
FAMIS WORKSHOP 2026

Building a Legally Sound AI Framework for ESE

Dr. M. Tricia Phillip-Magee

Director of Exceptional Student Education, Orange County Public Schools



Part I: The Adoption Surge

Mapping the Velocity of Classroom
Practice

National Adoption Trends: 2024–2026

User Demographic	2024 Baseline	2026 Current	Growth Delta
Teacher Adoption	48%	85%	+37 Percentage Points
Student Adoption	63%	86%	+23 Percentage Points

Sources: CDT "Mistrusted and Maligned" (2024) & Gallup/Walton Foundation (2026)

Part II: Compliance Mandates

Federal Deadlines and Digital Rights

The Compliance Bridge: April 2027



Firm Deadline

April 26, 2027 target for districts with 50,000+ residents to reach enterprise-wide accessibility.



WCAG 2.1 AA

Mandates strict adherence for all digital content, including outputs generated by AI tools.



Strict Liability

The burden of compliance shifts entirely to district infrastructure; "bad faith" is no longer a defense.

Beyond the Website: Classroom Duty



Digital Content Scope: AI-generated PDFs and handouts shared via LMS must be fully tagged for screen readers by the 2027 window.



Multimedia Mandate: Video content created via GenAI requires high-accuracy captions and audio descriptions.



Professional Duty: Staff must manually verify AI outputs for logical reading orders and Alternative Text (Alt-Text).

Part III: Sovereignty

The Privacy Redline for ESE

Data

Privacy & "Inference PII"

Redaction is not Anonymization.

The Inference Risk: AI can re-identify a student through "sparse data" points—disability category, grade level, and school—even without a name.

*"Example: Inputting a student's **disability category, grade level, and school** creates a contextual fingerprint identifiable to the local community."*

Legal Standard: §1002.222, F.S. (Florida Student Data Privacy Act)

Systemic Risks: Bias & Error



Hallucinations

AI may fabricate student progress data or legal citations, jeopardizing the defensibility of records.



Algorithmic Bias

Monitoring systems flag ESE students at higher rates due to non-typical behavior or speech patterns.



Genericism

"Boilerplate" AI text risks violating the IDEA mandate for highly individualized educational programming.



Part IV: Remediation

Clinical Judgment in the Age of
Automation

Case Study: The Bypass Trap

W.A. v. Clarksville–Montgomery Co. (2025)

District provided AI tools (ChatGPT/Grammarly) as home-based assistive technology for a student with dyslexia.

The Ruling: Federal court ruled a **denial of FAPE**. Technology allowed the student to bypass learning gaps rather than remediating them.

Citation: 127 F.4th 656 (6th Cir. 2025)



The Remediation Gap: FAPE vs. Efficiency

The Bypass Liability

AI efficiency becomes a legal liability if cognitive struggle is removed from the learning process.

Rule: Clinical judgment and instructional remediation cannot be automated.

The Instructional Standard

AI must differentiate the **delivery** of instruction, not substitute for the **mastery** of foundational skills.

Remediation is the core mandate of IDEA.

Part V: Infrastructure

Building the Enterprise

Shield

Reclaiming Sovereignty: The Vault

The Vault Principle

Managed Enterprise silos where data is contractually excluded from global model training.

- 📋 **Prompt Logging:** Essential audit trail for legal defensibility.
- 🛡️ **Managed Ecosystem:** Moving from Shadow IT to whitelisted silos.

"If we don't own the logs, we don't own the compliance."



The FTE–Compliance Connection

Boilerplate AI goals don't just violate **IDEA's individualization mandate** — they can **invalidate the matrix level that funds the seat.**

Legal Defensibility

IDEA requires IEP goals to be genuinely individualized. AI cannot substitute for clinical judgment.

=

FTE Accuracy

A generic goal can invalidate the matrix level — reducing the funding weight for that student's seat.

Operational Mandate: HITL



Human-in-the-Loop (HITL)

AI Augments; Professionals Authenticate.

0% AI Autonomy Redline:

Placement Decisions | Discipline Manifestations | Final Assessments

The PARTS Framework

P

Purpose
Goal
alignment.

A

Authorize
Permission.

R

Review
Audit bias.


T

Trust
Source validity.

S

Safety
PII shield.

The Legally Sound IEP Workflow



Same Goal, Two Outcomes

A dyslexic student needs help with a writing assignment — the same request, prompted two different ways.

✗ THE BYPASS PROMPT

“Write a five-paragraph essay on the causes of the Civil War for a 9th grader with dyslexia who is struggling to get started.”

What it produces:

- A complete, polished essay the student did not write
- Zero engagement with the actual skill gap — organizing ideas, sequencing paragraphs
- The student turns in work that isn’t theirs; no learning occurred

This is the Clarksville pattern: cognitive struggle removed entirely, not reduced.

✓ THE REMEDIATION PROMPT

“Act as an executive functioning coach for a 9th grader with dyslexia who struggles to start writing tasks. Break this essay prompt into a 3-step checklist the student can complete independently — do not write any of the content yourself.”

What it produces:

- A 3-step scaffold (e.g., “1. List 3 causes. 2. Pick one fact per cause. 3. Write one sentence per fact.”)
- The student still does 100% of the thinking and writing
- AI removes the executive function barrier, not the cognitive struggle

This is the 0% AI autonomy line in action — AI assists access, never substitutes mastery.

The difference isn’t the tool. It’s whether the student still has to do the work.

Before You Practice: Why This Is Safe

The Rule Has Not Changed. Today's Practice Follows It.

- 1. Every scenario in the toolkit is fictional.** “A student with an ED eligibility” is a category, not a person. No name, ID, or real record is ever entered.
- 2. You are practicing the Framework, not the data.** RICO is a structure for the prompt itself — it works identically whether the context is fictional or real.
- 3. The same Safety Shield applies back at your desk.** §1002.222 and FERPA govern what you put INTO a tool. A fictional category violates neither — and you will build the habit of writing it that way every time.

If it could identify a real student, it doesn't belong in the prompt — in the workshop or on Monday morning.

The RICO Framework

Role • Instructions • Context • Output

R

Role

Tell the AI who it is. A specific identity frames everything that follows.

"You are an experienced ESE compliance specialist..."

I

Instructions

State what you need before you explain the situation. Number the steps.

"List the steps before removing a student for 10+ days..."

C

Context

Load the situation. No names, IDs, or identifying details – ever.

"Student has ED eligibility, active BIP, 8 days removed..."

O

Output

Name the format. Checklist? Table? Word count?

"Numbered checklist, 10–14 steps. Flag legal deadlines in bold."

🕒 4-MINUTE PRACTICE

Use the RICO Framework

THE SCENARIO

A student with an ED eligibility has 8 cumulative days of removal.

Your principal just called you about a new incident.

Go.

Scan the QR code → open the toolkit • use RICO • write your prompt

We'll cold-call 2-3 people to share. What does your R-I-C-O look like?

Reminder: this is a fictional scenario. No real student is involved – that's what makes it safe to practice.



ESE AI Toolkit

PARTS • RICO • Prompts

Ethics: The Cost of Hallucinations

Professional Liability

Courts are setting a precedent that "unintentional misstatement" is no longer a valid legal defense.

- **Powhatan County School Bd. v. Skinger (2025):** Advocate sanctioned for AI-hallucinated citations in an ESE dispute.
- **Mata v. Avianca:** Federal fines for fake AI citations in legal proceedings.

Case Citation: No. 3:24-cv-874 (E.D. Va. 2025)



Let's Keep the Conversation Going



Dr. M. Tricia Phillip-Magee

marlenetricia.phillip-magee@ocps.net