



A N O I N T E D E D U
GOVERNANCE SCIENCE REPOSITORY

WHITESIDE RESEARCH GOVERNANCE SERIES
Presentation Script · WGRS-WP-2026-001-ES

Executive Summary

The Whiteside Literature Review Standard (WLRS) v2.0

Narration script for the sixteen-slide presentation

Jermaine E. Whiteside
Doctoral Candidate, Liberty University · ORCID 0009-0004-7136-6859

DELIVERY SCRIPT — COMPANION ASSET

16 slides · approx. 21 minutes · Version 1.0
Published by AnointedEDU

Delivery Notes

This script accompanies the sixteen-slide executive summary presentation. The same text is stored in the speaker notes of the presentation file, so the two cannot diverge.

SLIDES	16
WORD COUNT	2797 words
ESTIMATED DURATION	approximately 21 minutes at 135 words per minute, excluding questions
REGISTER	Formal spoken academic register. Written to be delivered aloud, not read silently.
COMPANION FILE	WGRS-WP-2026-001-ES_Presentation_Narrated.pptx

STANDING GUIDANCE

- The status of the work is stated on the opening slide and again on the limitations slide. Neither should be abbreviated under time pressure; if the presentation must be shortened, reduce the innovation slides rather than the qualifications.
- The canonical contribution statement on slide eleven is fixed wording. Read it as printed rather than paraphrasing it.
- No commercial product is named anywhere in this script. Where citation software is discussed it is described generically, consistent with the tool-independence principle the standard sets out.
- Timings shown against each slide are cumulative and assume an unhurried delivery. They exclude pauses for questions.

A NOTE ON TONE

The script concedes the limits of the work early and repeatedly. That is deliberate. A standard concerning evidence integrity is best presented by a speaker who names its gaps before an audience finds them.

Narration Script

Slide numbers correspond to the presentation file. Cues are stage directions and are not spoken.

SLIDE 01 Title

137 w · 1:01

Cue. *Title slide holds. Pause before speaking.*

Good morning. My name is Jermaine Whiteside. I am a doctoral candidate at Liberty University, and I will be presenting a governance standard for scholarly evidence: the Whiteside Literature Review Standard, version two point zero.

Two qualifications before I begin, so that the status of this work is clear from the outset.

First, this is an independently published working paper. It has not undergone peer review, and I will be specific later regarding what it does and does not demonstrate.

Second, what follows is a proposal concerning how research evidence ought to be governed. It is not a validated methodology, and I would ask that it be assessed as a proposal rather than as a result.

The problem it addresses, however, is one I expect will be familiar from your own research programs. I will begin there.

SLIDE 02 The publication behind this summary

138 w · 2:02

Cue. *Advance on "three documents are relevant".*

A brief note on the structure behind this presentation, since the series is deliberately layered.

Three documents are relevant. The standard itself, WLRS version two point zero. The constitution — the full working paper, research code W-G-R-S, working paper, twenty twenty-six, zero zero one — which sets out the standard in detail. And the executive summary, the companion document from which this presentation is drawn.

The status of the work is stated plainly on the slide. It is independently published. It has been circulated for scholarly comment. It has not been peer reviewed, and I do not present it as though it had been. Every claim I make this morning should be read in that light.

The persistent identifiers appear at the foot of the slide, and I will return to them at the close with scannable codes.

SLIDE 03

Citation managers organise. They do not govern.

145 w · 3:07

Cue. *Allow the two-panel contrast at right to settle before the closing sentence.*

Consider the situation to which the standard responds.

Any active research program now involves an extensive body of material: journal articles, books, government reports, datasets, the researcher's own annotations, increasingly the record of AI-assisted conversations, and successive manuscript drafts. Hundreds of items. Frequently thousands.

Citation management software handles the bibliographic dimension of that material competently. Author, title, year, digital object identifier — sorted, tagged, and retrievable.

What such systems do not provide is governance. There exists no defined mechanism for preserving the researcher's appraisal of a source, for evaluating it against a consistent standard, or for reusing that appraisal systematically across projects and across years.

Organisation and governance are not equivalent. Organisation records where a source is located. Governance records what was concluded about it, upon what evidence, and whether that conclusion remains current.

That distinction defines the gap the standard is intended to fill.

SLIDE 04

Six persistent governance challenges

177 w · 4:25

Cue. *Six cards. Do not read them; narrate across them.*

Six recurring failures. I expect several will be familiar.

Rediscovery. A foundational source is located, worked through carefully, and then located again two years later — and the work is repeated, because nothing preserved the first appraisal.

Fragmentation. The review is distributed across a citation manager, a notes application, a folder of documents, and institutional storage. No single location holds it.

Disconnection. Annotations, and increasingly AI-assisted dialogue concerning a source, drift from the source itself until the reading that produced a given claim can no longer be reconstructed.

Attrition. Intellectual work is lost between projects. The dissertation concludes, the articles commence, and the accumulated capital does not transfer.

Shallow memory. The library preserves references. It does not preserve the reasoning.

And ungoverned artificial intelligence. This is the most recent of the six and, in my assessment, the most consequential. AI-assisted research is now routine, and in most workflows there is neither traceability nor a clear line of human accountability.

The cumulative effect is reduced efficiency, duplicated effort, and scholarship whose value decays more rapidly than it should.

SLIDE 05 The distinction in one line

132 w · 5:24

Cue. *Statement slide. Slow the delivery. Pause after the second line.*

If a single sentence is to be retained from this presentation, I would propose this one.

Current literature review methods emphasise finding sources. The standard instead emphasises governing them.

The distinction is not cosmetic. It alters what a source is understood to be.

Under traditional practice a source is consumed. It is read, quoted, cited, and then effectively retired; its value expended upon a single manuscript.

Under this standard a source is an asset. It is processed once, to a defined standard, and thereafter supports the dissertation, the journal article, the conference paper, the monograph, the teaching, and the public scholarship without being reconstructed on each occasion.

The same reading. A different unit of value. The literature review ceases to be a task that is completed and becomes infrastructure that is held.

SLIDE 06 What governed evidence offers an institution

166 w · 6:38

Cue. *Draw attention to the conditional verbs before reading the five rows.*

Why should an institution take an interest in this, rather than leaving the matter to individual researchers?

Five reasons. I would draw attention to the verbs, which are deliberately conditional. These are structured possibilities, not measured outcomes. Nothing on this slide has yet been demonstrated empirically.

Research integrity. Where every claim traces to a verified quotation upon an identified page, the claim is more readily defended under challenge.

Doctoral organisation. A considerable proportion of doctoral delay is not intellectual difficulty but reconstruction — candidates rebuilding literature work already performed. Structured governance may reduce that burden.

Faculty productivity. A source processed once and carried across four projects is a materially different asset from one processed four times.

AI governance. This is the consideration I would press with any research office. Defined verification requirements are the mechanism by which human accountability is preserved for work an AI system produced.

Institutional memory. When a researcher departs, governed evidence assets remain legible to those who follow. Ungoverned notes do not.

SLIDE 07 Four principles govern the standard

190 w · 8:02

Cue. *Four cards. Close on the tool-independence band at the foot.*

The standard rests upon four principles. They are deliberately few, since a constitution that cannot be recalled governs nothing.

First, permanent intellectual assets. A source is durable rather than disposable. Once properly processed it should serve the dissertation, the journal article, the conference presentation, the monograph, the teaching, and the public scholarship without reconstruction.

Second, evidence before output. Knowledge capture precedes publication. One does not draft first and seek support afterwards. Outputs emerge from a governed evidence base, which is to say that every claim traces to a verified quotation.

Third, structured reuse. A single evidence package — the quotations, the connections, the critical appraisal, the derivative outputs — serves future work without the recreation of notes.

Fourth, living scholarship. A source is not a static archival entry. As interpretation develops, the note is updated through what the standard terms a Re-engagement Log, which preserves the original reading alongside the revised one. The development of the researcher's own thinking remains visible.

Beneath all four lies tool independence. No field name refers to a particular product. Should the software in use today be obsolete within five years, the governance survives it.

SLIDE 08 The Foundational Gate

201 w · 9:32

Cue. *Read the override question from the navy panel verbatim.*

The first of three innovations. It answers an evident objection.

If governing a source properly requires substantial effort — and it does — then applying that effort to everything read is not workable. The standard therefore begins by determining what merits it.

The first stage is an objective score. Eight criteria. Does the source support a major theoretical construct? Is it cited three or more times within a major work? Does it bear upon multiple projects? Does it shape a conceptual framework? Four criteria satisfied, and the source becomes a candidate for Tier A — full processing.

I do not, however, consider a checklist a suitable arbiter of scholarly significance. A second stage therefore follows.

The constitutional override is a single question. If this source disappeared tomorrow, would one or more of my major arguments materially weaken?

An affirmative answer classifies the source as Tier A irrespective of its objective score.

That provision is deliberate. The mechanism exists to support judgment, not to displace it. A source may fail the score and nonetheless prove load-bearing for an argument, and the researcher is the party positioned to know it.

Remaining sources receive lighter processing, and may be promoted should their importance grow.

SLIDE 09 The Citation Hub Architecture

204 w · 11:02

Cue. *Six groups. The stored-prompt provision at the foot is the substantive point.*

The second innovation concerns the structure into which a Tier A source is processed.

Thirteen notes, each serving a distinct governance function. I have grouped them into six functions here so that the logic is visible rather than the list memorised.

Evidence capture — the summary with its evidence register, and the key quotations.

Research integration — the connections to the wider research program, to specific dissertation chapters, and to related scholars.

Scholarly interrogation — the question bank one would face at a defence or from a reviewer, the critical appraisal, and the record of AI-assisted dialogue.

Knowledge translation — podcast and teaching applications, through which scholarship reaches those who will never read the journal article.

Future development — research ideas and identified gaps.

And traceability — the citation map linking quotations to concepts and to outputs.

The provision at the foot of the slide is, I would suggest, the most consequential. Every note from C point zero seven onward is stored together with the prompt that generated it.

An AI-assisted asset is therefore not opaque. The prompt may be rerun upon a different platform and the outputs compared. Without the stored prompt, AI-assisted scholarship is merely convenient. With it, the work becomes auditable.

SLIDE 10 The Research Maturity Model

194 w · 12:28

Cue. *Direct attention to the permitted-outputs column, then to the highlighted L3 row.*

The third innovation is a maturity model, and its purpose is restraint.

Seven levels, L zero through L six. A source is acquired, verified, core processed, integrated, synthesised into a manuscript, translated into teaching or public scholarship, and eventually re-engaged following publication.

The significant column is the one at the right. Each level carries permitted outputs, and these are deliberately restrictive at the lower levels.

At L zero and L one — acquired and verified — the permitted output is none. The source may not yet be cited. It has been catalogued; it has not been processed.

At L two the source is available for internal drafting only.

Only at L three, integrated, may it enter dissertation or journal writing.

This constrains a habit common to most of us: a promising abstract is encountered in the morning and appears within a draft the same afternoon. The maturity model renders that a breach of the standard rather than a matter of discretion.

The chapter processed under this framework reached L three. It is prepared for dissertation and journal use, and I shall be precise shortly regarding what remains outstanding before that is fully the case.

SLIDE 11 Canonical Contribution Statement

162 w · 13:40

Cue. *Statement slide. Read the canonical paragraph without paraphrase.*

This paragraph constitutes the canonical contribution statement. It appears without alteration across the website, the publication, the citation package, and the accompanying video, so that the claim does not drift according to where it is encountered.

The principal contribution is the reconceptualisation of foundational evidence as a governed, reusable, version-controlled intellectual asset rather than a disposable citation tied to a single manuscript.

I wish to be exact regarding what is and is not being claimed.

This is not a further literature review technique. Many exist, and several are sound. It does not compete with them upon the methods of searching, screening, or synthesis.

What it proposes is a layer beneath those techniques: constitutional principles governing how evidence is acquired, evaluated, preserved, traced, and reused across the whole lifecycle of inquiry.

Whether that layer justifies its overhead is an empirical question, and it has not yet been answered. The contribution claimed here is conceptual — a different account of what a source is.

SLIDE 12 A five-layer governance architecture

184 w · 15:02

Cue. *State the qualification before the audience infers otherwise from the diagram.*

The standard is intended as the first of five forming a governance ecosystem.

Knowledge creation comprises this standard together with a protocol for AI-assisted source interrogation: acquisition, processing, interrogation.

Knowledge quality would be governed by an audit standard addressing citation integrity, verification, and counterargument analysis.

Knowledge control would address version history, release management, and archival preservation.

Knowledge stewardship would address asset lifecycle, maturity, derivative outputs, and eventual retirement.

An integration paper would describe how the standards interlock.

The qualification is important, and I prefer to state it rather than permit the diagram to imply otherwise. Only the layer marked in gold exists. This standard is published. The remaining four

are manuscripts in preparation. They have not been released, they have not been reviewed, and anything I say regarding their operation is design intent rather than delivered capability.

I present the architecture because it explains why the standard is scoped as narrowly as it is. It deliberately does not attempt the auditing or the versioning, since those belong to standards that do not yet exist.

The lower four rows should be read as a roadmap.

SLIDE 13 Where this is intended to be used

160 w · 16:13

Cue. *Close with the direct address to research offices.*

Where is the framework intended to be applied, and by whom?

The environments appear at the left. Doctoral education is the evident case, and the one in which I have tested it. The same problem, however, arises within faculty research programs, institutional research offices, policy organisations and think tanks, government research agencies, and for independent scholars possessing no institutional memory beyond their own.

The audiences appear at the right. Doctoral researchers structuring a review and seeking to build capacity rather than merely complete a chapter. Faculty directing programs and supervising candidates. Research directors considering governance at institutional scale. Librarians developing evidence stewardship infrastructure — a conversation I would particularly welcome. And researchers in AI governance, since the traceability problem in AI-assisted scholarship will not resolve itself.

I would single out one group. If you direct a research office and are being asked how your institution maintains human accountability for AI-assisted work, this framework constitutes one concrete answer available for evaluation.

SLIDE 14 What this does not yet demonstrate

214 w · 17:48

Cue. *The central slide of the presentation. Do not hurry it.*

This is the slide I would ask you to weigh most heavily. I prefer to present it myself rather than have a reviewer discover it.

The framework has been demonstrated through precisely one reference implementation. A single chapter — chapter one of Thomas Groome's *Sharing Faith* — producing twenty-six intellectual assets. One case.

Four consequences follow, and I shall state each plainly.

The quotation verification remains in progress. The page numbers within that implementation derive from my reading notes and have not yet been checked against the original volume. Until that check is complete, those page references should not be treated as confirmed.

No independent inter-rater reliability assessment has been conducted. No other researcher has applied the framework to the same source in order to establish whether comparable results would follow.

Processing efficiency has never been compared against simpler note-taking. I believe the overhead is recovered across a research program. I have not demonstrated it.

And the four companion standards remain proposals rather than completed work.

The honest summary is therefore as follows. This is a single-implementation pilot. It is not an operational or generalisable standard, and I would ask that it not be adopted as one.

A stress test against a second chapter is under way. That is the purpose of version 1.1.

SLIDE 15 The publication ecosystem

176 w · 19:07

Cue. *Three status columns. Name the conditional column explicitly.*

One of the governance principles within the standard is that the extent of what exists should not be overstated. The ecosystem is accordingly presented in three columns reflecting actual status.

Available now: the full constitutional publication, this executive summary, an author overview video, a cinematic presentation, two podcast discussions, this presentation, the publication website, and the persistent identifier records — ORCID, Wikidata, and Google Scholar.

In development: an AI Companion, which will publish the prompts, the revision history, and the editorial decisions behind the work — the standard applied to itself. A scholarly citation package in each major format. An evidence package containing the source verification and quotation audit. And publication-level analytics.

Future or conditional: a digital object identifier, should one be obtained; and a repository record for this publication, following acceptance, if applicable. That final item is deliberately conditional. This paper does not at present hold such a record, and I will not imply otherwise.

I would rather present a column marked conditional than have it later discovered that something described as available was not.

SLIDE 16 **Where to find the work**

217 w · 20:43

Cue. *Closing slide. Hold the codes on screen through questions.*

Permit me to close where I began.

The claim is modest; the ambition is not. The claim is that one chapter has been processed under this framework and the workflow held. The ambition is that literature review should cease to be something completed and become something held — infrastructure that compounds across a career rather than dissipating at the conclusion of each project.

Between the two lies a considerable body of work not yet undertaken. A second implementation. Independent assessment. A comparison of efficiency. Four companion standards.

What I am requesting, therefore, is not adoption but scrutiny. If you supervise doctoral candidates, I would value your account of where this would fail under real conditions. If you direct a library or a research office, I would value knowing what would need to be true before you would pilot it. And should any of you be willing to apply the framework to a source within your own field and report the outcome, that would be the most useful contribution available from this room. A second reference implementation originating outside my own work is worth considerably more than anything I can produce alone.

The codes displayed will resolve to the publication, the collection, and the scholarly profiles. All materials are openly available.

Thank you. I would welcome your questions.