

G. ERIC PENDLETON

Selected Work Samples

AI for Business & Transformation

Technical L&D Materials from Pharma, Banking, and IT

Boston, MA (US Eastern) · 617-895-7640 · ericpendleton@gmail.com

August 2026

A Note on Confidentiality & Materials Several of these samples were produced under contract for pharmaceutical and enterprise-software clients. Sensitive client identifiers have been redacted, and no proprietary code, trade secrets, or regulatory correspondence are included. Where a document is referenced but omitted, that exclusion is deliberate; full artifacts are available upon request subject to appropriate discretion.

Sourced from legacy projects dating back 10 to 15 years, these materials reference obsolete technologies and pose zero commercial risk to former clients. They are featured here because while software architectures change, the underlying craft does not: these samples reflect rigorous audience analysis, structural clarity, and the ability to translate complex engineering into precise user and administrator documentation.

What is in this portfolio, and why

These days, an L&D instructor is asked to hold two things at once: fluency with AI, and credibility with learning-and-development professionals who will quickly discern whether the person teaching them has ever actually owned a curriculum. These eight samples are organized around that pair.

Samples 1–5 are instructional-design and course-authorship work from twenty years inside corporate L&D. Samples 6–7 are the AI half: published trade writing and enterprise onboarding documentation produced at an AI/NLP software company in 2016–17. Sample 8 covers documentation at scale.

#	Sample	Year	What it evidences
1	Learning for Competence Original white paper, 16 pp.	2018	Sole-authored instructional model: competency-based learning, document taxonomy, LMS requirements, implementation roadmap, ROI measurement. Course-author thinking at system level.
2	Enterprise Quality Training: Train-the-Trainer vILT course, 60 slides	2018–19	Full course authored and facilitated: ADDIE, measurable objectives with Bloom's verbs, evidence-based adult learning, trainer skills, 11 case-study scenarios, learner assessment.
3	Author Training System-training course, 124 slides	2019	End-to-end software training for document authors — objectives, workflow modelling, job aids, appendices. Described only; not reproduced.
4	Training Needs Analysis Worksheet 10-step instrument	2019	Front-end analysis: business outcome first, trainable-vs-untrainable competencies, delivery-format selection, cost-benefit, evaluation plan.
5	Training Development Guide Reusable authoring template	2019	Competency-based course skeleton: prerequisites → objectives → demonstration → practice → assessment, with a signed assessment matrix.
6	"5 Steps to Unlocking the Value in Unstructured Big Data" Quirk's Media, published	2017	Translating AI/NLP concepts into a decision framework for non-technical business readers — the exact task this role is built around.
7	Enterprise AI platform documentation Luminoso — 3 guides	2017	Teaching business users to operate an NLP analytics platform: workflow-first structure, key-term glossaries, publish/share governance.
8	Documentation at scale Manuals and knowledge systems	2007–17	Multi-hundred-page user manuals and a 500+ page role-based knowledge system serving ~2,000 employees.

Samples 1–5 come from regulated pharmaceutical and biotech environments, where training content is audited. That constraint shaped the method on display and is directly relevant to the governance portion of any AI-adoption curriculum.


Learning for Competence

Original white paper · 16 pages · February 2018 · Sole author

What it is. A proposal arguing that most corporate training is a compliance exercise — reading procedures and checking a box — and setting out a competency-based alternative in enough operational detail that an organization could actually adopt it. It defines competency-based learning, separates *documents to perform* from *documents to train*, specifies what an LMS must do to support the model, gives a stage-by-stage implementation roadmap, and closes on how to measure return. It was written for a process-driven industry migrating in from pharmaceutical practice — that is, it argues a method across domains rather than describing one company's setup.

WHY IT MATTERS FOR AI FOR L&D

- It is the closest thing here to a course author's blueprint: a full model, argued from first principles, with figures, tables, an appendix of learning stages, and a roadmap someone else could execute.
- Its central distinction — training that produces a completion record versus training that produces demonstrated capability — is exactly the distinction AI-for-work courses tend to blur. AI can generate learning content very fast; it cannot certify that anyone can now do the job.
- Section 7 (LMS requirements) and section 9 (measuring return) are the parts an L&D audience will want AI applied to, and the parts where an instructor needs to be specific about what is safe to automate.



A comprehensive strategy for achieving outcomes and evaluating progress through large-scale training initiatives.

6. Learning Assignment Structure

The key to structuring learning assignments in a meaningful way is role-based. We need to make the link clear between the tasks being performed and what learning is required. Also, since associates are often required to be qualified to perform a large collection of tasks, it is also considered a best practice to have these assigned as roles, in a manner that matches the appropriate learning progression.

Any employee's list of training modules could easily be several pages long. Therefore, it would be impossible for employees to accurately prioritize what training is crucial to their qualification (i.e., critical knowledge they must have to carry out their daily work). The end result can be more confusion around training priorities.

Figure 5 depicts a training assignment structure that supports a competency-based learning model.




Figure 5. Training Assignment Structure to support a competency-based model.

Figure 6 provides guidance for the level of training required and type of learning materials to be developed for each assignment, which should be role-based.

Starting with the outermost ring, the 'Know About' area indicates training that an employee must be assigned but not necessarily know about in absolute detail. In essence, they must know the basics of their industry, understand the purpose, and be able to locate further information when needed.

The second ring 'Be Trained On', indicates learning that must be mastered in more detail – i.e., to be trained to a certain degree, to be able to understand and perform basic tasks, but not in quite as much detail as the final ring.

In the innermost 'Be Qualified On' ring, the employee is expected to be qualified on this training in absolute detail and be able to demonstrate complete mastery of the specific job task or function.



Figure 6. Considerations for types of training assignments

Please consult the APPENDIX for details about learning stages, corresponding principles, support materials, resources, and instructors appropriate for each stage.

Section 6, with original figures. Role-based assignment structure, and the three-ring distinction between what a learner must know about, be trained on, and be qualified on — the mechanism that keeps a curriculum from expanding into an unreadable list.

The line I would keep in an AI-era course: “Qualification is not being able to produce a report that shows that an associate completed training. Rather, qualification is demonstrating an ability to perform.”

Enterprise Quality Training: Train-the-Trainer

Virtual instructor-led course · 60 slides · 2018–19 · Global biopharmaceutical client · Authored and facilitated

What it is. A complete train-the-trainer course for staff who would go on to train others on regulated processes. Seven sections: ADDIE and learning objectives; adult learning theory; regulatory expectations; case studies drawn from real enforcement letters; trainer interpersonal skills; eleven difficult-training-situation scenarios; and hands-on trainer practice, closing with assessment against the course's own stated objectives.

It was built to be delivered live — group breakouts, paired practice, discussion — not narrated over slides.

WHY IT MATTERS FOR AI FOR L&D

- **Instructor.** The course is structured around learner production. Roughly a third of it is people doing something and being corrected.
- **Course author.** Objectives are declared at the front and re-used verbatim as assessment criteria at the back — a discipline most courses claim and few actually execute.
- It takes a position against a popular idea (learning styles) on the evidence. An AI-for-L&D course will need the same nerve about several fashionable claims currently made for AI.

ADDIE BROKEN DOWN

- Analyze – What will learners be able to do after training? (=LOs)
- Design – Create assessments, choose training delivery method(s), decide sequence
- Develop – Create course materials and assessments (use/update existing materials? new materials?), train trainers
- Implement – Prepare learners, distribute pre-work, ready environment, hand out training materials, roll out instruction
- Evaluate – Assess learning based on learning objectives

VIDEO

[client mark removed]

ADDIE presented as a working framework rather than a rigid sequence.

LEARNING OBJECTIVES

Training development cannot begin until clear Learning Objectives have been established! LO action verbs:

Knowledge	Comprehension	Application	Analysis	Synthesis	Evaluation
define	classify	apply	analyze	arrange	assess
identify	compile	calculate	calculate	assemble	compare
label	conclude	demonstrate	categorize	compose	critique
list	discuss	develop	classify	construct	decide
match	describe	interpret	criticize	design	determine
name	explain	locate	compare	develop	establish
recall	express	operate	contrast	diagnose	evaluate
recognize	give examples	perform	determine	manage	judge
record	identify	practice	differentiate	organize	justify
relate	interpret	predict	distinguish	plan	measure
repeat	recognize	present	examine	propose	rate
select	summarize	report	outline	relate	recommend
state	translate	use	test	summarize	

[client mark removed]

Measurable objectives, with a Bloom's-taxonomy verb table the learner uses in the practice exercise later in the session.

LEARNING STYLES DEBUNKED

Learning Styles are POPPYCOCK

- The self-esteem movement of the 1980s and 90s gave way to the myth that people are either visual, verbal, reading, auditory, tactile, taste or my "favorite," *kinesthetic* learners! YIKES! (On this topic, ever notice how everyone's a *visual* learner?)
- Domestic and international Psychological research has repeatedly found that people do not learn differently (at least, not how learning styles profess)
- Authors of these scientific studies argue that the use of learning-style tests and teaching tools is a wasteful use of resources.

While it is essential to utilize different learning methods and engage all our senses, it's about WHAT we're learning, not HOW we learn.

[client mark removed]

Taking an evidence-based position against learning styles — in a corporate room where the idea is popular.

OJT TRAINER – “IDEA” MODEL

As the trainer:

1. Demonstrate the task
2. Review key points
3. Demonstrate the task a second time
4. Perform simple parts of the task
5. Guide learner to perform the whole task

The trainee performs the whole task...

1. ...monitored by trainer
2. ...on their own (# of times?)

[client mark removed]

The on-the-job training model the course drives toward: demonstrate, review, demonstrate again, partial performance, whole performance — monitored, then unmonitored. Practice is sequenced, not assigned.

What I would carry into an AI curriculum. This structure is the answer to the most common failure of AI training: a demo, a burst of enthusiasm, and no behavior change. The equivalent for an L&D cohort is a live session where each participant runs a real piece of their own work through a tool, gets it checked against source, and finds the error themselves.

COURSE ARCHITECTURE AT A GLANCE

§	Section	Mode
1	ADDIE and learning objectives	Instruction + verb-table exercise
2	Adult learning theory: train for competence	Instruction + discussion
3	Regulatory expectations and warnings	Instruction
4	Enforcement case studies	Group breakout, 5 cases
5	Trainer interpersonal skills	Instruction + practice
6	Difficult training situations	Group breakout, 11 scenarios
7	Forms, trainer practice, learner assessment	Paired practice + assessment

Four of seven sections put the learner in production. The course closes by assessing against the five objectives declared on slide 3 — the same five, unchanged.




Author Training

System-training course · 124 slides · v1.2, April 2019 · Global biopharmaceutical client · Not reproduced

What it is. A step-by-step course teaching document authors to work inside an enterprise controlled-document system, from first login through final approval workflow. It opens with six measurable objectives, previews the whole document lifecycle as a workflow diagram before teaching any step, then walks search, creation, templates, revision, translation coordination, training determination, and the four approval workflows — closing with five appendices covering rejection reasons, withdrawal, document relationships, tracked changes, and workflow notifications.

WHY IT MATTERS FOR AI FOR L&D

- It is the largest single course in this portfolio and the most procedural — evidence of stamina, structure, and version discipline across a long build (v1.2, with a documented comment-and-revision history).
- The design choice worth noting: the entire workflow is shown *before* any individual step is taught, so learners always know where they are. That orientation-first pattern is what most AI tool training lacks.
- Teaching authors to author — deciding when a new document is warranted rather than a revision, who must review, what triggers a training requirement — is close cousin to the judgment an *AI for L&D* cohort needs about when to generate, when to reuse, and when to escalate.



Training Needs Analysis Worksheet

Reusable analysis instrument · 10 steps, 3 pages · 2019 · Unbranded

What it is. A worksheet that forces a requester to answer the business question before any training gets built. It starts from a quantified business outcome, works back to the behaviors that would produce it, then applies a filter most needs analyses skip: which of these competencies are actually *trainable*, and which should be hired for instead.

STRUCTURE

1. Determine business problem and quantify desired business outcome
2. Link desired business outcomes to employee behavior
3. Identify trainable competencies — “It is more effective for an employer to hire people who are already extroverts than to attempt to train introverts to be more outgoing.”
4. Evaluate competencies against a five-point criticality scale
5. Determine performance gaps — before vs. after — and analyze why they exist
6. Prioritize training needs — top five only; more than five means a second module
7. Determine delivery format, ranked by retention
8. Conduct a cost-benefit analysis, including lost productivity
9. Plan for training evaluation against four questions
10. Write no more than five learning objectives

WHY IT MATTERS FOR AI FOR L&D

- Step 3 is the transferable idea. The AI-era version is: which parts of this problem are a knowledge gap, which are a process gap, and which are neither — because a tool will not fix a broken process any more than training fixes a hiring mistake.
- Steps 8 and 9 give an L&D team the vocabulary to make an AI business case to a CFO, which is usually where AI enthusiasm inside the function goes to die.
- An instrument like this is also the clearest place to demonstrate AI genuinely earning its keep: synthesizing stakeholder interviews into a first-draft gap analysis in minutes, then having a human check every claim against what was actually said.

Training Needs Analysis (TNA)

Before completing the worksheet, please answer the following questions:

- What is the business problem?
- What is needed? Why?
- Where is it needed / who needs it?
- How will it be provided?
- How much will it cost?
- What will be the business impact?

Step 1: Determine Business Problem and Quantify Desired Business Outcome

Examples of ultimate business goals are:

- Improve _____ by ____%.
- Increase _____ by _____ (dollars, time frame).

Step 2: Link Desired Business Outcomes with Employee Behavior

Identify the desired critical competencies (behaviors and associated knowledge, skills, abilities, as well as any personal characteristics linked to desired business outcomes).

- What learners need to do
- What capabilities do they need to do each action/task
- What motivates them to do it

Critical = at least a 4.0 average on the five-point rating scale.

Step 3: Identify Trainable Competencies

Not every competency can be improved through training. For example, a sales job may require sales people to be outgoing and initiate conversations with total strangers. It is more effective, then, for an employer to hire people who are already extroverts than to attempt to train introverts to be more outgoing. Similarly, it may be more effective to hire people with specialized knowledge than to educate and train them.

Step 4: Evaluate Competencies

Decision makers should evaluate each critical competency from Step 2 and determine if each one is something they expect employees to possess prior to job entry. Taken together, this should provide decision makers with a list of critical competencies that are amenable to training. How will learners be evaluated?

- Competency evaluations and
- Tests or assessments.

Training Development Guide

Reusable authoring template · 6 pages · 2019 · Unbranded

What it is. The skeleton every course in the system was built on. It fixes the sequence — purpose, prerequisites, measurable objectives, background, materials, regulatory context — then routes the learner through a four-step progression that ends in a signed assessment matrix. Fill it in and you have a defensible course; skip a field and the gap is visible.

THE PROGRESSION

1. **Overview and background** — context the procedure itself does not and should not carry
2. **Demonstration** — the trainee observes the qualified trainer, documented to a minimum count
3. **Practice** — the trainee performs under direct supervision, documented to a minimum count
4. **Final assessment** — performance against the stated objectives, evaluated step by step and initialed at the time of performance

The assessment table pairs each task with an explicit criterion in two registers: *explain the significance of* (understanding) and *demonstrate per the procedure* (capability). Both must be satisfied.

WHY IT MATTERS FOR AI FOR L&D

- Templates are what make a course author's judgment reusable by other people. This is the artifact that turns a method into an institution.
- It is also the natural unit for an AI exercise: give a cohort a real procedure and this template, have AI produce a first draft, then have them find what the model invented, softened, or quietly omitted. That exercise teaches the tool and the verification habit in one pass.
- The signature column is the point where AI stops. Someone with a name attests that a person can perform. No model can hold that.

Training Development Guide			
TRAINING COURSE NAME: _____ SOP/PRC: _____			
PURPOSE The purpose of this Training Guide is to _____.			
PREREQUISITES The following courses/procedures need to be completed prior to training: _____.			
LEARNING OBJECTIVES Upon completion of this training, the trainee will be able to (example actions): • Install a _____. • Place a _____. • Determine _____. • Calculate _____. • Identify _____. • Perform _____ task according to xyz standards.			
Step 1: Overview and Background This training guide provides background and best practice options for _____ in order to _____. <i>Include relevant background, images, workflow steps, diagrams, etc. (anything not included in SOP).</i> Materials for Training The QT obtains the following in order to conduct the training: • Item 1 _____. • Item 2 _____. • Software login credential for _____. • Tool x _____. <i>Explain how items and tools will be used properly in training. Account for why these tools are important, how to maintain them, any safety precautions, location-specific details, etc.</i> The proper method for using item 1 is _____.			
Regulatory Expectations There are various regulatory expectations and requirements that are related to _____. According to the FDA's _____ documentation: <i>"E pluribus unum."</i>			
Step 2: Demonstration The trainee will observe the task in the presence of the Qualified Trainer and document the minimum required demonstrations as indicated below prior to moving to Step 3 (Practice).			
Step 3: Practice The trainee will practice the task under direct supervision of the Qualified Trainer and document the minimum required practice sessions as indicated below prior to moving to Step 4 (Final Assessment).			
Step 4: Final Assessment Based on the Learning Objectives, the trainee will perform/demonstrate the process under the direct supervision of the Qualified Trainer. The Qualified Trainer will evaluate the trainee against the assessment criteria and will approve each step at the time of performance.			
Step	Task	Assessment Criteria	Qualified Trainer Initials/Date
1	Install _____.	Explain the significance of _____.	
		Demonstrate _____ per the procedures in this training guide.	
2	Place a _____.	Explain the significance of placing _____ on _____.	
		Demonstrate the _____ installation per the procedures in this training guide.	
3	Determine _____.		
4	Calculate _____.		
5	Identify _____.		

“5 Steps to Unlocking the Value in Unstructured Big Data”

Quirk's Media · published 24 April 2017 · byline: Eric Pendleton, Project Training Manager, Luminoso

What it is. A published trade article translating natural-language-processing capability into a five-step decision framework for business readers with no technical background: define the business objective first; remove data silos; empower analysts with access and executive cover; analyze unstructured and structured data together; measure and track return.

Live at quirks.com/articles/5-steps-to-unlocking-the-value-in-unstructured-big-data. A companion piece ran in *Customer Experience Magazine* the same year; that publisher's link no longer resolves.

WHY IT MATTERS FOR AI FOR L&D

- This is the earliest dated proof of the specific skill this role hires for: taking real AI work and turning it into something a non-technical professional can act on. **It was published in 2017** — before the current wave, when the translation was harder because the audience had no priors at all.
- The structure is instructive and still holds: it leads with *business objective*, not with the technology, and it ends on measurement. That is the same spine an AI-adoption curriculum needs.
- Steps 2 and 3 are organizational, not technical — silos and permissions. The people who took AI seriously early learned quickly that the blocker is rarely the model. An L&D cohort planning adoption across their organization needs that lesson on day one.

“While companies are still becoming aware of the massive amounts of unstructured data they have, they are not necessarily aware of the risks they face or the untapped value they’re missing out on. But experience shows that they should make this a priority. For the average Fortune 1000 company, increasing data accessibility by just 10 percent results in more than \$65 million in additional net income.” -Eric Pendleton, Training Manager (Luminoso AI) 2017

[Read full article](#)



Enterprise AI platform documentation

Luminoso Technologies · 2016 · Dashboards & Widgets (10 pp.), Analytics Quickstart, Beta Quickstart

What it is. Onboarding and reference documentation teaching business users — product, sales, marketing, and research teams — to operate an NLP text-analytics platform. Not API documentation: guidance for people whose job was insight, not engineering.

The structure is workflow-first. Before explaining a single button, it establishes what must be true in the underlying project: data prepared and subsetting, a hypothesis stated, topics built and named deliberately, the analysis actually run. Only then does it teach the interface. It also carries an explicit key-terms section and a governance section on publishing, snapshots, and sharing permissions.

WHY IT MATTERS FOR AI FOR L&D

- It is direct evidence of teaching an AI product to a non-technical business audience — the closest structural analogue in this portfolio to what learners will be asked to do inside their own organizations.
- “What’s your hypothesis? What types of insights do you seek?” appears before any tooling. That question is the single most useful correction to the way most people currently approach a chat interface.
- The publish-and-share section is early governance thinking: who sees which view, who can edit, what a frozen snapshot means when the underlying data has since changed. Version and audience control are unsolved problems in AI-assisted L&D right now.

Luminoso Introduction

There are five easy steps for managing a project end-to-end in Luminoso Analytics — from uploading your source data and leveraging the Analytics tool’s features, to reporting.

However, the data investigation / insight discovery process of your analysis consists of three iterative A, B, C steps (which are outlined below and further discussed on pages 3 & 4), called the **Luminoso Analytics Methodology**.

PRE-LUMINOSO: Hypothesis

- What are your reporting objectives for the analysis?
- Determine subsets (optional)
- Upload project

B. EXPLORE & DISCOVER

- Review the Auto-Topic Suggestions
- Examine Cloud > Clusters > Concepts
- Read top related verbatims

C. BUILD TOPICS

- Organization Name/Brand
- Positive & Negative Emotions
- Feature/Function/Service/Dept/Demographic

D. ANALYSIS EXPORTS

- Download metrics to measure:
 - a single concept’s frequency and importance within the data, or
 - its association strength between any two concepts.

The Luminoso Analysis Methodology (LAM) comprises a number of foundational data investigation & insight discovery techniques that apply to all projects.

The LAM is an iterative process.

POST-LUMINOSO: Present (reporting)

- Utilize Downloads to extract meaningful metrics
- Import into third-party visualization tools (not covered)

The Luminoso Analytics Methodology (LAM) is the basic investigation and discovery mechanism. Start with the workflow to become an expert at data investigation, and then learn more advanced discovery techniques to get the most out of your data.

We offer training and guidance on specific analysis methodologies according to project type:

Emotion Analysis, NPS, Trending/Timelines, Subset Comparisons, Differential Analysis, Predictions ... Please contact our training team for details ([Tessa Rowell](#) or [Eric Pendleton](#)).

We recommend that you learn these techniques before applying a specific methodology to your project or tackling more advanced practices.

1

Visit support.luminoso.com

LUMINOSO dashboards & widgets

What are Luminoso Dashboards?

Designed to communicate relevant insights to the right people, the Dashboards module provides the ability to visualize, package, and share findings from your data.

Data is never analyzed in a vacuum. Every piece of data has its intended audience, and insight-rich text analytics may have multiple audiences. Free-form customer feedback often contains critical information for your product, sales, marketing, and even research teams.

After analysis, you can create a dashboard and populate it with widgets representing key findings. Different dashboards can be tailored to, and shared with different audiences, depending on their organizational roles and priorities. These dashboards can be shared and updated as your business demands.

What should be done in my Luminoso Analytics project first?

To ensure a smooth workflow, you need to take the following steps before building a dashboard:

- 1) Build your project**
 - Prepare data: Determine subsets.
 - Upload data.
 - What’s your hypothesis? What types of insights do you seek?
- 2) Explore your project**
 - Before thinking about a dashboard, take some time to analyze your data in the Analytics Explorer. Run the auto-topic generator.
 - The Dynamic Concept Cloud view is a good guide to understanding the prominent themes and connections in your text.
 - Our data exports provide information on share of voice and conceptual relationships, which can be used to guide your analysis.
- 3) Create/edit topics**
 - As part of the analysis process, build topics from the concepts which are most relevant to your analysis and interests.
 - Use automatically suggested topics, edit them, build your own.
 - Topics can include (but are not limited to): sentiment concepts, named entities (such as products, brands, and organizations), actions, and descriptive terms. The “name” field of each topic is what feeds into your Luminoso Dashboards, so make sure that field reads as you would like it to.
- 4) Analyze your project**
 - To get an idea of what to showcase in your dashboard, analyze your project data.
 - “Topic Match Counts” will give you **share of voice metrics** for your most important topics, while
 - “Topic Association Scores” and “Topic Subset Association Scores” will show important relationships and insights.

Once you have an understanding of what your data says, create a dashboard to visualize your insights.

Page 1 of 10

Opening pages of different manuals. The prerequisites for a useful analysis come before the interface, not after it.

Documentation at scale


2007–2017 · Selected · Available on request

What it is. A group of large reference and knowledge systems, included here as evidence of volume and sustained structure rather than for any single idea.

- **Medici 3.9 User Manual** (2009) — full product manual for an enterprise financial platform.
- **FACT Manual** — long-form technical reference documentation for a large bank's internal commercial lending platform.
- **Commercial-lending knowledge system (PRE)**, Citizens Bank (2010–13) — 500+ pages of role-based support content for a platform used by roughly 2,000 employees, maintained across releases.
- **Training Plan, Q2 2016**, Luminoso — program-level planning rather than course-level (early AI company).
- **Training Guide: thermocouple installation** — a worked instance of the Sample 5 template applied to a real regulated task, carrying reviewer comments (Pharma).

WHY IT MATTERS FOR AI FOR L&D

- Scale is where AI's promise to L&D is most real and most dangerous. A 500-page knowledge base is exactly the kind of corpus a team will want to generate, migrate, or refresh with AI — and exactly the kind where a quiet fabrication survives review because nobody rereads page 340.
- Having maintained content at that size, I can be concrete with a cohort about what actually breaks: not the writing, but the consistency, the cross-references, and the change-trigger between a source procedure and the training built on it.



Citizens Financial Group, Inc. CRM Dashboard Manual (9.5)

The Credit Risk Management (CRM) Dashboard

The CRM Dashboard is a portal solution that provides a pathway to the Credit Risk Management applications.

Figure 1 CRM Dashboard applications (tab structure)

As highlighted above, CRM Dashboard is comprised of 10 interconnected applications, which are accessible by entering <http://crmdashboard/> into your browser. Colleagues with intranet access who have CRM Dashboard-specific permissions can access Dashboard. See

Risk Analyst (Moody's): Summary of key financial information delivered from the Moody's KVM Risk Analyst statement spreading application. Key ratios are used in other applications such as the rating models. Financial spreads cannot be added/edited through this module, but there is a link to Risk Analyst (external) application. Contact [Credit Services](#) for access (not Credit Services IT).

Online Rating Application (ORA): For performing PD (risk) rating and LGD (facility) rating models for all existing and proposed Wholesale customers.

Real Estate Risk Services (RERS): For ordering and viewing appraisal, environmental, flood and cost construction information.

Financial Reporting & Covenant Tracking (FaCT): For tracking Financial/Non-Financial Requirements and associated covenants for each of your customers and their associated entities.

Field Exam Tracking (FET): For ordering, tracking and reporting on Field Examinations within the Portfolio Credit Risk Group.

Public Firms: For identifying entities with publicly traded equity. Obtains and warehouses key metrics such as Moody's KVM EDF and S&P data to be used in other applications such as the ratings models.


Managed Asset Reporting System (MARS): For Specialized Lending Services (SLS) use only. Manages key financial data and metrics involving collections and recoveries. Transfers to SLS from the Line of Business are initiated through the MARS application.

Reports: Repository for customer and portfolio-level reporting for all CRM Dashboard applications including FaCT, Annual Reviews, PD and LGD past due/coming due reports, Multibank and Shared National Credits, Leveraged lending, etc.

Tracking: Contains tracking questions (many can only be set by the Credit Officer) about the customer in session including Annual Review Exemption, Targeted Franchise, Watchlist status, and others. This data is used in different areas of the CRM Dashboard including the PRE application.

Portfolio Review: The Portfolio Review application automates the Portfolio Review process through workflow, notification and electronic data capture of customer review information. Information is leveraged from the various CRM Dashboard applications and from previously completed customer reviews, avoiding the need to re-enter data.

Proprietary & Confidential Page 6 of 129



Citizens Financial Group, Inc. CRM Dashboard Manual (9.5)

Reports

The CRM Dashboard was created and is utilized for two overarching purposes:

- 1) As an instrument for Lenders to rate customers and track their financials and
- 2) As a Management tool to monitor compliance using Reports.

Click **Reports**, expand the section/application from which you wish to generate a report, and then select the specific report you wish to generate.

Figure 2A Reports, ORA Annual Review PD/LGD Status Report (example)

Click the link for a specific report (illustrated above), a special browser will pop up (below), giving you the ability to select different criteria. When you are done, click **View Report**.

Figure 2B Reports, ORA Annual Review PD/LGD Status Report (selection criteria prior to report generation)

Proprietary & Confidential Page 7 of 129

What these samples do and do not show

What they show. Twenty years of owning the L&D function from the inside: needs analysis, competency models, course authorship, live facilitation, train-the-trainer, LMS structure, and assessment design — most of it in regulated environments where the work had to survive an auditor. Plus, the earliest half of the AI record: published translation of NLP for business readers, and onboarding documentation for an enterprise AI platform, both from 2015–16.

What they do not show. The current AI practice. Since the end of 2022, my daily work has been a source-grounded, human-verified AI methodology applied to contracts, financial records, and counsel-facing analysis — multi-model prompting across Claude, GPT, and Gemini, with every material claim checked against source before use. That work is confidential to its principals and does not belong in a portfolio. Samples can be provided under appropriate terms.

If it is useful, here is the module I would want to build. Not a tour of prompts, and not a lecture about disruption. One session in which each participant brings a real piece of their own work — a needs analysis, a set of objectives, a bank of assessment items — runs it through a model and then is walked through finding what the model invented, softened, or dropped. They leave with two things: a sense of where AI genuinely compresses the design cycle, and a verification routine they can defend to a compliance officer. That second thing is what makes the first one safe to use.



[CONTACT US / ADVERTISE](#)
[SUBSCRIBE](#)




By **Eric Pendleton**

Eric Pendleton is a Product Training Manager on the Customer Success team at Luminoso. In this role, he works with Fortune 1000 companies to implement or improve their unstructured data programs, so that they can uncover actionable insights from their data, in order to achieve their business objectives. Eric has more than 20 years of experience in software support roles, including training, technical documentation, software localization, and change management in a variety of industries from martech to biotech in the US and Europe. Bilingual in French (and shamefully proficient in Spanish), Eric holds a BA in French and Psychology from Colby College, a MA in Literature from Washington University in St. Louis, and a MS in Business from Emmanuel College.



February 08, 2017

Surveys: Time for a Major Upgrade

Let's face it. Surveys suck.
They suck to take and they really suck to design. Even a well-designed survey tends to make assumptions (in the form of leading questions) and is limited.

Authoring a comprehensive, concise, neutral, and unbiased list of survey questions is quite an art. For global organizations, having a comprehensive, worldwide survey that is offered in several languages is a must. Unfortunately, this means that by the time you get the responses back, the information is stale.