EXODUS Collaborative Writing Assignment

The Case

Alex Chen is a quantum computing engineer at Quantum Frontier, a leading AI company with government contracts in defense and intelligence. After five years of high-impact work on Project Daedalus, the company's flagship quantum machine-learning initiative, Chen accepts a position at NexGen Cybernetics, a competitor known for an open-source, collaborative approach to AI development.

Before departing, Chen copies portions of the quantum AI algorithms she worked on at Quantum Frontier to personal cloud storage. Her stated intent, later recorded in her exit interview, centers on two threads: using some material as a portfolio piece, and continuing to develop algorithms she feels personally invested in. Quantum Frontier detects the exfiltration. The FBI investigation that follows produces charges under the Computer Fraud and Abuse Act (CFAA) and the Economic Espionage Act (EEA).

The case itself consists of primary source artifacts rather than a summarized narrative: internal memos, encrypted chat logs, FBI investigation reports, exit interview transcripts, legal strategy documents, and news coverage. Your task is to reconstruct meaning across these artifacts and arrive at a defensible position on the ethical, legal, and policy dimensions of what happened.

Before the Collaborative Session

Read the entire EXODUS case independently, including all artifacts. As you read, compose brief notes identifying key evidence and your initial analytical observations. These notes are low-stakes and ungraded, but your group will rely on them during the collaborative session. Students who arrive having only skimmed tend to become passengers in their groups rather than contributors.

Phase 1: In-Class Collaborative Synthesis (75 minutes)

Form self-selected groups of three to four students. Over the course of a single 75-minute in-class session, work through the six analytical tasks below. You do not need to tackle them in order, but do not tackle Task 6 (policy development) until at least three of the prior tasks are substantially engaged. Policy recommendations produced without grounding in ethical and legal analysis tend to be superficial.

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

The Six Analytical Tasks

1. **Ethical Framework Application.** Apply consequentialist, deontological, and virtue ethics frameworks to evaluate Chen's actions. Explain how each framework might yield competing conclusions depending on interpretive emphasis. Resist the temptation to treat frameworks as answer keys.
2. **Stakeholder Analysis.** Conduct a comprehensive stakeholder analysis identifying all affected parties. Consider Chen's former colleagues at Quantum Frontier, the two companies, Chen herself, the broader public who might benefit from or be harmed by AI applications, and any relevant institutional actors (government agencies, the scientific community, investors).
3. **Legal Analysis.** Determine potential legal violations under CFAA, EEA, ITAR, and EAR. Evaluate whether Chen's actions meet the statutory elements of each, attending to ambiguities in "authorization," "trade secret," and intent. Note where the facts meet the statutory elements cleanly and where they do not.
4. **Ethics of Military AI.** Analyze the ethics of military artificial intelligence development, as raised in Chen's exit interview concerns. Consider benefits, risks, and the ethical responsibilities of AI researchers and the companies that employ them.
5. **Corporate Response Evaluation.** Evaluate the corporate response to the alleged data exfiltration through deontological and social contract theory lenses. Examine whether institutional actions reflect agreements that rational individuals would accept as fair.
6. **Policy Development.** Develop three organizational policy recommendations (addressing institutional practice, not regulatory or legislative proposals) with implementation analysis. Each recommendation should name the ethical framework grounding it, describe challenges in implementation, and address how it balances competing stakeholder interests.

The session ends with brief oral presentations of one to two minutes per group, sharing findings and exposing your group to alternative interpretations from others. These presentations are formative and ungraded.

Phase 2: Collaborative Written Analysis

Your group produces a **collaborative written analysis of 900 to 1200 words, due at the end of the module.** The written analysis must demonstrate integrated mastery of multiple disciplinary writing conventions:

- Legal analysis citing specific statutory provisions
- Ethical argumentation using philosophical frameworks
- Policy writing with implementation considerations

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

Work exclusively with the case artifacts. External research is neither required nor expected. The analytical challenge is documentary interpretation, not literature review.
Your group organizes the writing process autonomously. Some groups divide by framework, others by task, others draft collectively. There is no required approach, but there is a required outcome: the final document should read as an integrated argument rather than stitched-together individual sections.

Collaborative Process Reflection
Alongside your analysis, each group submits a brief **250-word reflection** describing:

- Your collaborative process: how responsibilities were sorted, how disagreements were navigated
- What each member contributed
- How the collaboration shaped your group's thinking

This reflection is not counted toward the 900 to 1200 word limit of the analysis. Write it honestly. It helps the evaluator understand your process and serves as a check against uneven contribution.

Assessment Rubric
The written analysis is assessed across six weighted dimensions:

Dimension	Weight	What the evaluator is looking for
Multi-stakeholder Analysis	25%	Depth and breadth of stakeholder identification Sophistication of handling competing legitimate interests.
Synthesis and integration of perspectives	20%	Integration of ethical, legal, and policy analysis as woven argument rather than discrete sections. Recognition that frameworks yield competing conclusions depending on interpretive emphasis.
Policy development with implementation pathways	20%	Quality of organizational policy recommendations. Grounding in identified ethical frameworks. Implementation analysis including challenges and trade-offs.
Evidence of genuine collaborative process	15%	Reflection quality. Evidence of productive disagreement, shared intellectual engagement, and equitable contribution.

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

Understanding of technology transfer and strategic protection concepts	15%	Accurate grasp of CFAA and EEA elements as applied to ambiguous facts. Recognition that decades-old statutes may not anticipate modern technology practices.
Writing quality including proper citation practices	5%	Clarity, organization, mechanics. Correct citation of artifacts and statutes. Distinction between descriptive summary and analytical argumentation.

What distinguishes excellent work from adequate work is sophistication, not accuracy. Adequate submissions apply frameworks mechanically 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. Excellent submissions also show evidence of genuine deliberation: passages where the group visibly considered an alternative position and explained why they moved past it, rather than presenting the final position as the only defensible one.

A Note on the Work

This assignment is demanding. It asks you to hold three frameworks, two statutes, six analytical tasks, and a full stakeholder landscape in productive tension at once. The case is designed to resist easy convergence. If your group reaches a unified position early, something has probably gone wrong: you have likely settled on a single framework and are deriving your other answers from its logic rather than treating the tasks as independent probes.

The gap between what is legal and what is ethically defensible is not a bug in the case but a feature. Professionals in cybersecurity governance navigate contested interpretations rather than apply predetermined formulas. This assignment asks you to practice that navigation.

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