

TOKEN-WISE AI FOR BAU READINESS

Use AI where structured context can reduce reconstruction and expose change - without asking a model to impersonate accountable professional judgment.

By Scott Norton / PCI DSS QSA | Johanson Group, LLP

PRACTITIONER POSITION Spend tokens where readiness gains memory, not where conclusions gain confidence theater.

AI can make compliance work faster while making the compliance story worse. It can summarize a policy, draft a response, reorganize evidence, and produce language that sounds complete long before the underlying population, source, or responsibility has been established. Fluency compresses uncertainty into attractive prose. If the operating model cannot distinguish a retrieved fact from an inferred explanation and a professional conclusion, the tool has added speed to the wrong layer.

A token-wise approach treats model context as a designed resource. Do not pour the entire program into a prompt and hope coherence appears. Supply the smallest trustworthy context that preserves provenance: the relevant asset relationships, current control description, source records, change events, known exceptions, and unresolved questions. Then ask the model to perform bounded work whose output a named person can inspect.

USE AI TO MAINTAIN READINESS MEMORY

Business-as-usual readiness fails when the reasoning behind a decision disappears. A ticket records that something closed. A spreadsheet records that something applies. A screenshot records that something once appeared true. AI can help connect those fragments by extracting named entities, comparing versions, identifying missing owners, proposing relationship edges, and summarizing what changed between two known states.

The valuable output is not an autonomous compliance conclusion. It is a better-prepared human review: a concise delta, a list of unsupported assertions, a proposed evidence lineage, or a set of questions routed to the right owner. The model saves attention by organizing material. The practitioner spends judgment on what the material means.

DESIGN CONTEXT PACKETS, NOT PROMPT AVALANCHES

Token-wise does not simply mean using fewer tokens. It means increasing the ratio of decision-relevant context to noise. A strong packet states the task, authoritative inputs, time boundary, asset or control scope, known limitations, prohibited conclusions, and required citations back to supplied records. This makes the model's operating boundary inspectable before anyone reads its output.

Context packets should be generated from maintained sources where possible. A current inventory slice is better than an old narrative about the inventory. A diff between approved and observed configuration is better than hundreds of unlabeled screenshots. A responsibility edge with named parties is better than an entire provider contract pasted without a question. Structure creates leverage.

SEPARATE MECHANICS, ANALYSIS, AND AUTHORITY

AI is well suited to repetitive mechanics: normalization, classification suggestions, cross-reference checks, version comparison, extraction, formatting, and draft generation. It can also assist analysis by surfacing inconsistencies and plausible interpretations. Authority is different. A model does not own the control, accept the risk, approve the scope, or sign the assessment conclusion.

Every workflow should label which layer produced each statement. Source facts should point to their origin. Model observations should remain observations. Practitioner positions should identify the accountable reviewer. Assessment conclusions should stay inside the applicable professional process. The boundary is not ceremonial; it tells the next reader how much trust each statement deserves and who can correct it.

SPEND TOKENS ON DELTAS AND UNANSWERED QUESTIONS

The highest-value recurring use is often change detection. What assets appeared or disappeared? Which responsibility language changed? Which evidence recipe no longer resolves? Which exception passed its decision date? Which control description disagrees with the current configuration source? These questions align model effort with the places where readiness decays.

A token-wise program tracks usefulness, not output volume. Did the workflow reduce time spent reconstructing context? Did it expose a real gap earlier? Did a reviewer accept, revise, or reject its proposed classification? Did it preserve citations? If no decision improved and no uncertainty became clearer, the output may be inexpensive to generate and expensive to govern.

FIVE QUESTIONS FOR THE NEXT WORKING SESSION

- Which recurring readiness task is mostly normalization, comparison, or routing?
- What is the smallest trustworthy context packet for that task?
- Which conclusions must the model be prohibited from making?
- How will every output point back to its source record?
- Who accepts, revises, or rejects the model's contribution?

ONE MOVE TO MAKE NOW

Build one bounded AI context packet for a recurring readiness task: purpose, authoritative inputs, time boundary, limitations, prohibited conclusions, output schema, citations, and named reviewer.

THE PRACTITIONER POSITION

My practitioner position is that AI should improve the memory and preparation of the compliance program while leaving authority visible. The model can help the program notice, compare, organize, and ask. It should not make unsupported certainty cheaper to manufacture.

Spend tokens where judgment gains context. Preserve the provenance. Name the reviewer. Let accountable people make the conclusion.

AUTHORITY AND SOURCES

Independent practitioner perspective. These dispatches are Scott Norton's practitioner analysis and recommendations, not PCI SSC terminology or guidance, not a formal position of Johanson Group, LLP, and not an assessment conclusion about any entity. AI assistance is recorded as production provenance; accountable professional judgment remains human.

Authority classification: PRACTITIONER_POSITION / PRACTITIONER_METHOD / AI_ASSISTED. Authorship and authority classification do not determine legal intellectual-property ownership.

- PCI DSS v4.0.1 - PCI Security Standards Council: https://www.pcisecuritystandards.org/document_library/?class=pcidss&doc=pci_dss