Chapter 5: The Spite-Driven Will? Restorying Elon Musk and xAI

Abstract

This chapter examines Elon Musk's AI leadership through xAI's public story, its products, and the material infrastructure required to build and deploy them. I use PERVIEW's seven-step Restorying process as a provisional organizational-development-and-change (ODC) approach. This is an archival analysis and a proposal for future inquiry; I do not claim that xAI participated in a Restorying intervention. I also revise the earlier title's "spite-driven" characterization into a question rather than a finding about Musk's motives. The public record supports an analysis of competing accounts: xAI's stated aim to advance understanding, its self-description of Grok as rebellious, the company's rapid integration of compute, models, software, and X, and counter-stories about environmental burden and harmful image generation. The current workbook baseline remains provisional: X=5, Y=1, Z=2, A=NE; M is not calculated. I specify the evidence needed to revise these positions and identify a possible path toward answerable infrastructure governance.

Keywords

Elon Musk, xAI, Restorying, PERVIEW, organizational development and change, Concord Tesseract™, XYZA, AI infrastructure, answerability

1. Why Restorying Is the Candidate ODC Approach

This chapter asks how organizational development and change might help an AI corporation and the wider AI industry move toward Concord: a condition in which power, production, and answerability are brought into a more accountable relationship. The case is Elon Musk and xAI, now situated inside SpaceX's larger corporate structure. I select Restorying as the candidate approach because the case turns on competing stories that are also material arrangements: a story about understanding the universe, a story about speed and vertical integration, and counter-stories from people who experience the consequences of the systems and infrastructure being built. The method makes room to examine those stories without assuming that one executive's account represents the whole organization or the public interest.

I use the PERVIEW site as the current procedural description of the seven-step embodied Restorying method and my Restorying page as the earlier organizational formulation. The common movement is from characterizing an organization and its received identity, through

externalizing a problem pattern and examining its payoffs and consequences, toward reclaiming exceptions, rehistoricizing the account, and publicizing a new story with witnesses and support networks (Boje, 2019; PERVIEW Inc., n.d.). The sites describe a method; they are not evidence that xAI has used it. My use here is analytic and prospective. The case is based on public records and company statements, so the embodied and dialogic elements of an actual intervention cannot be supplied by archival analysis.

The two pages are complementary but do not use every label identically. The PERVIEW page uses "Characterize" and "Rehistoricize"; my earlier Restorying page uses "Recharacterize" and "Restory," describing the sixth step as rehistoricizing the received account (Boje, 2019; PERVIEW Inc., n.d.). I retain PERVIEW's current seven-step sequence in the table below while keeping the older terminology visible. For this archival case, the VIEW cycle's embodied questions are prompts for a future facilitated process; the records reviewed here do not provide access to participants' bodily experience or internal states.

I also revise the title. "The Spite-Driven Will" appeared in the earlier working file, but a title that asserts an executive's inner motive exceeds the evidence I have assembled. The question mark leaves the phrase as a hypothesis about a public narrative that future inquiry could examine. The chapter does not diagnose Musk, infer a private psychological state, or treat adversarial rhetoric as proof of spite. The analytic unit is a pattern of publicly observable leadership claims, organizational choices, and consequences. If the evidence does not sustain "spite-driven," the phrase should be removed from the final title.

2. Case Boundaries and Evidence

The corporate boundary changed while this chapter was being prepared. xAI announced Grok in November 2023 and described it as a conversational AI with real-time knowledge through X and a "rebellious streak" (xAI, 2023). SpaceX later reported that it completed its acquisition of X.AI Holdings Corp. on February 2, 2026, after xAI had acquired X and SpaceXAI Corp. in March 2025. In its September 2026 Form 10-Q, SpaceX described a vertically integrated Super Intelligence segment spanning infrastructure, Grok, Cursor, and X (Space Exploration Technologies Corp., 2026). I therefore use "xAI" for the company and product history under examination, while naming SpaceX when discussing the current parent-company structure and disclosures.

This shift matters to the analysis. The object is not a stand-alone start-up in the same sense as it was in 2023. The filings describe a corporate system that links computing infrastructure, a frontier model, a software-engineering product, and a real-time information platform. That integration can shorten coordination paths and give the organization control over more of the technical stack. It also concentrates decision rights and makes it harder to assign responsibility to a single product boundary when harms or benefits cross infrastructure, model, and platform layers. The filing is a formal corporate account of the business structure; it does not settle how power is experienced by workers, users, or surrounding communities.

The evidence base is deliberately bounded. I give priority to official xAI and SpaceX statements, the SEC filing, regulator announcements, and court decisions. A company description establishes what the company says it is doing; it does not independently verify the social value of the activity. A complaint establishes what a plaintiff alleges; it does not establish liability. A regulator's decision to investigate identifies questions under examination; it is not a final enforcement determination. I use these distinctions throughout, because the case loses accuracy if allegations, procedural rulings, and findings on the merits are collapsed into one story.

The chapter's provisional cutoff is November 11, 2026. It should be updated before publication. Later developments may change corporate structure, product safeguards, regulatory status, litigation posture, and the evidence available for scoring. The material assembled here does not include interviews with xAI employees, affected residents, users, regulators, or Musk. Accordingly, the scores are not a representative measure of the full organization, and they should not be read as settled judgments about a person's character.

3. Company Story and Counter-Stories

3.1 The promise of understanding and the attraction of speed

xAI's public mission is framed as helping humanity pursue understanding and knowledge. Its first Grok announcement also presented a recognizable character: witty, willing to answer broadly, and possessed of a "rebellious streak." The product's access to real-time information through X was described as a distinctive advantage (xAI, 2023). This is a story of inquiry coupled to immediacy. It places the AI system close to the flow of public information and casts unconventional behavior as part of the product's identity.

The company's own milestones make the acceleration narrative concrete. xAI described a $6 billion Series B in May 2024 and said the funds would support its first products, advanced infrastructure, and research. In December 2024, it announced a further $6 billion Series C and described Colossus as operating within 122 days, with plans to double the cluster to 200,000 GPUs (xAI, 2024a, 2024b). Those are company-reported funding and construction claims. They show how xAI presents compute as a strategic advantage and speed as an organizational accomplishment; they do not independently establish how costs and benefits were distributed.

That story has a serious and potentially productive side. AI research and deployment depend on computing capacity, technical talent, infrastructure, and iterative learning. Building infrastructure quickly can reduce delays and enable new experiments. A leader who describes a large ambition can align resources around difficult technical work. In Restorying's first step, I therefore begin with the organization at its best rather than defining it through its most troubling episode. A credible "best" account might include scientific curiosity, technical competence, willingness to challenge settled assumptions, and the capacity to build systems that make useful knowledge more accessible. Those are candidate strengths to be tested with people inside and outside the company, not conclusions inferred from branding.

At the same time, speed can become a story filter. When rapid construction, model release, and platform integration are treated as proof of progress, questions about consent, environmental load, moderation, repair, and community voice can be recoded as obstacles to momentum. This is not a claim that all xAI decisions are made in this way. It is the organizational pattern I propose examining: whether speed and scale help the company learn, or whether they narrow the field of stories that can influence a decision before deployment.

3.2 The organization becomes infrastructure

The September 2026 SpaceX filing describes the Super Intelligence segment as a vertically integrated platform. The company lists infrastructure, Grok, Cursor, and X within a single segment and describes the segment as a linked offering to enterprise and consumer customers (Space Exploration Technologies Corp., 2026). This corporate arrangement extends the case beyond model behavior. It brings together the compute needed to train and operate models, the model interface, a software-development product, and a large public communications platform. A Concord analysis must follow that chain. If power is exercised

through infrastructure, an interface-only account will miss where energy, access, and external costs are organized.

The Memphis-area Colossus build-out makes the material layer visible. The company's own site described rapid construction and later updates about turbine removal and engagement with public authorities. These are company-authored accounts and should be read as such (SpaceXSI, 2026). Community organizations and the NAACP, by contrast, alleged that xAI and its subsidiary operated gas turbines without required permits and sought court intervention. The district court's October 30, 2026 order denied xAI's motion to dismiss, concluding at that procedural stage that the NAACP had statutory standing, the turbines were not "mobile sources" under the relevant Mississippi rules, and the citizen-suit provision was not unconstitutional. The court allowed the claims to proceed; it did not issue a final judgment establishing that every alleged violation occurred or determine ultimate liability (National Association for the Advancement of Colored People v. X.AI Corp., 2026).

This distinction does not reduce the importance of the dispute. It clarifies what can be said now. The complaint and the court's ruling make public health, permitting, and community participation central to the case. The company's later public statements and regulatory filings add its side of the account. A full Restorying process would invite affected residents to describe their experience directly, identify what decisions they could influence, and specify what repair would mean to them. A public archive cannot stand in for that participation.

3.3 Product risk and regulatory attention

The same discipline applies to Grok and image generation. In January 2026, the European Commission opened formal proceedings under the Digital Services Act to examine whether X had assessed and mitigated risks associated with Grok functionalities on its platform, including risks concerning manipulated sexually explicit images and material that might amount to child sexual abuse material (European Commission, 2026). The Commission announcement marks an investigation, not a final finding. In October, Ofcom reported that its investigation into X was examining risk assessment, reporting and complaints processes, and moderation decisions. It also described restrictions that SpaceXSI had implemented for UK users on image generation and pornographic content in the standalone Grok service after regulatory engagement (Ofcom, 2026). As of the chapter cutoff, Ofcom had not announced a final determination in the investigation.

These episodes raise an organizational-development question that is broader than whether a safeguard exists. How early did the organization identify foreseeable harms? Who had authority to slow or alter deployment? Could people represented or affected in generated images obtain redress? Did the company change the system because of its own learning, user reports, public criticism, regulator engagement, or some combination? What evidence would show that safeguards work across products and jurisdictions? The regulator's account and the company's implemented restrictions are relevant evidence of change, but they do not by themselves establish that the risk is resolved or that affected users shaped the redesign.

The two material domains—energy and content—should not be collapsed into a single moral verdict. They are different systems with different affected publics and legal rules. Their connection lies in an organizational question: whether the company's power to scale technology is matched by mechanisms that make consequences visible, give affected stakeholders voice, and allow decisions to be revised before harm becomes entrenched. Restorying supplies a method for bringing those different accounts into a shared inquiry without pretending that one story can substitute for all the others.

4. A Provisional Restorying Design for xAI

The following table adapts the seven steps on PERVIEW and David Boje's Restorying page to xAI's public record. It is a proposed ODC design, not a report of work conducted inside the company. In a live process, the participants would include people who build and deploy the systems, workers whose roles change, users, communities near infrastructure, regulators, and people affected by model outputs. Participation would need independent facilitation, informed consent, protection against retaliation, and a way for participants to see how their contributions affected decisions. Public relations exercises or executive storytelling sessions alone would not satisfy that requirement.

Step	Question for this case	Evidence or activity to carry forward
1. Characterize	What does xAI do at its best, and who can confirm that account?	Document useful scientific and technical contributions; interview workers, users, partners, and affected communities.
2. Externalize	What recurring pattern is the problem: speed without adequate feedback, fragmented responsibility, or another pattern?	Name a system-level pattern rather than diagnosing Musk or blaming an individual; map decision rights and story flows.

Step	Question for this case	Evidence or activity to carry forward
3. Sympathize	What payoff keeps the pattern in place?	Examine time-to-compute, control over the stack, competitive positioning, financing, and the appeal of a bold product identity.
4. Revise	What costs and consequences have been excluded or minimized?	Trace community, labor, privacy, consent, environmental, and legal consequences using independent evidence and affected-party accounts.
5. Strategize	Where are the Little Wow Moments when the pattern changed?	Identify documented cases of restraint, repair, safety improvement, shared authority, or community-responsive design; test whether these were durable.
6. Rehistoricize	Can the company's past be retold without erasing either achievement or harm?	Construct a timeline with competing accounts, corrections, unresolved allegations, decisions, and learning loops.
7. Publicize	How will a new story be witnessed and held accountable?	Publish commitments, owners, deadlines, indicators, and a response process; invite external review and report missed commitments.

The first step is not praise detached from evidence. PERVIEW begins with an organization-at-its-best account while also asking about unsafe conditions, draining work, and stories that remain unspoken (PERVIEW Inc., n.d.). Applied here, the "best" story should be developed with employees, users, and communities, not written only by executives. If the company's scientific ambitions produce public value, a Restorying process should be able to name and document that value. If the current account leaves out cost-bearers, the process should bring those accounts into the same inquiry.

Externalizing the problem is especially important in a case organized around a highly visible founder. The question is not "What is wrong with Musk?" but "What pattern is the organization reproducing, what does it reward, and how could decision processes change?" A pattern such as compressed deployment cycles without sufficient downstream feedback is only a provisional label. Participants may identify a different pattern, or find that the pattern occurs in one domain and not another. Externalization keeps the inquiry open to a structural explanation and reduces the risk that a leader-centered story will obscure systems, incentives, and collective agency.

Sympathizing asks what the current story does for the organization. Vertical integration may provide technical control, reduce coordination delays, and create a distinct product

experience. A rebellious brand may help attract employees and users who distrust conventional technology companies. These payoffs can be real even when the pattern also produces costs. A useful intervention preserves what is valuable—ambition, engineering speed, and willingness to test assumptions—while changing the arrangements that make harm invisible or difficult to contest.

The revise and strategize steps require a demanding standard for counterevidence. An exception is not a Little Wow Moment merely because the company announced a new safeguard or published a community-engagement statement. The exception must be observable, affect a decision, include relevant stakeholder voice, and persist beyond the immediate publicity cycle. Examples could include a product restriction that measurably reduces a documented risk, a permitting decision that responds to residents' evidence, or a governance change that gives affected stakeholders continuing influence. Those examples are criteria for future investigation, not claims that xAI has already met them.

Rehistoricizing should preserve conflict instead of resolving it by authorial choice. The 2023 "understanding and knowledge" mission, the company's "rebellious" product identity, the later merger sequence, the current vertically integrated segment, the Colossus dispute, and the regulator-engaged restrictions belong in one chronology. The chronology should distinguish what the company said, what outside actors alleged, what regulators investigated, and what courts decided. Publicizing a new story would then require more than a revised mission statement. It would require commitments with named owners, dates, evidence of performance, access to independent review, and public explanation when a target is missed.

5. Provisional Concord Tesseract™ Reading

The four coordinates use the book's behaviorally anchored seven-point scales: X moves from transactional control toward transformational development; Y moves from power over others toward serving others; Z moves from a single dominant voice toward polyphonic interaction among voices; and A moves from an alibi in Being toward moral answerability. A score is a provisional coding position on its axis, not a stand-alone measure of virtue. The dimension definitions and scoring rules remain those established in the Preface and method chapter.

I carry forward the latest workbook baseline for Musk/xAI: X=5, Y=1, Z=2, and A=NE. These are provisional archival positions for the book's APP demonstration, not validated

measures of the whole organization. The Evidence Log currently has N=0 for this case. Under the project's rule, A=NE is not zero; because A is not scored, the multiplicative composite M is not calculated. I do not substitute earlier additive totals or earlier A scores from other drafts. The A-cube's three graphic dimensions—content and sense, situated eventness and ownership, and integration into action—are visual organizing dimensions only; they are not three additional scored variables.

Dimension	Provisional position	Reason and limit
X: Transactional ↔ Transformational	5/7	The public account emphasizes ambitious technological change and large-scale system building. This provisional score does not establish that transformation is participatory, sustainable, or beneficial to all affected groups.
Y: Power over ↔ Serving others	1/7	The low-end position reflects concentrated founder control in the public and corporate record, including integration of platforms and infrastructure. It should be tested against internal governance evidence and accounts from workers, users, and affected communities.
Z: Monophonic ↔ Polyphonic	2/7	The public company narrative is highly legible and leader-centered; counter-stories are visible in regulatory and community records but their influence on decisions is not yet established. This score describes the available record, not the number of people with opinions.
A: Alibi in Being ↔ Moral answerability	NE	The current evidence log does not yet contain a qualifying, fully specified answerability episode under the project's evidence gate. NE is not zero and is not a claim that answerability is absent.
M: Multiplicative composite	Not calculated	M is calculated only when all four dimensions are evidence-supported and scored. No substitute arithmetic is used.

These scores should be revised when the evidence log contains dated episodes, source provenance, coding decisions, confidence notes, and counterevidence for each dimension. The Y and Z positions are particularly sensitive to missing voices. A corporate filing and executive statement can show formal structure and public positioning; neither can tell us whether employees or affected communities experienced decision-making as dialogic or answerable. The first independent-coder pass should therefore include different source classes and explain which

observations move a coordinate. The score should remain provisional until that record is assembled.

Figure 5.1 belongs here: the book's grayscale ten-leader XYZA navigation map, with Musk's current point highlighted at the workbook position and the A status shown as NE. The figure should not imply that the three A-cube drawing axes are scored. The map should be updated from the cumulative workbook, retaining the other leaders' agreed coordinates and clearly distinguishing a provisional starting position from any future evidence-supported movement. I have not generated a new final graphic in this pass because the expert image-generation prompt is still forthcoming.

6. What Would Count as Organizational Change?

Restorying is useful here only if the new story can be connected to changes in work and governance. The intervention would need to be evaluated through observable outcomes: whether risk review precedes release; whether people affected by product or infrastructure decisions can shape them; whether complaints produce timely remedies; whether energy and emissions claims are independently verifiable; whether decision-makers can explain and revise choices; and whether commitments survive pressure to move quickly. A narrative shift may be a necessary part of change, but it is not sufficient evidence that the material conditions have changed.

A pilot could begin with a bounded decision episode rather than a whole-company transformation claim. For example, a cross-functional group could reconstruct one product-safety decision or one infrastructure-permitting episode from the first proposal through review, deployment, complaint, and correction. The inquiry would map who spoke, who was absent, what evidence moved the decision, what alternatives were considered, and which consequences were recognized. A second cycle would test a redesigned process and compare its timing, participation, and outcome. Any such study would require company access and voluntary participation; no such access or intervention is claimed in this chapter.

A restorative ODC approach also needs safeguards against capture. If the company sponsors the process, participants must know who controls the record and whether dissent can be published. Residents and users should not be asked to contribute unpaid stories that only enrich corporate learning. Confidentiality, compensation, consent, and data retention should be negotiated before collection begins. Independent facilitation and participant review of the

resulting account would help prevent the Restorying process from becoming a new corporate master narrative with more inclusive language but unchanged authority.

7. Evidence Gates and Self-Correction

This chapter is an invitation to revise the theory and the scores. I would change the analysis if internal governance records showed that affected stakeholders routinely influenced high-stakes decisions, if independent performance evidence demonstrated timely repair across the relevant risk areas, or if additional source comparison showed that my current reading overstates the unity of the xAI/SpaceX system. I would also revise the X, Y, or Z positions when higher-quality evidence warrants it. A qualifying A score requires a specific episode in which an accountable actor addresses a particular other, owns a consequential choice, and integrates that answerability into action; symbolic language alone is not enough.

The chapter also needs further research before publication. The public record is uneven: company sources describe products and structure; regulators describe duties and investigations; plaintiffs and community organizations describe harms and seek relief; court orders determine legal questions at particular procedural stages. The next research pass should add Musk's dated statements on xAI's purpose and governance, xAI's documented safety and release decisions, first-party company responses to the turbine dispute, and accounts from residents and workers. The author's additional AI-generated research notes can serve as leads, but each factual claim must be checked against the underlying record and disagreements preserved.

My conclusion is therefore bounded. Restorying is the strongest present ODC candidate for xAI because the case needs a method that can surface a dominant corporate story, externalize system patterns rather than pathologize a leader, recover exceptions, and bring counter-stories into a public process of change. Whether it would work in this organization remains an empirical question. The current evidence supports a provisional analysis of concentrated control, material scale, and contested answerability; it does not support a definitive account of Musk's private motives or a completed assessment of xAI's internal culture. The project's purpose is to make these interpretations testable and revisable.

References

Boje, D. M. (2019, February 7). *Welcome to restorying*. David M. Boje. https://davidboje.com/restorying/

European Commission. (2026, January 26). *Commission investigates Grok and X's recommender systems under the Digital Services Act*. https://digital-strategy.ec.europa.eu/en/news/commission-investigates-grok-and-xs-recommender-systems-under-digital-services-act

National Association for the Advancement of Colored People v. X.AI Corp., No. 3:26-cv-00074-DMB-JMV (N.D. Miss. Oct. 30, 2026) (Opinion and Order). https://docs.justia.com/cases/federal/district-courts/mississippi/msndce/3%3A2026cv00074/52261/139

Ofcom. (2026, October 14). *Ofcom update: Monitoring X's response to hate and terror content, and Grok investigation latest*. https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/ofcom-update-monitoring-xs-response-to-hate-and-terror-content-and-grok-investigation-latest

PERVIEW Inc. (n.d.). *PERVIEW Inc.: 7-step Restorying method*. https://perview.org/

Space Exploration Technologies Corp. (2026). *Quarterly report on Form 10-Q for the quarter ended September 30, 2026*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1181412/000162828026066880/spcx-20260930.htm

SpaceXSI. (2026, July 30). *Greater Memphis area site updates*. https://x.ai/memphis/updates

xAI. (2023, November 3). *Announcing Grok*. https://x.ai/news/grok

xAI. (2024a, May 26). *Series B funding round*. https://x.ai/news/series-b

xAI. (2024b, December 23). *xAI raises $6B Series C*. https://x.ai/news/series-c