

PUBLIC PROTOCOL SPECIFICATION

WaveFront Protocol

Public Draft 0.9

By Németh Attila

Contents

0. Document Metadata	14. Facilitator Intervention
1. Abstract	15. Integration Proposal
2. Status of This Document	16. Human Ratification
3. Problem Definition	17. Versioned Next Front
3.1 Originating Interaction	18. Propagation Ledger
3.2 Why WaveFront?	19. Meeting Lifecycle
4. Protocol Scope	19.1 Before
4.1 In Scope	19.2 During
4.2 Out of Scope	19.3 End of Meeting
4.3 Intended Meeting Context	19.4 Next Relevant Meeting
4.4 What WaveFront Is Not	20. Governance and Authority
5. Design Principles	21. Privacy, Transparency and Monitoring Boundaries
6. Actors and Roles	22. Protocol Conformance
7. The WaveFront State Model	22.1 Normative Language
7.1 Current Front	22.2 WaveFront Core Requirements
7.2 Statement	23. What Does NOT Constitute WaveFront
7.3 Dependency	24. Abstract Data Model
7.4 Divergence	25. Worked Protocol Traces
7.5 Next Front	25.1 Product concept / Driftline
8. Contextual Participant Intervention	25.2 Positioning / LedgerNest
8.1 Core Intervention Set	25.3 Research synthesis / CareLoop
EXACTLY	25.4 Investor narrative / GridTrace
CLARIFY	25.5 Product review / LoopCart
CHALLENGE	25.6 Founder and board session / SignalPort
PARK	26. Relationship to Existing Approaches and Prior Art
9. Extended Intervention Model	27. Known Risks and Failure Modes
10. System Processing States	28. Versioning
11. Significance Assessment	29. Glossary
12. Dependency Analysis	30. References
13. Convergence Point	

0. Document Metadata

Field	Value
Title	WaveFront Protocol
Status	Public Draft 0.9
Author	Németh Attila
Publication date	2026-09-23
Canonical URL	https://wavefrontprotocol.com/protocol/0.9/
Contact	protocol@wavefrontprotocol.com
Rights	All Rights Reserved
Copyright	Copyright © 2026 Németh Attila.

1. Abstract

WaveFront is a protocol for turning meaningful participant interventions during a live meeting into a traceable, human-ratified next shared state. It addresses a recurring choice in consequential working sessions: interrupt the speaker, wait and risk losing the point, or place a comment in a parallel channel where it may never affect the reasoning it concerns.

WaveFront is not primarily a transcript, meeting summary, chat, whiteboard, or AI notetaker. Its core loop is:

context-attached participant intervention → semantic and dependency
interpretation → significance assessment → natural convergence point →
integration proposal → human ratification → versioned next shared state

The system may capture, classify, interpret, connect, and propose. It may not silently make participant input canonical or create the official shared state. Humans ratify changes. A ratified state can preserve an unresolved disagreement; convergence means a controlled transition to a usable shared state, not forced consensus.

2. Status of This Document

Public Draft 0.9 is a working public specification. It defines the current protocol vocabulary, state model, normative core, and examples sufficiently for review and implementation experiments. It is not a claim that every user-interface or governance question has been settled.

This draft may be superseded by later versions. Once a version-specific document is published, that document remains an immutable historical record even when a later version supersedes it. The publication date of this working draft is **2026-09-23**.

The key words **MUST**, **MUST NOT**, **SHOULD**, **SHOULD NOT**, and **MAY** have the meanings defined in Section 22.

3. Problem Definition

Important corrections frequently enter a meeting badly. A participant can interrupt and break the presenter's logical flow; wait and return to the point after its context has faded; or write in parallel chat, splitting attention and leaving the comment detached from the statement and downstream reasoning it affects. Less dominant or slower-to-respond participants face an additional airtime disadvantage.

These failure modes are not equivalent:

- **Interruption** preserves immediacy but imposes a coordination cost and may turn a correction into an unstructured debate.
- **Waiting** preserves speaking order but risks memory loss, a late context switch, or reasoning continuing from a defective premise.
- **Parallel chat** enables contribution without taking the floor, but creates a second stream that presenters and participants must monitor. The important item may remain socially acknowledged but logically unused.
- **Loss of contextual corrections** occurs when an observation is retained as text but its target, scope, evidential status, or effect on later claims is not retained.

Recording and summarization do not by themselves solve this problem. A recording preserves what happened; a summary compresses it after or during the fact. Neither necessarily determines which live statement a correction affects, whether later reasoning depends on that statement, when to revisit it, who may authorize a revision, or how the resulting shared state differs from its predecessor. WaveFront therefore treats the meeting output as a governed state transition rather than merely a better record.

3.1 Originating Interaction

This section is non-normative. The concept began with two recurring human interaction patterns, not with market validation. In the first, **CLARIFY**, there is no real disagreement: the current wording sounds as though it contradicts the prior point even though the participants mean substantially the same thing. In the second, **EXACTLY**, a new formulation says more clearly what a participant was already trying to express.

Both patterns expose the same interaction problem. A useful correction or confirmation may arrive while someone else still has the floor, when speaking up would break the flow but waiting may detach the contribution from its context. The originating promise of WaveFront is therefore: **a participant should not need to interrupt the meeting in order to change its eventual outcome.**

3.2 Why WaveFront?

In wave physics, a **wavefront** is the set of points on a wave that are in the same phase. The direction of propagation is normal—perpendicular—to the front. Under Huygens's principle, every point on a wavefront can be treated as a source of secondary wavelets; the envelope of those wavelets forms the next wavefront.^[11]

WaveFront uses this as a conceptual and naming model for collaborative reasoning:

Physical model	WaveFront protocol
Current wavefront	Current shared reasoning state
Point on the front	A Statement
Local displacement / secondary wave	A participant Intervention
Interference	Reinforcing, weakening, or conflicting participant effects
Next wavefront	The integrated, human-ratified Next Front
Damping	A rejected, irrelevant, resolved, or otherwise non-propagating Intervention
Refraction / change of direction	A material correction changes the trajectory of the reasoning

The metaphor does not prove or scientifically validate a collaboration model. It makes the protocol's transition legible: a Current Front $F(n)$ receives local interventions; their effects are interpreted and integrated only through human ratification; the result is the Next Front $F(n+1)$. The formal transition in Section 17 retains this notation.

The metaphor is also a protocol boundary. A system that only stores transcript, comments, and summaries but never produces a versioned next shared state does not implement the defining WaveFront transition.

4. Protocol Scope

4.1 In Scope

WaveFront specifies:

- a shared, live representation of the meeting's current working state;
- statement-level or logical-unit anchoring for participant interventions;
- a small core intervention vocabulary and an extensible semantic model;
- interpretation of significance, scope, evidence, and downstream dependencies;
- silent-by-default collection and surfacing at a natural convergence point;
- integration proposals and explicit human ratification;
- versioned Current Front to Next Front transitions; and
- provenance linking source statements, interventions, proposals, decisions, and remaining divergence.

4.2 Out of Scope

The protocol does not prescribe a database, language model, speech recognizer, video platform, visual design, scoring algorithm, or proprietary significance threshold. It does not define employee performance analytics, rank participants, decide who is correct, guarantee consensus, or replace an organization's legal and decision-authority structures. It does not require transcription, although an implementation may use it with notice and appropriate controls.

4.3 Intended Meeting Context

The primary context is a working session of roughly 4–12 people in which reasoning changes live and the outcome matters: product-concept definition, positioning, research synthesis, product review, strategy or board work, and similar decision-forming sessions. Pure status broadcasts, simple information delivery, routine one-to-one conversations, and meetings with no shared state to form may not justify another interaction layer.

4.4 What WaveFront Is Not

WaveFront is not a mandatory replacement for a video call, chat service, document repository, or meeting platform. It is a stack-augmenting protocol. It is also not synonymous with an AI-generated transcript, a comment rail, a graph of conversation, or a polished summary. Those can support the protocol but do not constitute it without the complete loop in Section 1.

5. Design Principles

1. **Flow before interruption.** Preserve the main thread where a point can be handled silently; return an intervention earlier when subsequent reasoning or an imminent decision materially depends on a challenged, corrected, or unverified Statement.
2. **Participant agency without airtime competition.** A participant must be able to affect the shared state without first winning the floor.
3. **Silent by default.** Anchored interventions enter a side layer and do not automatically seize the main conversation.
4. **Human ratification.** AI may propose; humans ratify. No generated interpretation becomes official merely because it looks canonical.
5. **Convergence is not forced consensus.** Convergence produces a usable, explicit state. It may record competing interpretations or minority disagreement.
6. **Preserve unresolved divergence.** A material unresolved objection remains visible and linked to its subject.
7. **Traceability without employee surveillance.** Explain how shared reasoning changed without turning contribution traces into individual productivity or speaking-time scores.
8. **Shared state over transcript.** Optimize for what the group can responsibly continue from, while retaining history as provenance.
9. **Minimal participant friction.** Keep the common intervention set small; offer richer types contextually rather than as a permanent control panel.
10. **Augment the stack.** Permit integration with existing meeting and document tools instead of requiring organizational migration.
11. **Reversibility.** A later Front can revise an earlier conclusion while preserving the chain of reasons and authority.
12. **Visible uncertainty.** Scope, evidence limits, unverified attribution, and measurement confounders must not be polished away.

6. Actors and Roles

- **Participant:** attends the meeting, reads the Current Front, and creates anchored interventions. A Participant has contribution rights but does not thereby have integration or decision authority.
- **Facilitator:** manages flow, selects convergence points, combines related interventions, requests clarification, and presents or edits integration proposals. Facilitation is procedural authority, not automatic subject-matter authority.
- **Decision Owner:** has recognized authority for a defined decision or state transition. The Decision Owner ratifies where organizational rules assign that authority, or records the room's ratification where authority is collective.

- **WaveFront System:** maintains representations, performs supported interpretation and dependency analysis, suggests handling, and records provenance. It has no independent ratification authority.
- **External Guest:** a Participant with limited access, often joining by a low-friction link. Guest identity, visibility, intervention rights, and retention must be explicit.

Roles may overlap: a presenter may be a Participant and Facilitator; a Facilitator may also be the Decision Owner. Their authorities remain distinct even when held by one person. An implementation must show whether an action is a participant contribution, a facilitation action, or an exercise of decision authority.

7. The WaveFront State Model

7.1 Current Front

The **Current Front**, denoted $F(n)$, is the presently ratified or explicitly designated shared working state from which the meeting reasons. It is a concise, structured representation of current facts, assumptions, interpretations, proposals, questions, and decisions relevant to the session. It is not a verbatim transcript and does not claim that every displayed statement is timeless truth. Each statement retains its type, status, scope, and provenance where those matter.

At a meeting's opening, the initial Current Front may be an agreed pre-read state or a presenter's declared working hypothesis. If it is not yet jointly ratified, that status must be visible; presentation must not masquerade as agreement.

7.2 Statement

A **Statement** is the smallest logical unit in a Front to which an intervention can be meaningfully anchored. Depending on context, it may be a:

- **fact** — asserted observation or datum;
- **assumption** — premise treated provisionally as true;
- **interpretation** — meaning inferred from observations;
- **conclusion** — claim derived from prior reasoning;
- **proposal** — possible action, wording, experiment, or solution;
- **question** — uncertainty whose answer may affect the state; or
- **decision** — authorized choice with an applicable scope.

A single sentence may contain multiple logical Statements and should be separable when interventions affect only one. Conversely, a table row, diagram node, or spoken proposition may be a Statement even when it is not a grammatical sentence.

7.3 Dependency

A **Dependency** is a directed reasoning relationship in which a later Statement relies on, narrows, explains, evidences, or operationalizes another. Dependencies need not form a perfect formal proof. They must be sufficient to recognize when a correction to an earlier premise can materially affect later conclusions or decisions.

7.4 Divergence

A **Divergence** is a material difference in interpretation, evidence assessment, preference, or proposed action that has not been resolved. Divergence is part of the shared state when it remains relevant. It includes the competing positions, their anchors, known evidence, ownership or next check where applicable, and status. It is not a discarded comment or a euphemism for failure.

7.5 Next Front

The **Next Front**, $F(n+1)$, is the new shared working state produced after interventions have been interpreted, surfaced at an appropriate convergence point, and ratified by authorized humans. It differs from a draft Integration Proposal: the proposal is offered; the Next Front is accepted. Once active, $F(n+1)$ becomes the Current Front for subsequent reasoning.

8. Contextual Participant Intervention

An **Intervention** is participant input attached to a specific Statement or identifiable logical unit. The anchor preserves what prompted the input and permits dependency analysis. If exact anchoring is temporarily impossible, the input must remain unintegrated until its target is clarified; proximity in chat is not a reliable substitute.

An intervention can add text, evidence, a source, or structured metadata. Creating one does not interrupt the speaker, validate its contents, revise its target, or make it canonical. The expected default is capture now, assess continuously, and return at a suitable convergence point.

8.1 Core Intervention Set

The participant-facing core contains exactly four common actions:

EXACTLY

- **Purpose:** identify wording or reasoning that accurately expresses the participant's view or the emerging direction.
- **Semantic meaning:** strong confirmation, not merely applause.
- **Expected effect:** may increase confidence that a Statement should remain or be emphasized and may combine with similar signals.
- **Does not automatically:** prove the Statement, establish unanimity, or ratify it.

CLARIFY

- **Purpose:** correct precision, meaning, cause, scope, data, or wording without necessarily opposing the underlying direction.
- **Semantic meaning:** "this needs qualification or a more accurate formulation."
- **Expected effect:** may produce a revised Statement, narrower scope, separated claims, or a request for clarification.
- **Does not automatically:** reject the Statement or imply adversarial disagreement.

CHALLENGE

- **Purpose:** materially contest data, an assumption, inference, conclusion, or proposed action.
- **Semantic meaning:** "reasoning should not continue as though this point were settled."
- **Expected effect:** triggers significance and dependency review and may require earlier convergence if downstream reasoning is at risk.
- **Does not automatically:** make the challenger correct, stop the meeting, or delete the challenged Statement.

PARK

- **Purpose:** preserve a relevant side issue that would currently disperse the meeting.
- **Semantic meaning:** "retain this, but do not resolve it in the present decision thread."
- **Expected effect:** records the item with an owner or return condition where appropriate and excludes it from the immediate integration.
- **Does not automatically:** dismiss the issue, mark it unimportant, or guarantee action without ownership.

9. Extended Intervention Model

Extended types express distinctions repeatedly needed in the source scenarios. They need not appear as a permanent full button set; the interface may offer them contextually or the system may suggest a type for human confirmation.

- **EVIDENCE:** attaches data, observation, customer language, research, or counterexample. It should retain source, sample, time window, and confidence when known. Evidence may strengthen, weaken, or constrain a claim.
- **OPEN QUESTION:** records an unanswered question that blocks closure or a forthcoming decision. It should identify what depends on the answer.
- **SCOPE:** states that a claim holds only for a segment, period, use case, or condition, or that two people are reasoning at different scopes.
- **ASSUMPTION:** identifies an unproven premise or testable hypothesis and, where possible, what result would support or disconfirm it.
- **RISK:** records a credible adverse consequence or boundary that should affect the proposal or decision.

- **MEASUREMENT LIMIT:** identifies inadequate sample size, unsuitable metric, confounder, incomparable cohorts, attribution problem, or insufficient observation period.
- **PROPOSAL:** offers alternative wording, action, experiment, or solution. It remains a candidate until assessed and ratified.
- **VERIFY:** marks uncertain content, authorship, source, quotation, number, or attribution. A VERIFY item must not be integrated as fact until a human or trusted source verifies it.
- **DECISION POINT:** identifies a choice that needs an owner, options, missing evidence, threshold, deadline, or review condition.

These types may coexist with the core semantics. For example, a participant can CLARIFY a Statement by attaching EVIDENCE and a SCOPE limitation.

10. System Processing States

The following are processing and outcome states, **not participant reaction buttons**:

- **FACILITATOR SUGGESTION:** a noncanonical recommendation to combine items, ask a question, distinguish scopes, pause, or choose a convergence point. It may originate from software or a human facilitator.
- **INTEGRATION PROPOSAL / INTEGRATION:** INTEGRATION PROPOSAL is a candidate revision showing what would change and why. INTEGRATION is the resulting accepted revision after human ratification. Implementations must not blur the proposed and ratified states.
- **DECISION:** an authorized meeting outcome. It should identify owner, scope, duration, success or stop conditions, and review point where applicable.
- **NOT ELEVATED:** the intervention remains in the record but was not proposed for the main Front at that time because it was low-significance, duplicative, unverified, irrelevant to the active decision thread, or better parked.

NOT ELEVATED does not mean deletion. The item remains retrievable, linked to its anchor, and open to later human review or elevation. The reason for non-elevation should be available.

11. Significance Assessment

Significance assessment determines whether, how, and when an intervention should return to the main reasoning. It is a conceptual judgment, not a prescribed numerical score. An implementation should consider:

- whether the intervention affects a central premise;
- whether later Statements or an imminent decision depend on the affected Statement;

- whether multiple independent interventions point to the same issue;
- whether apparent disagreement is real or only a wording mismatch;
- whether the intervention changes scope;
- whether it introduces new evidence or changes the evidential strength;
- whether it identifies blocking uncertainty;
- whether it exposes a measurement limitation or confounder; and
- whether it changes the available options, authority, reversibility, or impact of a decision.

The system may prioritize and explain. It must not invent confidence from repetition, treat status or seniority as truth, or conceal a material minority objection. Public Draft 0.9 defines no scoring weights, thresholds, or proprietary algorithm.

12. Dependency Analysis

Comments cannot be handled as isolated annotations because meeting reasoning is sequential. If a market conclusion depends on a sample estimate and that estimate is challenged, merely attaching the comment to the number is insufficient: claims about positioning, product scope, and a proposed test may also require review.

Dependency analysis identifies the affected subgraph of reasoning. It asks: Which subsequent Statements use this premise? Does a narrower scope leave some conclusions valid? Can mixed hypotheses be separated? Is a decision about to rely on an unverified attribution? The result informs urgency and an Integration Proposal; it does not autonomously invalidate or rewrite downstream Statements.

When dependencies are uncertain, the proposal should say so and invite human confirmation. A system should prefer an explainable “these later claims may depend on this premise” over false formal precision.

13. Convergence Point

A **Convergence Point** is a natural breakpoint at which one or more significant interventions can be brought into the main meeting, interpreted together, and offered for human ratification without unnecessarily destroying conversational flow.

Typical points include:

- the end of a logical block;
- a topic transition;
- a presenter's summary;
- immediately before a decision;

- before solution design begins from a materially challenged problem statement; and
- before reasoning continues from a materially challenged premise.

The last two may justify an earlier pause than the presenter planned. Silence remains the default, but a Convergence Point may need to occur earlier when subsequent reasoning or an imminent decision materially depends on a challenged, corrected, or unverified Statement.

Convergence does not require agreement. A valid outcome can be: adopt a correction; preserve two interpretations; narrow a claim; record an open question; park a branch; or decide under explicit uncertainty.

14. Facilitator Intervention

The Facilitator turns the side layer into a manageable choice without becoming an invisible editor. Supported patterns include:

- **Combining similar comments:** group interventions aimed at the same semantic issue while retaining individual provenance.
- **Wording mismatch:** identify apparent disagreement where parties mean the same thing at different levels or use different professional language.
- **Scope mismatch:** state the segment or condition for which evidence holds rather than generalizing.
- **Separating mixed hypotheses:** split a proposed solution or explanation so its components can be tested independently.
- **Outcome versus replaced workflow:** distinguish the benefit a buyer seeks from the present behavior or tool the product actually replaces.
- **Wedge versus platform:** distinguish an initial repeatable use case from longer-term expansion without allowing either to impersonate the other.
- **False binary:** preserve conditional or layered explanations when “A or B” is not supported—for example, execution weakness and an unclear ICP can coexist.
- **Measurement confounder:** park a delivery issue from the product decision while accounting for it in evaluation.
- **Uncertain attribution or source:** block integration until verified rather than attaching authority to generated content.
- **Choosing a natural integration point:** wait until the end of a block unless subsequent reasoning or an imminent decision materially depends on the affected Statement.

Facilitator actions must be visible and contestable. Combining comments must not erase differences among them.

15. Integration Proposal

An **IntegrationProposal** describes a candidate state change. It should contain:

1. the original Statement or Front fragment;
2. the interventions and evidence considered;
3. the interpretation, including relevant dependencies;
4. the proposed replacement, qualification, addition, split, or divergence;
5. why the change matters now; and
6. unresolved questions, sources requiring verification, and downstream Statements requiring review.

A proposal may recommend no textual change—for example, retain the Statement and mark a challenge NOT ELEVATED with reasons. It must distinguish quotation from system paraphrase and participant authorship from machine inference.

16. Human Ratification

AI may propose. Humans ratify.

The person or group with the applicable authority may:

- **accept** the proposal;
- **edit** it before accepting;
- **reject** it with an optional reason;
- **ask the room / request clarification**;
- **preserve as unresolved divergence**; or
- **park** it with an owner or return condition where appropriate.

Ratification must identify the human authority and time. Silence, lack of objection, an AI confidence score, or visual placement in the main pane is not ratification. If the Facilitator lacks integration rights, they may prepare the proposal but cannot accept it. If authority is collective, the implementation should record the meeting's declared rule rather than infer consensus.

17. Versioned Next Front

The state transition is:

$$F(n) \rightarrow \text{interventions} \rightarrow \text{interpretation} \rightarrow \text{convergence} \rightarrow \text{ratification} \rightarrow F(n+1)$$

$F(n)$ remains in history. Interventions do not directly mutate it. Interpretation and an Integration Proposal form a candidate transition. Ratification authorizes $F(n+1)$, which then becomes the Current Front.

A valid Next Front may contain explicit unresolved disagreement. For example, it may record that the current evidence supports a product effect in one subgroup while the cause of the wider phenomenon remains disputed. This is often more accurate and actionable than a prematurely unified statement.

18. Propagation Ledger

The **Propagation Ledger** is the traceability record for state transitions. It connects:

- the original Statement and Front version;
- each anchored Intervention and its author;
- Evidence and source, where applicable;
- the Facilitator or system suggestion;
- the Integration Proposal and its rationale;
- the human Ratification, authority, and disposition;
- the resulting Statement and Next Front; and
- any preserved Divergence, parked branch, or NOT ELEVATED item.

The Ledger answers not only “what changed?” but “what prompted the change, what reasoning propagated, who authorized it, and what remained unresolved?” It is not an employee scorecard. Access and retention may be restricted without breaking the conceptual chain, provided authorized users can inspect the provenance required for their work.

19. Meeting Lifecycle

19.1 Before

The organizer states the meeting purpose, expected output, roles, decision authority, participant access, transcription status, and retention policy. A pre-read or working hypothesis may seed $F(\emptyset)$, visibly labeled as ratified, provisional, or presenter-proposed. The protocol should require no lengthy organizational migration, especially for External Guests.

19.2 During

The Current Front follows the reasoning in concise logical units. Participants attach silent interventions. The system assesses meaning, significance, and dependencies, and the Facilitator handles significant clusters at Convergence Points. Humans ratify any canonical change. Lower-significance items remain traceable without repeatedly disrupting the meeting.

19.3 End of Meeting

The meeting closes with the latest ratified Front, Decisions, owners and conditions, open questions, parked items, and preserved Divergences. The transcript or recording remains history; it is not interchangeable with the ratified state. If the room has not ratified a proposed synthesis, the export must label it as proposed rather than official.

19.4 Next Relevant Meeting

The next session begins from the latest applicable ratified Front, not from a newly generated summary that silently replaces it. New evidence can reopen prior Statements. The Propagation Ledger makes reversals intelligible and preserves the historical Fronts rather than rewriting them.

20. Governance and Authority

- **Integration rights:** must be declared. Participant contribution rights do not imply permission to alter the Front.
- **Facilitator power:** is bounded by visible procedure. A Facilitator can sequence and propose but cannot hide material input or assume decision ownership.
- **Decision ownership:** attaches to a defined question and scope. A board may approve a testing envelope while management owns tactics; a product owner may own a product choice without owning factual attribution.
- **Minority disagreement:** remains in the Front or Ledger when material to the decision, its risks, or later review. Inconvenience is not a deletion criterion.

- **Reversibility:** later evidence may produce a new Front. The earlier version and its authority remain visible.
- **Visible authorship and ratification:** participant words, system interpretations, facilitator edits, and ratifying human actions must be distinguishable.
- **No invisible AI authority:** generated language must never acquire authority merely through fluency, layout, or automation.

An organization may define local quorum, voting, or escalation rules. Those rules must be explicit and cannot transfer ratification to the WaveFront System.

21. Privacy, Transparency and Monitoring Boundaries

Participants should know whether transcription or recording occurs and what the live shared representation contains. Access, External Guest permissions, retention, export, and deletion policies should be explicit at meeting level or through a clearly applicable organizational policy.

Traceability must not require employee-performance surveillance. Personal speaking-time scores, participation rankings, sentiment-based performance judgments, and individual productivity scoring are not part of the protocol. Implementations should minimize collected personal data and prevent a contribution record intended for shared reasoning from silently becoming an employment-evaluation feed.

The system may record authorship necessary for provenance. It should also support appropriate access controls and retention limits. An unresolved disagreement must not disappear because it is inconvenient, politically costly, or expressed by a minority. Privacy controls and intellectual honesty must be designed together: restricting access is different from falsifying the state.

22. Protocol Conformance

22.1 Normative Language

The capitalized keywords **MUST**, **MUST NOT**, **SHOULD**, **SHOULD NOT**, and **MAY** in this document are to be interpreted as described in RFC 2119[12] and RFC 8174[13] when, and only when, they appear in all capitals.

- **MUST / MUST NOT:** an absolute requirement or prohibition for WaveFront Core conformance.
- **SHOULD / SHOULD NOT:** a strong recommendation; deviation requires a documented reason appropriate to the context.
- **MAY:** an optional capability consistent with the protocol.

22.2 WaveFront Core Requirements

WF-1 Shared Current Front. A conforming implementation **MUST** maintain a participant-visible Current Front distinct from the transcript and **MUST** show whether its Statements are ratified or provisional.

WF-2 Context anchoring. Every Intervention considered for integration **MUST** be anchored to a Statement or identifiable logical unit. An unclear anchor **MUST** be resolved before integration.

WF-3 Silent participant intervention. Participants **MUST** be able to contribute an anchored Intervention without automatically interrupting the main conversation.

WF-4 No automatic canonicalization. Participant comments and generated interpretations **MUST NOT** automatically become part of the canonical Front.

WF-5 Significance/dependency handling. The implementation **MUST** support assessment of significance and relevant downstream dependencies and **SHOULD** expose an understandable rationale for elevation.

WF-6 Convergence Point. Significant interventions **MUST** be returned to the main reasoning at a declared natural Convergence Point rather than merely left in a parallel stream. The Convergence Point **SHOULD** occur earlier when subsequent reasoning or an imminent decision materially depends on a challenged, corrected, or unverified Statement.

WF-7 Human ratification. Every canonical transition **MUST** be ratified by an identifiable authorized human or explicitly governed human group. The system **MUST NOT** ratify.

WF-8 Versioned Next Front. Ratification **MUST** produce a distinct, versioned Next Front without overwriting the prior Front.

WF-9 Preserved unresolved divergence. A material unresolved disagreement **MUST** be representable in the shared state and **MUST NOT** be reported as consensus.

WF-10 Traceable propagation history. A conforming implementation **MUST** maintain a Propagation Ledger linking original Statement, Intervention, proposal, Ratification, resulting Statement, and Divergence where applicable.

WF-11 NOT ELEVATED is not deletion. A NOT ELEVATED item **MUST** remain retrievable with its anchor and **SHOULD** retain the non-elevation reason. Non-elevation **MUST NOT** silently delete it.

WF-12 No required employee-performance surveillance. Conformance MUST NOT require speaking-time, productivity, sentiment, or contribution scoring of employees.

WF-13 Authority visibility. The implementation MUST distinguish participant authorship, system inference, facilitation, and human ratification.

WF-14 Verification boundary. Content marked VERIFY, including uncertain attribution, MUST NOT be integrated as established fact until verified by a human or trusted source.

WF-15 Lifecycle export. The meeting output MUST distinguish transcript/history from the latest ratified state, decisions, open questions, parked items, and Divergences.

WF-16 Transparency. Participants MUST be informed when transcription occurs, and applicable access and retention conditions MUST be available.

Conformance concerns protocol behavior, not interface similarity. A text interface, document extension, meeting-platform integration, or other product may conform if it satisfies these requirements.

23. What Does NOT Constitute WaveFront

Each item below may be useful and may support an implementation, but alone is insufficient:

- **Transcript + comments:** stores speech and annotations but need not interpret dependencies, reintegrate comments, or ratify a Next Front.
- **AI summary:** compresses content but can silently choose meaning and lacks participant-anchored propagation and authority.
- **Shared notes:** enable joint editing but do not necessarily distinguish proposals from ratified state or preserve provenance.
- **Decision extraction:** finds apparent decisions but does not establish decision ownership, ratification, or the state from which the decision arose.
- **Meeting chat:** supports parallel participation but leaves important contributions in a competing stream.
- **Editable collaborative document:** provides real-time coauthoring but not the intervention-to-convergence governance loop.
- **Conversation graph:** represents relationships among utterances or claims but need not provide natural reintegration, human authority, or a versioned shared state.

A conforming implementation requires the complete state-transition loop, including anchored intervention, interpretation, significance and dependency handling, a Convergence Point, an Integration Proposal, human Ratification, a versioned Next Front, and traceable propagation.

24. Abstract Data Model

This model is conceptual and does not prescribe storage technology.

- **Meeting:** a bounded working session with purpose, access rules, Actors, and a sequence of Fronts.
- **Actor:** a human or system identity with one or more declared roles. Only human Actors can ratify.
- **Front:** a versioned shared working state composed of Statements, Dependencies, Decisions, and Divergences.
- **Statement:** an anchorable logical unit with type, status, content, scope, authorship, and provenance.
- **Dependency:** a directed relationship between Statements indicating relevant reasoning reliance or qualification.
- **Intervention:** an Actor-authored, time-stamped input anchored to a Statement or logical unit, with core and optional extended semantics.
- **Evidence:** sourced material attached to an Intervention or Statement, including known limitations.
- **IntegrationProposal:** a candidate transformation from part of one Front toward a possible next state, linked to supporting and opposing inputs.
- **Ratification:** a human-authority disposition of an IntegrationProposal: accept, edit-and-accept, reject, clarify, preserve divergence, or park.
- **Divergence:** a structured unresolved difference linked to affected Statements and evidence.
- **Decision:** a ratified choice with owner, scope, conditions, and review information where applicable.
- **PropagationLedger:** the ordered provenance graph connecting all of the above across Front versions.

A Meeting has one active Current Front and historical Fronts. A Front contains Statements connected by Dependencies. Interventions target Statements; Evidence can support an Intervention or Statement. One or more Interventions inform an IntegrationProposal. A Ratification disposes of that proposal and, if accepted, produces Statements, Decisions, or Divergences in a Next Front. The PropagationLedger connects the complete chain.

25. Worked Protocol Traces

The following compact traces are derived from the six canonical, fictional meeting simulations. Their company names, participant data, figures, and quoted meeting facts are simulation material, not claims about real companies.

25.1 Product concept / Driftline

- **Situation:** Driftline must choose a six-week product test for field-service dispatching.

- **Initial Front:** "Driftline automatically optimizes service companies' daily dispatch plan," presented as a working hypothesis rather than agreement.
- **Relevant Statement:** the product core is automatic daily replanning after disruption.
- **Interventions:** customer success CLARIFY says dispatchers ask not to phone everyone after a technician drops out; sales CHALLENGE says full automation creates fear; EVIDENCE limits a 47-minute estimate to three manual logs; engineering CLARIFY says the first version can rank three explained alternatives, not rewrite the day automatically; design CHALLENGE reframes the dispatcher from bottleneck to absorber of uncertain exceptions.
- **Why they matter:** together they challenge the central problem, promise, evidence, and feasible scope on which solution design would depend.
- **Facilitator/system handling:** combine the comments into "exception resolution rather than full automation," separate two solution hypotheses—alternative generation versus a unified impact view—and park platform potential and automated customer communication.
- **Convergence point:** before moving from problem definition into solution ideas, then again before committing to the experiment.
- **Integration proposal:** focus on multi-visit daily exceptions; show affected parties and consequences plus ranked alternatives without removing dispatcher control; validate through a concierge workflow.
- **Human ratification:** Luca accepts the revised working hypothesis; Nóra, the decision owner, approves the six-week validation rather than building an "AI dispatch platform."
- **Next Front:** make the moment after the daily plan breaks less chaotic by helping a dispatcher reach an acceptable decision; test an impact card plus three ranked alternatives at three high-exception-frequency pilots, measuring time to approved resolution, contacts required, and decision confidence.
- **Remaining divergence:** whether this can become a standalone wedge and which interface ultimately fits remain open; full automatic replanning stays a longer-term possibility.
- **Protocol lesson:** a cluster of anchored corrections can change the problem definition without turning the presentation into continuous interruption.

25.2 Positioning / LedgerNest

- **Situation:** LedgerNest needs precise positioning for its month-end finance workflow.
- **Initial Front:** "LedgerNest uses AI to automate modern finance teams."
- **Relevant Statement:** position it as AI-powered month-end close automation, comparable to close-management platforms.
- **Interventions:** finance CHALLENGE says the product does not automate ledger, consolidation, or audit; customer success CLARIFY says daily use centers on missing documents and stalled reconciliations; sales distinguishes desired outcome (faster close) from replaced behavior (Excel, Slack, email, and shared-inbox chasing); product EVIDENCE reports 2.7-times higher weekly activity where local operations span at least three countries; later evidence shows a median 31% reduction in unowned open items in the last three close days, while total close duration does not always fall.
- **Why they matter:** the existing claim misstates product scope, comparison, ideal customer, and proof.
- **Facilitator/system handling:** separate outcome from replaced workflow; identify "every item" as an apparent wording dispute rather than a substantive one; integrate the narrower "cross-functional exception" scope; prefer distributed finance teams over a simple company-size label.
- **Convergence point:** after the comparison and customer analysis, before final copy is assembled.
- **Integration proposal:** describe a month-end exception-resolution workspace that keeps a blocked item, missing proof, owner, and decision in one traceable case; do not claim to replace communication channels or the ERP.
- **Human ratification:** Áron integrates the changes; Réka approves a 30-day website and ten-sales-call test.
- **Next Front:** "LedgerNest is a month-end exception resolution workspace for distributed finance teams. It keeps every blocked item, missing proof, owner and decision in one traceable case—so controllers know what is holding up close before the final days turn into a chase."
- **Remaining divergence:** the broader "finance orchestration" category remains a future option rather than the opening claim.
- **Protocol lesson:** the protocol should distinguish genuine disagreement from different professional wording and distinguish the promised outcome from the workflow actually replaced.

25.3 Research synthesis / CareLoop

- **Situation:** CareLoop is interpreting appointment research and deciding what to test about no-shows and late cancellations.
- **Initial Front:** patients mainly miss or cancel appointments late because they forget them, and additional reminders may reduce the problem; this is the starting hypothesis, not a ratified conclusion.
- **Relevant Statement:** a narrow observation about late cancellation is at risk of becoming a general explanation for all no-shows.
- **Interventions:** research and domain participants use SCOPE and CLARIFY to distinguish late cancellation from no-show and to note that different causes require different interventions; MEASUREMENT LIMIT challenges overgeneralization; an OPEN QUESTION asks whether expectation-setting and an easier uncertainty signal should be tested together.
- **Why they matter:** a real but bounded insight would otherwise propagate into an unsupported universal product thesis.
- **Facilitator/system handling:** identify scope mismatch between the presenter's inference and the evidence, and split one narrative into two connected but separately measurable hypotheses.
- **Convergence point:** after the interpretation block and before prototype design.
- **Integration proposal:** for some first-appointment patients, late cancellation is reinforced by two frictions: expectations become clear too late, and there is no convenient, consequence-free way to signal uncertainty. Do not claim this explains every no-show.
- **Human ratification:** Eszter integrates the scoped interpretation; Tamás approves a four-week test.
- **Next Front:** test an expectation card and a structured "may not be able to attend" path for first ultrasound and dental appointments, measuring earlier signals and successful rebooking separately, with receptionist workload as a guardrail.
- **Remaining divergence:** the wider causes of no-shows remain unresolved; the small two-clinic context cannot support a broad statistical uplift claim.
- **Protocol lesson:** a useful synthesis protects the boundary between evidence, interpretation, and product hypothesis instead of merely producing a smoother summary.

25.4 Investor narrative / GridTrace

- **Situation:** GridTrace is forming an ambitious but supportable narrative about industrial energy savings.
- **Initial Front:** "GridTrace can use AI to reduce industrial energy consumption by up to 20%, without hardware installation," presented as the startup's working hypothesis rather than a ratified fact.
- **Relevant Statement:** the company can present a general 20% savings opportunity and a broad industrial platform story.
- **Interventions:** technical and finance participants CLARIFY that identified savings are not the same as realized savings, human verification remains necessary, and technical reuse is not equivalent to gross margin; SCOPE separates a food-plant wedge from cross-industry expansion; EVIDENCE reports that three pilots shared 70% of measurement configuration and 60% of recommendations came from the same pattern library, while onsite validation remains service-intensive.
- **Why they matter:** later traction, scalability, "why now," and market claims depend on these distinctions.
- **Facilitator/system handling:** separate wedge from expansion; separate identified savings, implemented measures, and results still under measurement; require onsite-hours and gross-margin trends rather than treating technical reuse as commercial proof.
- **Convergence point:** before finalizing the deck narrative and headline.
- **Integration proposal:** GridTrace turns hidden industrial energy loss into a prioritized, verifiable savings plan; begin with mid-sized food plants' compressed-air and process-heat losses, while presenting broader use as conditional expansion.
- **Human ratification:** the team removes the 20% headline and approves separate presentation of identified 8–14% opportunities, implemented actions, and still-measured results.
- **Next Front:** a credible purpose, problem and why-now statement; a specific food-plant wedge; no guaranteed 20% saving or universally hardware-free deployment; scalability evidenced by declining expert effort, reusable configurations and patterns, and deployment gross-margin trend.
- **Remaining divergence:** cross-industry expansion remains a hypothesis; predictive maintenance and carbon accounting are parked outside the first five slides.
- **Protocol lesson:** corrections should attach before downstream narrative hardens, preserving ambition while exposing its evidential boundaries.

25.5 Product review / LoopCart

- **Situation:** LoopCart reviews poor activation in its returns-policy builder.
- **Initial Front:** users drop because the builder is complex; a five-step wizard with three preconfigured templates will solve it.
- **Relevant Statement:** interface complexity is the primary activation barrier.
- **Interventions:** support CHALLENGE says merchants often have not agreed the policy across legal, finance, and warehouse teams; product data EVIDENCE says 38% return to the same draft on multiple days; an OPEN QUESTION proposes treating shared policy approval as the first task; engineering raises a 9% EU Shopify re-authentication error; analytics explains this is a MEASUREMENT LIMIT even if it is outside the product concept; a cosmetic illustration comment is low-significance.
- **Why they matter:** the first cluster changes the problem definition, the authentication error confounds measurement, and the illustration comment does neither.
- **Facilitator/system handling:** split interface complexity from organizational decision uncertainty; constrain collaboration to a validation prototype rather than a workflow engine; park the re-auth fix as delivery while excluding failed connections from the cohort; mark the illustration comment NOT ELEVATED. When the system falsely attributes a comment to "Lili," the Facilitator rejects the attribution and blocks its use pending a verified author or source.
- **Convergence point:** before detailing the wizard solution, and again before adopting the measurement plan.
- **Integration proposal:** use a template-based draft plus a restricted review link through which a reviewer can approve or comment on three critical choices; measure draft creation, reviewer invite, review completion, and live policy.
- **Human ratification:** Hanna revises the thesis; Péter approves a two-week prototype plus four-week trial test, including explicit thresholds from the simulation.
- **Next Front:** the first value moment is not completing a wizard but creating, discussing, and approving a workable returns policy. A fuller collaboration flow is considered if at least 25% invite a reviewer and comparable trials with review show at least a 15-percentage-point higher publish rate.
- **Remaining divergence:** whether collaboration is broadly needed remains testable; a full role, permission, and workflow engine is excluded; the attribution error contributes no fact.
- **Protocol lesson:** significance handling must elevate organizational insight, account for a parked confounder, retain a low-impact note without surfacing it, and prevent unverified AI attribution from becoming authority.

25.6 Founder and board session / SignalPort

- **Situation:** SignalPort must explain weaker growth after moving from regional 3PLs toward strategic shippers.
- **Initial Front:** the strategy is right and the gap is primarily inconsistent enterprise sales execution.
- **Relevant Statement:** falling win rate is evidence of an execution problem.
- **Interventions:** a board member CHALLENGE notes lower buyer urgency and usage; sales CLARIFY identifies mixed ICPs; customer success EVIDENCE contrasts 22% weekly active case-owner share at large customers with 61% at 3PLs; the CTO warns user activity is not directly comparable but says the share of handled exceptions is also lower; finance later contrasts estimated 31-month versus 14-month CAC payback; other evidence shows regional 3PL new ARR grew 38% in 12 months despite reduced marketing.
- **Why they matter:** the comments challenge the meeting's central causal premise and any immediate personnel or sales-process conclusion.
- **Facilitator/system handling:** stop before sales actions; separate selling badly to a good ICP from selling well to a weak ICP; reject the false binaries of execution versus ICP and enterprise versus niche; park the final total-addressable-market and fundraising-narrative discussions.
- **Convergence point:** after the CEO's opening diagnosis and before personnel or broad campaign decisions; again when setting the test and authority.
- **Integration proposal:** core product-market fit remains strong among high-exception-frequency regional 3PLs; strategic shippers are an unproven expansion where buyer-user distance, lower workflow penetration, and incomplete qualification interact with execution weakness.
- **Human ratification:** András accepts the combined explanation; the board approves a six-week proof framework, spending cap, and thresholds while leaving tactical ownership with management.
- **Next Front:** pause the broad enterprise campaign; run at most eight strictly qualified shipper opportunities, requiring at least 200 relevant weekly exceptions, an operational owner, a live workflow within 14 days, at least 60% of in-scope cases handled, and two active owners. Continue the narrow motion if at least three opportunities reach usage thresholds and two commit to commercial expansion; otherwise pause the dedicated motion for at least two quarters. Return two SDRs to the 3PL core.
- **Remaining divergence:** enterprise potential and the larger vision remain open; neither "return to small customers" nor "just sell enterprise better" is ratified.
- **Protocol lesson:** a Next Front can encode a conditional, falsifiable strategy and distinct board/management authority without manufacturing consensus about the long-term market.

26. Relationship to Existing Approaches and Prior Art

WaveFront assembles capabilities whose components have substantial prior art. Google Wave combined communication and collaborative document work in a real-time space;^[1] the Apache project is now retired.^[2] Research on parallel meeting chat describes both participation benefits and attention, overload, and reintegration problems.^[3] MeetScript studied transcript-based interaction for active participation.^[4] Wikum+ interleaved discussion and summarization in online collaboration.^[5] Grain provides comments on live transcript moments.^[6] Microsoft Teams Facilitator provides real-time AI-supported meeting notes.^[7] Room Clarity describes a decision layer for Zoom;^[8] Arbiter describes real-time conversation intelligence;^[9] and MeetMap explored real-time collaborative dialogue mapping.^[10]

Accordingly, this specification does **not** claim that WaveFront invented real-time collaboration, transcript annotation, AI meeting summaries, shared meeting notes, or decision extraction. It also does not claim that any single component is unique.

The conceptual distinction asserted here is the complete reintegration and state-transition protocol:

anchored participant intervention → semantic/dependency interpretation → significance assessment → natural convergence → integration proposal → human ratification → versioned shared state

That distinction is descriptive, not a patent, trademark, or legal novelty claim. The cited sources document relevant antecedents and neighboring capabilities; they do not establish legal freedom to operate.

27. Known Risks and Failure Modes

- **Dual attention / cognitive overload:** reading a live Front and intervention layer can reproduce chat's split-attention cost. Mitigation: concise Statements, minimal core actions, notification restraint, and convergence rather than constant surfacing.
- **AI misunderstanding:** irony, scope, attribution, and dependencies may be misread. Mitigation: explainable proposals, VERIFY status, visible authorship, and mandatory human ratification.
- **Premature consensus:** a polished central text can appear more settled than it is. Mitigation: explicit provisional status and first-class Divergence.
- **Facilitator power:** selection and timing can suppress input. Mitigation: visible processing states, retrievable non-elevated items, bounded authority, and contestable proposals.
- **Monitoring perception:** provenance can feel like evaluation. Mitigation: explicit purpose, access and retention controls, data minimization, and exclusion of performance scoring.
- **Platform dependency:** deep coupling to one meeting vendor can impair portability or access. Mitigation: treat WaveFront as an augmenting protocol with exportable Fronts and Ledger history.
- **Easy feature copying:** transcript comments and summaries can be replicated. Protocol integrity therefore depends on the complete governed loop, not possession of a button or model feature.

- **Participant friction:** anchoring and labeling can be too costly. Mitigation: the four-action core, contextual extended types, guest-friendly access, and measuring repeat voluntary use.
- **Excessive facilitator load:** a queue of interventions may overwhelm the meeting lead. Mitigation: clustering, significance and dependency assistance, controlled Convergence Points, and delegation—without transferring authority to AI.

No mitigation eliminates the need for validation in real meetings. A technically functional interface can still fail socially if participants revert to speech or chat, treat the Front as prematurely official, or do not voluntarily reuse the protocol.

28. Versioning

WaveFront specifications use explicit version identities. A version-specific document is immutable after publication: corrections or changes are issued in a subsequent version or separately labeled errata rather than silently editing history. A future version may supersede Public Draft 0.9 but does not erase it.

“Public Draft” means the text is available for public review and implementation learning but remains subject to substantive revision. It does not reduce the force of its normative language for an implementation claiming conformance to this particular draft. The publication date remains **2026-09-23** in this working content master and must not be inferred from repository or source-file dates.

29. Glossary

- **Convergence Point:** a natural breakpoint at which significant interventions return to the main reasoning for possible integration; it may occur earlier when dependent reasoning or an imminent decision requires it.
- **Current Front:** the currently active shared working state, distinct from transcript history.
- **Divergence:** a material unresolved difference preserved with its context rather than hidden as consensus.
- **Integration Proposal:** a noncanonical proposed transformation of the Current Front, with reasons and provenance.
- **Intervention:** participant input anchored to a Statement or logical unit that may affect shared interpretation.
- **Next Front:** the human-ratified, versioned state that succeeds the Current Front.
- **NOT ELEVATED:** a retained Intervention not proposed for the main Front at the present time; not deletion.
- **Propagation Ledger:** the provenance record linking original Statement, Intervention, Evidence, proposal, Ratification, result, and remaining Divergence.
- **Ratification:** an authorized human decision on an Integration Proposal.
- **Statement:** the smallest logical unit in a Front that can be meaningfully anchored and may represent fact, assumption, interpretation, conclusion, proposal, question, or decision.
- **WaveFront System:** the software and AI assistance that represents and proposes but cannot ratify.

30. References

References cited by this specification are listed below.

1. Google Developers Blog. "Hello World, Meet Google Wave" (2009-05-28). <https://developers.googleblog.com/en/hello-world-meet-google-wave/>
 2. Apache Software Foundation. "Wave Incubation Status" (project retired 2018-01-15). <https://incubator.apache.org/projects/wave.html>
 3. Sarkar et al., Microsoft Research. "The Promise and Peril of Parallel Chat in Video Meetings for Work" (CHI 2021). <https://www.microsoft.com/en-us/research/publication/the-promise-and-peril-of-parallel-chat-in-video-meetings-for-work/>
 4. Chen et al. "MeetScript: Designing Transcript-based Interactions to Support Active Participation in Group Video Meetings" (2023). <https://arxiv.org/abs/2309.12115>
 5. Tian, Zhang, and Karger. "A System for Interleaving Discussion and Summarization in Online Collaboration (Wikum+)" (2020). <https://arxiv.org/abs/2009.07446>
 6. Grain. "Live meeting experience — inline comments on transcript moments." <https://support.grain.com/en/articles/14042480-live-meeting-experience>
 7. Microsoft Support. "Facilitator in Microsoft Teams meetings." <https://support.microsoft.com/en-us/teams/copilot/facilitator-in-microsoft-teams-meetings>
 8. Room Clarity. "The decision layer for Zoom." <https://roomclarity.com/>
 9. Arbiter. "Real-time conversation intelligence." <https://www.witharbiter.com/>
 10. University of Michigan Deep Blue. "MeetMap: real-time collaborative dialogue mapping." <https://deepblue.lib.umich.edu/items/cb1be5a5-699b-4a48-a454-c3a77ad0ece1>
 11. OpenStax. "Huygens's Principle," *University Physics Volume 3*. <https://openstax.org/books/university-physics-volume-3/pages/1-6-huygenss-principle>
 12. Bradner, S. "Key words for use in RFCs to Indicate Requirement Levels." RFC 2119 (1997). <https://www.rfc-editor.org/rfc/rfc2119>
 13. Leiba, B. "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words." RFC 8174 (2017). <https://www.rfc-editor.org/rfc/rfc8174>
-

Publication Information

Feedback and publication inquiries: protocol@wavefrontprotocol.com

Errata: <https://wavefrontprotocol.com/errata/>

Provenance: <https://wavefrontprotocol.com/provenance/>

Rights: <https://wavefrontprotocol.com/rights/>

Cite this specification

Németh, Attila. *WaveFront Protocol*. Public Draft 0.9.

2026-09-23.

<https://wavefrontprotocol.com/protocol/0.9/>