: techniques should be chosen and combined deliberately, on the basis of data flow, threat model, and validated scope, not treated as interchangeable.

10. Conclusion

Protecting cardholder data is both a security and a compliance problem, and the techniques that address it do so in different proportions and under different conditions. Encryption, in its point-to-point and format-preserving forms, makes intercepted data non-usable to an attacker without the key; tokenization replaces the PAN with a surrogate from which it cannot be recovered. The compliance benefit of the latter is real but conditional: tokenization can reduce PCI DSS scope, and a segmented token-only system may fall out of scope, only when the PAN is not recoverable from tokens, the environment is isolated from detokenization and keys, no other cardholder-data channels exist, and that segmentation is validated, while the vault, detokenization, and key functions remain in scope and high-value tokens may too. The breach-cost figures that frame the problem are observed sector averages, not measured returns on any one technique. Chosen under an explicit threat model and validated against real data flows, these techniques meaningfully reduce the usefulness of stolen card data and the burden of protecting it, which is the most that can be claimed, and the most that should be.

References

[1] PCI Security Standards Council, "Payment Card Industry Data Security Standard (PCI DSS), v4.0.1," 2022.
[2] PCI Security Standards Council, "Information Supplement: PCI DSS Tokenization Guidelines," 2011, and "Tokenization Product Security Guidelines," 2015.
[3] EMVCo, "EMV Payment Tokenisation Specification — Technical Framework," v2.x, EMVCo, 2021.
[4] M. Dworkin, "Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption (FF1, FF3)," NIST Special Publication 800-38G, National Institute of Standards and Technology, 2016.
[5] ANSI X9.119-2, "Retail Financial Services-Requirements for Protection of Sensitive Payment Card Data- Part 2: Implementing Post-Authorization Tokenization Systems," American National Standards Institute.
[6] IBM Security, "Cost of a Data Breach Report," 2025 edition (financial-sector average ~USD 5.56M) and 2026 edition (~USD 6.29M; ~26% above the ~USD 4.99M cross-industry average), IBM Corporation.
[7] Verizon, "2024 Data Breach Investigations Report (DBIR)," Verizon Business, 2024.
[8] E. Barker, "Recommendation for Key Management-Part 1: General," NIST Special Publication 800-57 Part 1 Rev. 5, National Institute of Standards and Technology, 2020.
[9] S. Díaz-Santiago, L. M. Rodríguez-Henríquez, and D. Chakraborty, "A Cryptographic Study of Tokenization Systems," International Journal of Information Security, vol. 15, no. 4, pp. 413–432, 2016.
[10] PCI Security Standards Council, "Point-to-Point Encryption (P2PE) Standard," v3.x, 2021.
[11] PCI Security Standards Council, "Information Supplement: Guidance for PCI DSS Scoping and Network Segmentation," 2016.
[12] National Institute of Standards and Technology, "Post-Quantum Cryptography Standards (FIPS 203, 204, 205)," NIST, 2024.