DDRP

Deterministic Document Review Protocol

Comparative Analysis: Industrial Analogues and Alignment

Quantum Inquiry | quantuminquiry.org

1. What Is DDRP?

The Deterministic Document Review Protocol (DDRP) is an open, non-commercial protocol for reconstructing procedural care through the systematic identification, tracking, and resolution of obligation-creating language in professional documents.

DDRP addresses a fundamental evidentiary problem: in regulated, contractual, or legally consequential environments, it is not sufficient to have reviewed a document. The review must be reconstructible. DDRP provides the architecture for that reconstruction.

The protocol operates on three core questions, applied to any artifact under review:

Q1	Did obligation-creating language appear in the document?
Q2	Was that language detectable at the time of production?
Q3	Was each obligation structurally resolved, or left open?

Answering these questions does not require interpreting the meaning of an obligation, assessing its legal validity, or making compliance recommendations. DDRP deliberately stops at the structural layer. Whether an identified obligation is valid, appropriate, or satisfied in substance is a question for human judgment. DDRP creates the evidence base on which that judgment operates.

2. How DDRP Works

2.1 Design Constraint: Determinism Over Coverage

DDRP trades linguistic coverage for inspectability. A probabilistic system might surface an obligation phrased in unusual or metaphorical terms that a rule-based system would miss. But a probabilistic system cannot guarantee it will find that same obligation again if conditions change, and its reasoning cannot be decomposed into auditable steps.

DDRP accepts this coverage trade-off. The protocol defines a known lexical dictionary of obligation-creating terms. Every instance of a defined term is captured with 100% precision. Every gap in the dictionary is a known unknown, not a hidden failure. This is the architectural basis for producing an auditable record of procedural care.

2.2 The Four-Stage Pipeline

Stage 1 — Ingestion and Layout Zoning:

The document is parsed with coordinate-aware text extraction. Physical structure is preserved: headers, footers, appendices, and tables are identified and zoned. This ensures that the location of an obligation within the document is part of the evidentiary record.

Stage 2 — Lexical Extraction:

The protocol scans the artifact for terms in the obligation dictionary. The default lexicon is drawn from IETF RFC 2119: MUST, SHALL, REQUIRED, SHOULD, RECOMMENDED, MAY, OPTIONAL. Custom dictionaries extend coverage to domain-specific terms such as INDEMNIFICATION, LIQUIDATED DAMAGES, or TERMINATION FOR CONVENIENCE. No inference is performed. A term is either present or absent.

Stage 3 — Structural Resolution Tracking:

Each identified obligation is logged as an open item. Resolution tracking requires that a corresponding fulfillment token exist in the responsive artifact. The system does not assess whether the response is adequate; it records whether a response exists. Unresolved obligations remain flagged until a human reviewer closes them.

Stage 4 — Audit Trail Generation:

The protocol produces a machine-readable record for each obligation: the search token, document coordinates, artifact hash, timestamp, and resolution status. This record constitutes the evidence of procedural care. It proves that the act of checking occurred and that the result was visible at the time of production.

2.3 What DDRP Does Not Do

- It does not assess legal compliance or validity.
- It does not make recommendations.
- It does not interpret the meaning or intent of any clause.
- It does not replace human judgment on substantive questions.
- It does not use probabilistic or generative AI for any verification function.

These constraints are not limitations. They are the design. A tool that interprets meaning introduces variability. DDRP introduces none. Its value is precisely the auditable proof that specific language was found, or not found, through a process that a third party can verify and reproduce.

3. Industrial Analogues

Four existing systems share substantial architectural alignment with DDRP. Each was independently developed for a specific domain, yet each arrived at the same core design choices: lexical detection over semantic inference, structural observation over interpretation, and human judgment as the required terminus of the review process.

3.1 VisibleThread VT Docs

Overview

VisibleThread is the closest commercial analog to DDRP and the most widely deployed system in this category. It is used by 64% of Fortune 1000 Aerospace and Defense companies and nine of the top fifteen US government contractors. Its core function, document shredding, maps directly to DDRP's obligation extraction stage.

How It Aligns with DDRP

- Dictionary-based lexical scanning: VT Docs scans artifacts for user-defined obligation-creating terms, including SHALL, MUST, WILL, and REQUIRED, with no semantic interpretation.
- Compliance matrix generation: Extracted obligations are structured into a traceable Excel-format output, providing the DDRP-equivalent of a resolution tracking register.
- Version comparison: Side-by-side document analysis pinpoints added, removed, or modified language using deterministic logic, enabling reconstruction of what was visible at a specific point in production.
- Zero hallucination guarantee: VisibleThread's own documentation explicitly frames this as the distinguishing value over generative AI. The system finds exactly what is in the dictionary, every time.

- Human judgment required: VisibleThread's CEO has publicly emphasized that the software eliminates mechanical extraction work but does not displace human assessment of whether an obligation is valid or appropriately addressed.

Comparison Table

Dimension	DDRP	Comparator
Primary Function	Obligation extraction and resolution tracking	Document shredding and compliance matrix generation
Detection Method	Lexical dictionary with coordinate logging	User-defined dictionary scan with Excel output
AI/ML in Verification	None	None (VT Docs core)
Version Tracking	Artifact hashing and delta logging	Deterministic side-by-side comparison
Human Judgment Role	Required for resolution decisions	Required; explicitly emphasized by vendor
Domain	General professional documents	GovCon, Aerospace, Defense, Federal

DDRP diverges from VisibleThread in scope and formality: DDRP is a protocol specification, not a product. VT Docs implements DDRP-equivalent logic within a commercial platform but does not expose the underlying protocol for independent implementation or auditor verification at the specification level.

3.2 IETF RFC 2119 Keyword Lexicon

Overview

RFC 2119, published by the Internet Engineering Task Force, defines a standard vocabulary for expressing obligation levels in technical specifications. It is the de facto lexical foundation for requirements documents across engineering, legal, and federal procurement contexts. DDRP's default obligation dictionary is built on RFC 2119 terms.

Keyword Hierarchy

Keyword	Structural Meaning	DDRP Flag Level
MUST / SHALL	Absolute requirement — non-negotiable obligation	Critical
SHOULD / RECOMMENDED	Strong expectation — deviation requires justification	Major
MAY / OPTIONAL	Discretionary — no structural impact if absent	Informational

How It Aligns with DDRP

RFC 2119 does not perform document review. It defines the language that makes deterministic review possible. Without a standardized obligation vocabulary, lexical scanning produces variable results depending on how individual authors choose to express requirements. RFC 2119 solves this by creating a single, authoritative mapping between words and structural obligation levels.

DDRP's reliance on RFC 2119 is not incidental. It is the reason DDRP can make the claim that a MUST clause triggers a mandatory resolution entry. The claim is valid only because RFC 2119 has defined MUST as an absolute requirement independent of context. When DDRP logs a MUST obligation, the flag level is not a judgment; it is a direct read from a published standard.

The alignment is foundational rather than architectural. RFC 2119 is to DDRP what a unit system is to a measurement instrument: it does not do the measuring, but without it the measurements mean nothing.

3.3 FedRAMP RFC-0024 and OSCAL Telemetry

Overview

FedRAMP RFC-0024, published January 2026, proposes transitioning all FedRAMP Rev5 cloud service providers from human-written compliance narratives to machine-generated deterministic telemetry, using OSCAL (Open Security Controls Assessment Language) as the structured data standard. This is DDRP applied at scale to federal security authorization.

How It Aligns with DDRP

- Procedural care as machine-generated record: RFC-0024 explicitly states that human narratives cannot constitute a factual record of system state. Telemetry must be collected from authoritative sources and must be reproducible. This is the same evidentiary logic that underlies DDRP.
- Deterministic over probabilistic: The RFC explicitly prohibits the use of generative AI outputs as telemetry. Probabilistic inferences are disqualified on the same grounds DDRP disqualifies them: they cannot be decomposed into inspectable, repeatable observations.
- Artifact hashing and coordinate traceability: OSCAL records tie each compliance assertion to a specific, versioned artifact. The combination of hash, timestamp, and structured data field corresponds directly to DDRP's audit trail schema.
- OSCAL as the structural layer: OSCAL models the entire compliance lifecycle from control catalog through assessment results and plan of action, all in machine-readable JSON or XML. This is a domain-specific implementation of DDRP's requirement that obligations be structurally tracked, not narratively described.

Comparison Table

Dimension	DDRP	Comparator
Core Claim	Procedural care must be reconstructible	Compliance posture must be verifiable, not narratable
Detection Mechanism	Lexical scan of obligation-creating terms	Machine-readable control assertions tied to authoritative sources
Prohibition	No AI/ML in the verification layer	Generative AI outputs explicitly disqualified as telemetry
Audit Trail	Token, coordinates, hash, timestamp, status	OSCAL: control ID, timestamp, artifact reference, evidence link
Human Role	Required for resolution and closure	Required for assessment; automation handles evidence collection
Status	Open protocol specification	Draft RFC; final rules due June 2026

The FedRAMP/OSCAL framework operates at a larger architectural scale than DDRP and is domain-constrained to federal security authorization. DDRP is domain-agnostic. But the underlying philosophy, that compliance is a data problem requiring deterministic evidence rather than a documentation problem requiring persuasive narrative, is identical.

3.4 Proposal Quick Start (PQS)

Overview

Proposal Quick Start, developed by Golden Goose AI, is a deterministic toolset targeting the Government Contracting market. Its primary differentiator is deterministic mapping of Federal Acquisition Regulation (FAR) clause notations to their full authoritative text, and cross-referencing of requirements across RFP sections by pattern matching rather than semantic inference.

How It Aligns with DDRP

- FAR clause lookup as deterministic obligation detection: Federal solicitations frequently reference FAR clauses by notation (e.g., FAR 52.212-4) without including full clause text. PQS uses a deterministic lookup engine to resolve these notations to their underlying obligations. This is DDRP's Stage 2 extended to handle indirect obligation references.
- Requirement cross-referencing by pattern matching: PQS links obligations across RFP sections (Statement of Work, Instructions to Offerors, Evaluation Criteria) through structural pattern matching. This directly implements DDRP's structural resolution tracking logic.
- No semantic inference for mapping: Where broader AI-assisted tools use probabilistic discovery to surface thematic connections, PQS uses fixed lookup tables and cross-reference rules. The result is reproducible and auditable.

- Domain specificity: PQS is purpose-built for GovCon solicitations, where the FAR clause structure creates a regularized obligation vocabulary. This makes deterministic lookup more tractable than in unstructured commercial contracts.

Comparison Table

Dimension	DDRP	Comparator
Obligation Source	Any lexical marker in the document	FAR clause notations and direct requirement language
Lookup Mechanism	Dictionary scan with coordinate logging	Deterministic FAR clause table with cross-reference engine
Cross-Section Tracking	Resolution status per obligation	Alignment of Sections C, L, and M by pattern matching
AI/ML in Verification	None	None (mapping layer is deterministic lookup)
Domain	Domain-agnostic	Federal government contracting
Human Role	Required for resolution	Required; GovCon expert review emphasized

PQS narrows DDRP's domain to a specific regulatory vocabulary (FAR/DFARS) and extends the obligation detection stage to handle indirect clause references. The core protocol logic, find every obligation by defined pattern, track its resolution, require human judgment at the close, is the same.

4. Synthesis: Where DDRP Sits

The four comparators above were developed independently, by different organizations, for different markets. None of them reference DDRP. Yet each converged on the same architectural principles: lexical detection, structural observation, no probabilistic inference in the verification layer, and a hard requirement for human judgment at the point of resolution.

This convergence is not coincidental. It reflects a functional constraint that any system attempting to reconstruct procedural care must respect: the evidentiary record must be inspectable, and to be inspectable it must be deterministic.

Property	DDRP	VisibleThread	RFC 2119	FedRAMP/ OSCAL	PQS
Lexical Detection Only	Yes	Yes	Defines lexicon	Yes	Yes
No AI/ML in	Yes	Yes	N/A	Yes	Yes

Verification				(mandated)	
Structural Resolution Tracking	Yes	Yes	No	Yes (OSCAL)	Yes
Audit Trail Output	Yes	Partial	No	Yes	Partial
Human Judgment Required	Yes	Yes	Yes	Yes	Yes
Domain Agnostic	Yes	No (GovCon/A&D)	Yes	No (FedRAMP)	No (GovCon)

DDRP occupies a distinct position in this landscape. It is the only domain-agnostic, open-protocol specification in the group. The commercial and regulatory systems above are implementations of DDRP-equivalent logic within constrained domains. DDRP provides the generalizable protocol from which domain-specific implementations can be derived and against which they can be evaluated.

The practical implication: organizations operating outside the GovCon and federal security domains, where VisibleThread, PQS, and FedRAMP/OSCAL are native, can implement DDRP directly as their document review architecture. RFC 2119 provides the default obligation lexicon. The four-stage pipeline provides the processing model. The audit trail schema provides the evidentiary output. No commercial platform is required.

5. Conclusion

DDRP is not a response to the rise of AI. It is the articulation of a review discipline that has always been necessary in high-stakes document environments and that probabilistic AI makes more, not less, urgent.

The systems reviewed here, VisibleThread, RFC 2119, FedRAMP RFC-0024, and Proposal Quick Start, confirm that the DDRP design philosophy is not idiosyncratic. It is the architecture that practitioners across multiple high-assurance domains have independently arrived at when the requirement is not to summarize documents but to prove that they were reviewed.

The distinction matters. Summarization produces a description of what is present. DDRP produces a record of the process by which presence was determined. The first is useful. The second is what an auditor, a judge, or a contracting officer requires when procedural care is in question.

> *The goal is not to understand documents. The goal is to prove, reconstructibly, that they were examined.*