[This file is supplemental material to Straight, *Grey Areas and Egress: Artifact-Based Collaborative Learning in Cybersecurity Ethics*, prompt 10.2 (2026), doi: 10.31719/pjaw.v10i2.280]

EXODUS Collaborative Writing Assignment

Instructor Guide

Ryan M. Straight

Table of contents

Before the Assignment 1
The In-Class Session 2
The Written Analysis 3
Assessment 3
Common Student Difficulties 4
Cross-Case Connections 4
Adaptation to Other Domains 5
A Closing Note 6

This instructor guide accompanies Ryan M. Straight, "Grey Areas and Egress: Artifact-Based Collaborative Learning in Cybersecurity Ethics," Prompt: A Journal of Academic Writing Assignments, Issue 10.2 (2026), and the companion Project Overview supplement describing the assignment as taught. This guide is written for adopters (instructors considering running the EXODUS Collaborative Writing Assignment in their own courses) and addresses facilitation decisions, common difficulties, and adaptation pathways drawn from iterative teaching of the assignment in CYBR 329: Cyber Ethics.

Before the Assignment

The assignment works best when students have spent roughly the first half of the semester developing comfort with three things: the consequentialist, deontological, and virtue-ethics framework triad, the procedural workings of relevant federal statutes (in this iteration, the CFAA and EEA, with ITAR and EAR available as extension material), and the experience of writing shorter individual analyses that combine normative argument with documentary evidence. By the time EXODUS arrives, students have done at least two prior case analyses. The assignment is demanding. It asks students to hold three frameworks, two statutes, multiple analytical tasks, and four stakeholder perspectives in productive tension at once. Sending them into that task cold, without the prior work of simpler cases, invites collapse into single-framework thinking.

Groups of three to four work best. I let students self-select, which produces occasional homogeneity but also produces investment. Assigned groups can be more analytically diverse but tend to generate more procedural friction in the 75-minute window. When I have thirty-six students, I end up with nine or ten groups. I have not found larger groups (five or six) productive here, because the

analytical tasks require each member to hold a distinct thread, and larger groups tend to silo contributions rather than synthesize them. Requiring this in-class work encourages students to come prepared.

The In-Class Session

The 75-minute collaborative session is the pedagogical center of the assignment and also the place where facilitation decisions matter most. I begin with roughly ten minutes of framing: a brief reminder of the three ethical frameworks and a short note on the interpretive posture I want students to take (that frameworks are analytical tools, not answer keys). The framing matters less for information transfer than for signaling what kind of work the next hour is meant to do.

What students then do across the remaining time is structured but not rigid. Groups work through the required analytical tasks (in the EXODUS case: ethical framework application, stakeholder analysis, legal analysis, the ethics of military AI, corporate response evaluation, and policy development) at their own pace. Some groups work linearly, some jump around. I do not insist on any particular sequence, though I do encourage them to take the tasks generally in the order presented because policy recommendations produced without grounding in the ethical and legal analysis tend to be superficial.

My own role during this session is not to answer. It is to move among groups and watch for the specific failure mode of consensus forming too quickly. When I hear a group converging on a unified position after fifteen minutes of discussion, I know something is wrong. The case is designed to resist easy convergence. Quick agreement almost always means the group has settled on a single framework and is deriving its other answers from that framework's logic rather than treating the tasks as independent analytical probes. When I observe this pattern, I pose a question that complicates the emerging consensus. "What would a virtue ethicist say about Chen's choice to act secretly rather than openly?" works well when a group has consolidated around a consequentialist reading of beneficial outcomes. "How might this same evidence look from the perspective of the colleagues who stayed?" works when a group has fixated on Chen as the only stakeholder whose motives matter. These questions do not provide answers. They model the kind of multi-perspective interrogation the written analysis will require.

A common adopter concern at this point is what to do about groups that get stuck on the legal analysis. Students without prior law exposure can read the CFAA and EEA and then freeze, assuming they need technical legal expertise to proceed. I try to depressurize this quickly, reminding them the ethical analysis does not wait on legal certainty, and in fact part of what the assignment teaches is that legal frameworks designed decades ago may not resolve the ambiguity the case presents. I encourage groups to articulate what the statutes seem to prohibit, where the facts meet the statutory elements and where they do not, and what the remaining ambiguity looks like. That is usually enough. Groups that try to resolve the legal question before moving on tend not to finish.

The session concludes with brief oral presentations of one to two minutes per group. These serve two purposes: they expose students to interpretations other than their own, and they provide

formative feedback without the performance anxiety of formal grading. I do not assign a grade to these presentations. I comment briefly on what I found analytically productive in each and, where useful, name a question the group's presentation opened up that they might take into the written analysis.

The Written Analysis

Following the in-class session, groups produce a 900 to 1200 word collaborative written analysis, due at the end of the module, generally 3-4 days after the in-class writing. The word count is chosen deliberately. Shorter, and groups cannot adequately integrate the tasks. Longer, and they begin padding or writing in fluffy loops that weaken the analytical argument.

Groups organize the writing process autonomously. I have experimented with assigning roles (one person on legal analysis, one on framework application, one on policy) and not only is the work produced by such groups is consistently weaker than the work produced by groups that sort responsibilities organically, but the students expressly state they dislike the format. The role-assignment approach produces writing that reads as stitched-together sections, and invariably one student will dislike their role. The organic approach, messier and more variable, produces the integrative writing the assignment actually wants. I now simply tell students that their group will produce one document, that they should decide among themselves how to produce it, and that I expect evidence of genuine synthesis rather than divided individual work.

The companion reflection (250 words, submitted alongside the analysis, not counted toward the word limit) is the mechanism that makes this work. In it, groups describe their collaborative process, identify who contributed what, reflect on how the collaboration shaped their thinking, and identify strengths, weaknesses, and questions they still have even after the analysis. This reflection is the single most useful artifact I receive from this assignment. It tells me whether one member carried the analytical weight, whether the group genuinely synthesized perspectives or merely averaged them, and whether productive disagreement occurred. The reflection also signals to students that I am attending to their process, which tends to improve their process.

Assessment

The written analysis is assessed across six weighted dimensions, namely: multi-stakeholder analysis (25%), synthesis and integration of perspectives (20%), policy development with implementation pathways (20%), evidence of genuine collaborative process (15%), understanding of technology transfer and strategic protection concepts (15%), and writing quality including proper citation practices (5%). The weights reflect a deliberate judgment about what the assignment is actually teaching. Stakeholder analysis weighs heaviest because it is the analytical move most often performed poorly. Synthesis and policy development tie for second because they are the move that distinguishes sophisticated from mechanical work. The understanding-of-technology-concepts category is weighted less than one might expect because this is an ethics course, and technical detail serves ethical argument rather than the reverse. Writing quality is weighted lightly because, in my

experience, writing quality improves dramatically when the analytical work is good, and very little when the analytical work is superficial regardless of how carefully sentences are constructed.

What distinguishes excellent work from adequate work is *sophistication* rather than accuracy. The case studies are intentionally sprawling and dense, full of ostensible contradictions and blind alleys. Adequate submissions apply frameworks mechanically (consequentialism says X, deontology says Y, virtue ethics says Z, conclusions follow) and treat the tasks as independent questions with independent answers. Excellent submissions demonstrate nuanced understanding of how a single framework can support competing conclusions depending on which stakeholders and duties the analysis foregrounds. They also show evidence of genuine deliberation in passages where the group visibly considered an alternative and explained why they moved past it, rather than presenting their final position as the only defensible one. This distinction reflects how ethical expertise actually functions in cybersecurity governance.

Common Student Difficulties

Generally speaking, I've noticed a few patterns that appear reliably across semesters. First, students from technical backgrounds often expect the legal analysis to resolve the ethical analysis, and find it destabilizing when it does not. (I repeatedly remind students: just because something is legal doesn't mean it's ethical, and vice versa.) I address this directly during the in-class session by noting that the legal framework's inability to resolve the question is itself part of what the assignment is teaching. Second, students often overreach on the policy recommendations, proposing sweeping regulatory changes rather than organizational practices. The assignment asks for institutional policy (what a company like Quantum Frontier might actually implement) and I gently redirect groups toward this scope when I see them drafting omnibus federal legislation. Third, a small number of students each semester find the collaborative process itself frustrating, particularly when their group includes members they perceive as less prepared. The reflection component gives these students a place to voice this experience, and I take their feedback seriously when it appears. In practice, the assignment produces stronger writing from mixed-preparation groups than from homogeneous high-preparation groups, but the experience is not always comfortable.

Cross-Case Connections

For instructors using the full *Cyber Dimensions* textbook, EXODUS connects productively to several other cases. AVALANCHE explores whistleblower ethics and the tension between professional obligations and constitutional concerns, offering a useful contrast with EXODUS's framing of Chen's actions as IP-protection violations rather than whistleblowing. CO(G)NECT addresses ethics and legislation in emerging technology (brain-computer interfaces) with dual-use concerns and provides a bridge from EXODUS's quantum-computing specificity to broader AI governance. TECHGIANT takes up AI ethics and the innovation-versus-control tension at the corporate-culture level. BELLATOR extends EXODUS's professional-ethics thread into the international cybersecurity-operations domain. Running EXODUS in sequence with one or two of these cases, rather than as a standalone, tends to deepen student engagement with how the same ethical questions recur under different tech-

nical and organizational conditions. In actuality, each accompanies an entire module in the class: seven modules, seven case studies.

Adaptation to Other Domains

The structural template transfers readily. What EXODUS provides (primary source artifacts plus multi-task analytical synthesis plus collaborative writing plus process reflection) can be reproduced in any domain where professionals must interpret documentary evidence, apply contested frameworks, and develop defensible recommendations amid uncertainty. A public-health course might use epidemiological data, interview transcripts from affected community members, internal communications from health officials, and media coverage of an outbreak. Groups could evaluate institutional response through consequentialist, deontological, and virtue-ethics lenses while developing policy recommendations for future incidents. A privacy-regulation course might use internal memoranda, regulatory filings, user complaints, and journalism to examine a company's data-handling practices. A digital-rights course might use court filings, technical specifications, and advocacy-group documentation to analyze a content-moderation controversy.

The artifacts do not *need* to be fictional, but invite bias and prejudice otherwise. What they *do* need to have is enough documentary weight to support multiple plausible readings, enough specificity that students cannot resolve ambiguity by retreating to abstraction, and enough ethical texture that the multi-task structure generates genuinely different analyses in different groups. Cases drawn from public court records, journalism, FOIA releases, and institutional transparency reports can serve these purposes well, with the caveat that cases involving still-living parties require careful attention to how the assignment is framed and how student analyses are handled.

In the AVALANCHE case, clearly a reference to the Snowden revelations of 2013, students are *not* presented with Edward Snowden. Instead, they're presented with an analog and a similar event and sequence. When students hear "Snowden" (especially cybersecurity and intelligence students), they bring considerable amounts of pre-existing knowledge and feelings to the case. Making it fictional relieves them of this and allows students to analyze and make decisions based on the content of the case, itself.

For instructors developing their own artifact-based cases, two open resources support the work. The Cyber Dimensions OER Toolkit provides the pedagogical foundations, authoring guidance, and templates for building artifact-based cases in any domain. The Cyber Dimensions Case Study Demo is a complete worked implementation (the CIRCUIT case, a cybersecurity incident at a municipal utility) that shows the full production setup in rendered form. The source repository contains all Quarto source files, the YAML worldbuilding system, and the rendering pipeline for inspection or forking. Adopters who want to see how a case gets built from the ground up should start there.

A Closing Note

The assignment works best when the instructor treats it as a genuinely open inquiry rather than a structured vehicle for reaching a predetermined conclusion. Students are quick to detect facilitation that leads them toward an answer the instructor has already decided is correct, and when they detect this, the pedagogy breaks. My own practice has been to enter the in-class session with no preferred reading of Chen's actions. I have held positions I later revised on the basis of student arguments. The assignment is most valuable when the instructor, like the students, is willing to be moved by the evidence.

Ryan M. Straight University of Arizona ORCID: 0000-0002-6251-5662