AI SUMMIT HANDS-ON LAB

Exercise 2: From Inherited Policy to Defensible Policy

Source Hierarchy • Policy Review • Source Augmentation • Scenario Testing • Optional: Applicability

PARTICIPANT PACKET

The official exercise is 20 minutes. Use Sources A1 and A2 first. Use Sources B1–B3 only for Round 2. Sources C and D are optional early-finisher extensions; do not use them with the core materials.

Suggested prompts are provided, but you are free to experiment with different approaches.

Workshop data rule

Do not enter real client, case, employee, personnel, or other confidential information into the AI tool used for this exercise. The organization, policy, leadership memo, and scenarios are fictional. External reference sources are public.

Recommended time 20 minutes + optional early-finisher extensions

Best group size 3–6 participants per table

Participant outcome A revised, source-grounded policy excerpt, a three-item source-impact log, and a completed scenario stress test — built without confidential organizational information.

Start here

Today's constraint

Use only the sources provided in each round. Do not ask the AI to independently research employment law, legal ethics, nonprofit law, or other rules. Do not enter real organizational or client information.

Scenario

The Chesapeake Community Advocacy Center (CCAC) is a fictional Maryland 501(c)(3) nonprofit that provides civil legal services, community education, and policy advocacy. It has about 65 employees, including attorneys and non-attorneys, and no dedicated communications department. Its 2021 communications policy has been amended informally over time. Leadership wants a policy that staff can actually use — without turning every external interaction into an emergency call to the Executive Director.

Core 20-minute work product

- A Round 1 issue map and targeted revision using only the fictional existing policy and leadership direction.
- A revised Round 2 policy excerpt using the external sources released later.
- A three-item log identifying what changed, which source caused the change, and whether the source is binding, persuasive, or contextual for this organization.
- A stress test of the revised policy against three realistic situations.

What good work looks like

☐ The AI does not assume that statements in the old policy are legally correct simply because they appear in a policy.

☐ The revised policy distinguishes client confidentiality from the fact that information may already be public.

☐ The revised policy distinguishes organizational campaign intervention from issue advocacy and from an individual employee's personal speech.

☐ A source that applies directly to Maryland attorneys is not silently described as a legal requirement for every employee

☐ The AI does not invent an answer to a leadership decision that the source set leaves unresolved.

☐ The policy remains usable for a small organization rather than becoming a list of absolute prohibitions and approval gates.

SOURCE A1 Round 1 — Existing policy

FICTIONAL INTERNAL SOURCE

Chesapeake Community Advocacy Center — External Communications, Media & Social Media Policy • Adopted March 2021. Treat the legal claims in this inherited policy as unverified statements, not authoritative sources.
Open link to fictional policy document

Round 1 reminder

Do not “fix” this policy from memory or general legal knowledge. Your first job is to compare an inherited document against the organization’s actual decisions and identify where outside authority is still needed.

SOURCE A2 Round 1 — Leadership direction

FICTIONAL INTERNAL SOURCE

Leadership meeting notes • August 2026 • These are the decisions CCAC has actually made. Where leadership has not made a decision, the policy should say so or flag it for resolution rather than letting the AI decide.

- CCAC wants one coordinated organizational voice for crisis communications, formal institutional statements, and responses about internal organizational matters.
- CCAC also wants attorneys and other subject-matter experts to be available for routine interviews about their areas of expertise when appropriate. Leadership has NOT decided whether those interviews can be approved by program directors or must be approved by the Executive Director.
- Client confidentiality must be protected. CCAC wants documented client authorization before using a client’s identity, case facts, quotations, image, or story for publicity or fundraising. Leadership does not want “it was already public” to serve as the organization’s consent policy.
- CCAC wants to continue advocating on housing, consumer, family, benefits, and other public-policy issues, including during election seasons, while complying with the rules applicable to a 501(c)(3) organization.

- Employees may engage in lawful political activity in their personal capacity. CCAC does not want to treat an employee's personal views as an organizational statement merely because the employee identifies CCAC as their employer.
- CCAC wants reasonable rules protecting client information, nonpublic internal information, and the integrity of official organizational accounts, but does not want broader restrictions on personal speech than are actually necessary.

Task

Baseline prompt — use this first

You are reviewing an inherited communications and social-media policy for a fictional Maryland nonprofit.

Use ONLY Sources A1 and A2. Do not browse the web and do not rely on outside legal knowledge. Treat legal assertions in the existing policy as UNVERIFIED unless Source A2 independently establishes the organization's decision.

Return:
A. An issue map with FIVE important conflicts, ambiguities, or outdated assumptions. For each, identify the existing provision, what Source A2 says, and whether you would KEEP, REVISE, REMOVE, or FLAG FOR DECISION.
B. A targeted revision of Sections 3–6, no more than 500 words total. Preserve organizational decisions that are actually settled. Do not invent an answer where leadership has not decided.
C. TWO statements in the old policy that require an external legal/professional source before you would rely on them. Label them NEEDS SOURCE.

Human review — before you get external sources

☐ Find one direct contradiction inside the old policy that the AI should have caught without any outside source.

☐ Find one statement that sounds cautious or "safe" but is really an unsupported legal assumption at this stage.

☐ Check whether the AI invented who may approve a subject-matter interview.

SOURCES B1 + B2 Round 2 — Add professional authority

SOURCE B1

Maryland Rule 19-301.6 — Confidentiality of Information (1.6)

Maryland Attorneys' Rules of Professional Conduct • Current with amendments received through June 1, 2026 • Focus on section (a) and Comments [3]–[4].

Source type: Maryland professional-conduct rule. It applies directly to Maryland attorneys; do not silently describe every provision as a legal duty of every employee.

Open current Maryland rule

SOURCE B2

ABA Formal Opinion 480 — Confidentiality Obligations for Lawyer Blogging and Other Public Commentary

American Bar Association Standing Committee on Ethics and Professional Responsibility • March 6, 2018 • Focus on pp. 1, 3–4, and the conclusion on p. 6.

Source type: ABA ethics opinion interpreting the Model Rules. The opinion itself notes that jurisdiction-specific rules and authorities control.

Open ABA Formal Opinion 480

What these sources are for

Do not ask the AI to summarize the sources generally. Ask what, if anything, they change about the inherited policy's treatment of public records, anonymized stories, client consent, and public commentary.

Source-hierarchy question

If a source speaks directly to attorneys, should the final policy extend the same practical restriction to all staff? It may — but that broader rule should be identified as an organizational policy choice rather than misdescribed as a professional rule binding every employee.

SOURCE B3 Round 2 — Add nonprofit tax guidance

SOURCE B3

IRS guidance — 501(c)(3) campaign intervention, issue advocacy, and personal speech

Internal Revenue Service • Current webpages • Use only the two short FAQs below.

Source type: federal tax guidance. Focus on the distinction among candidate campaign intervention, organizational issue advocacy, and individuals speaking in their personal capacity.

FAQ 1 — Organization position on issues

FAQ 2 — Constitutional considerations / individual speech

FAQ 1

FAQ 2

Task

Revision prompt

I am adding three external sources of DIFFERENT authority and scope: a Maryland professional-conduct rule, an ABA ethics opinion, and IRS guidance for section 501(c)(3) organizations.

Re-evaluate the Round 1 policy using ALL supplied sources. Do not use outside law. Do not assume every source applies to every employee or resolves every policy choice.

Return:
A. A revised version of Sections 3–6, no more than 500 words total. Revise only where the new sources or Source A2 warrant a change.
B. A THREE-ITEM source-impact log. For each item identify: (1) the Round 1 wording/idea, (2) the revised wording/idea, (3) the new source that caused or supported the change, (4) the source type/scope, and (5) why the change matters.
C. Up to TWO issues that the source set still does not answer. Do not invent a rule to fill the gap.

Important: If the organization chooses to apply a restriction more broadly than the external source legally requires, say that it is an organizational policy choice rather than presenting it as the source's command.

Save both versions

Do not overwrite Round 1. The comparison is part of the exercise: the question is not whether Round 2 sounds better, but whether the new sources caused defensible substantive changes.

Compare Round 1 and Round 2

☐ Which change in the three-item log most materially changes what a staff member may do?

☐ Did the model correct the idea that public-record status automatically resolves confidentiality?

☐ Did it distinguish a 501(c)(3) organization taking a position on an issue from supporting or opposing a candidate?

☐ Did it treat an ABA opinion as persuasive guidance rather than silently calling it Maryland law?

☐ Did it convert any unresolved organizational choice into a made-up “requirement”?

☐ Did Round 2 become more restrictive in places where the sources did not actually require more restriction?

Source-selection question

Was Round 2 better because the AI “reasoned harder,” or because you gave it sources that addressed assumptions the old policy had previously treated as settled?

CORE: Scenario stress test — does the policy work?

Give the three scenarios on the next page to the AI in the SAME conversation. For each scenario, require the model to apply only the revised policy and the supplied sources.

Stress-test prompt

Apply the revised policy to each scenario I provide.

For each scenario return:
1. Your policy result: PERMITTED, PROHIBITED, APPROVAL REQUIRED, or POLICY UNCLEAR.
2. The specific revised policy provision supporting that result.
3. Any external source that materially supports the result.
4. Classify the basis as: SOURCE-BACKED CONSTRAINT, ORGANIZATIONAL CHOICE, or UNRESOLVED.
5. One uncertainty or fact that could change the answer.

Do not use outside knowledge. If the policy or sources do not answer the question, say POLICY UNCLEAR rather than filling the gap.

Scenario stress test — three situations

SCENARIO 1 — The public court opinion

A CCAC staff attorney represented a tenant in a case that resulted in a published court opinion. On the attorney's personal LinkedIn account, the attorney wants to write: "Proud of a big win for a tenant facing eviction this week," and link to the public opinion. The post does not state the client's name, but the client has not been asked for consent.

SCENARIO 2 — Issue advocacy during election season

CCAC's official account plans to post: "HB 999 would reduce access to counsel for low-income Marylanders. Maryland lawmakers should reject it." Two candidates for statewide office have publicly taken opposite positions on the bill. The proposed CCAC post does not name, picture, praise, or criticize either candidate.

SCENARIO 3 — The reporter calls a subject-matter expert

A reporter asks a CCAC benefits attorney to explain, on the record, how a newly announced public agency rule generally affects low-income Marylanders. The reporter is not asking about any client or case and wants a quote within 30 minutes. No crisis or internal CCAC issue is involved.

Human audit — do not accept the labels automatically

Scenario	AI result	Source/policy support	Constraint / choice / unresolved	Human verdict

The most important question

When the AI says something is "prohibited," can you tell whether the prohibition came from an authoritative source, from a leadership decision, from the inherited policy, or from the model's instinct to be cautious?

FINISHED EARLY? KEEP EXPLORING

The official 20-minute exercise ends after the scenario stress test and facilitator debrief. If your table finishes early, choose an optional source below to explore this exercise topic further.

OPTIONAL ADD-ON C — New source can expose new risk

SOURCE C

Social Media — Employee Rights and Protected Activity

National Labor Relations Board • Current webpage • Focus on the explanation of work-related social-media activity and group action.

Applicability is fact-dependent. This exercise asks whether the source creates a reason to revisit a blanket policy rule — not whether the AI can conclusively decide an employment-law case.

Open NLRB social-media guidance

New scenario

SCENARIO C — Workload discussion

On a personal Facebook account, an employee posts: “Our caseloads are impossible. Coworkers, who wants to sign a letter asking management to hire more staff?” Several coworkers respond that they want to participate.

New-source prompt

Add Source C to the source set. Re-evaluate the revised policy’s rule on employees criticizing management, compensation, workload, staffing, or workplace conditions.

Do NOT decide whether this specific employee is legally protected unless the supplied source is sufficient to do so. Instead return:
1. The policy language you would reconsider.
2. What Source C says that creates the concern.
3. A narrower policy approach that protects legitimate organizational interests without simply prohibiting all discussion of workplace conditions.
4. Any applicability or factual question that should be reviewed by HR or counsel.

Debrief question

Did the original “protect our reputation” rule look safer before you had this source? What does that tell you about asking AI to make a policy merely “more conservative”?

OPTIONAL ADD-ON D — Useful source ≠ automatically binding source

SOURCE D

Social Media Accessibility Guidance

Section508.gov • Current webpage • Focus on alt text, captions, readable content, and accessible social-media practices.

Section 508 directly governs federal agencies and federal ICT. The fictional CCAC is not a federal agency. Treat this as an applicability exercise, not an invitation to misstate federal requirements.

Open Section508.gov guidance

Task

Applicability prompt

Add Source D to the source set.

CCAC wants its official social-media content to be accessible. Review Source D and propose up to FOUR practical standards that CCAC could voluntarily adopt.

For each standard identify:
A. What the source recommends or requires for its intended audience.
B. Whether the supplied source establishes that the same rule is legally binding on fictional CCAC.
C. How you would phrase the standard in CCAC's policy without overstating the source's authority.

Do not use outside disability law. If the source does not establish a legal requirement for CCAC, say so.

Why this add-on matters

- Good policies often borrow from strong external guidance even when the guidance is not legally binding on the organization.
- AI systems frequently erase that distinction and rewrite "recommended practice" as "required by law."
- Source literacy includes asking not only "is this a good source?" but also "good source for what proposition, and applicable to whom?"

Core debrief — what actually made the policy better?

- What could the AI improve using only the organization's own policy and leadership decisions?
- Which Round 2 source caused the biggest substantive change?
- Did anyone's model preserve the public-record exception from the old policy after reading the Maryland confidentiality rule and ABA opinion?
- Did any model turn "501(c)(3)" into a blanket ban on issue advocacy or personal political speech?
- Where did the model distinguish a professional obligation that applies directly to attorneys from an organizational rule applied to everyone?
- Which issue remained a genuine leadership choice rather than a research problem?
- Would a more restrictive policy necessarily have been a better policy?

The lesson to land

AI can make an inherited policy clearer in minutes, but clarity is not the same as correctness. The high-value workflow is: establish the organization's actual decisions → identify unsupported assumptions → add sources with the right authority and scope → revise with a source-impact log → test the policy against real situations.

Use-it-Monday pattern

Start with an existing policy or draft → compare it to current leadership decisions → identify statements that need outside authority → add targeted sources → require source-driven revisions → stress-test the result before adoption.

Your transfer idea

Name one internal document from your own organization that could support a similar workflow without using client or sensitive personnel information:

Document / policy	What internal decision source would you pair with it?	What external source might sharpen it?	What would a human still decide?