

C.A.S. Full Assessment Report

“X” — x.com

Generated: September 14, 2026, 9:58 AM EDT · Assessment ID: 6aa7e5d3-4a07-b277-b365-5ba2

Author	Delante Clark
Affiliation	University of Central Florida
System	Cognitive Attentional Standards (C.A.S.) Platform
System Designer	Delante Clark
System Components	Attention Safeguard Models (A.S.M.), Sociological Impact Assessment (S.I.A.)
Assessment Year	2026
Subject	“X” (x.com)
Mode	Research Mode (Non-Certifying)

Suggested citation: Clark, D. (2026). *Cognitive Attentional Standards Platform (C.A.S.) Assessment Report: “X.”* Attention Safeguard Models and Sociological Impact Assessment Output.

1. Assessment Details

Organization	“X”
System / Product	x.com
System Type	Platform
Primary URL	https://x.com
Assessment Type	Research
Status	Pending Approval
License Status	Inactive

Scope Description

This assessment evaluates X.com through the Cognitive Attentional Standards (C.A.S.) framework to measure how platform architecture, algorithmic systems, engagement mechanics, and user-experience design influence attention, cognition, behavior, information consumption, and digital well-being. The assessment does not focus solely on content-moderation outcomes; it examines the underlying systems and design choices that shape user interaction and

behavioral patterns across the platform.

The assessment analyzes X.com's recommendation ecosystem, including the "For You" feed, trending topics, search, algorithmic ranking, content amplification, promoted content, monetization, creator incentives, and engagement-optimization strategies. Particular attention is given to how these systems influence attention allocation, session duration, interaction frequency, content-exposure diversity, and behavioral reinforcement cycles.

The review also examines design features that may contribute to attention extraction, compulsive engagement, emotional activation, information overload, or cognitive fatigue: notifications, infinite scrolling, engagement metrics, reposting mechanics, quote posting, media presentation, algorithmic prioritization, community discussions, livestreams, subscription features, and real-time event amplification. The assessment evaluates whether these systems encourage informed engagement or incentivize reaction-driven behavior that increases user dependence on the platform.

Assessed Scope

- Algorithmic feeds and recommendations
- Content ranking and amplification systems
- Trending topics and search functions
- Posts, reposts, quote posts and replies
- Engagement features (likes, bookmarks, follows)
- Notifications and user-retention mechanisms
- Advertising and sponsored content
- Creator monetization and premium features
- Video, media and live-content systems
- Communities, Spaces and messaging features
- User controls, safety tools and content filters
- Platform transparency and account-management features
- Mobile and web user experience
- Information distribution and virality mechanisms
- Behavioral reinforcement and attention-retention features

Explicitly Excluded from Scope

Internal source code and proprietary algorithms not publicly accessible; employee communications and internal business operations; third-party systems outside X.com's direct control; financial performance, revenue, or stock-valuation analysis; individual user-conduct investigations; law-enforcement or government activity on the platform; external linked websites; post-by-post content fact-checking; cybersecurity penetration testing; infrastructure and network security architecture; legal-compliance or regulatory-enforcement determinations; employment and HR operations; features unavailable during the assessment period; and activity occurring outside the X.com ecosystem.

Assessment Contributors

delante@asmcas.com (Admin) — finalized SIA, saved SPI/AIS/AEA scores, created assessment.

2. Executive Summary

This research assessment applied the Cognitive Attentional Standards (C.A.S.) framework to X.com, combining an automated 200-page preliminary scan with formal AEA/AIS/SPI scoring, a Sociological Impact Assessment (SIA), ASM pressure simulation, an independent algorithm-detection scan, an economic impact analysis, and an EU Digital Services Act (DSA) systemic-risk module.

Overall, X.com presents a **high-pressure, low-equity** attention profile — the most severe pattern of the platforms assessed under this program to date. Automated and assessor-reviewed scoring converge on the same conclusion: attention-supportive design is largely absent, attention burden is unevenly distributed across users, and engagement-extractive pressure is high. The SIA further finds meaningful risk for age-sensitive cohorts, lower-bandwidth and lower-resource users, and racialized structural harm through surveillance-weighted and centralized-access design.

ATTENTION EQUITY (AEA)	ATTENTION INEQUITY (AIS)	SAFEGUARD PRESSURE (SPI)
16.5 High Inequity	67.9 Moderate–Significant	75.8 High Coercive Pressure

Reading the scores: for AEA, higher is better (≤ 40 = High Inequity); for AIS, higher is worse (41–70 = Moderate Inequity, and 67.9 sits only 2.1 points from the Significant Inequity threshold); for SPI, lower is better, and scores above 70 indicate High Coercive Pressure.

Preliminary Automated Scan Findings

An automated crawl of 200 pages (generated Sep 14, 2026, 9:14 AM EDT) identified five structural signals prior to formal scoring:

Domain	Signal	Finding
AIS	Mobile vs. Desktop Differences	200 of 200 pages missing the viewport meta tag, potentially degrading the mobile experience and causing unequal access to content.
SPI	Missing Clear Instructions	200 of 200 pages lack visible help or instructional content for navigation or core functionality.
SPI	Unpredictable Navigation	200 of 200 pages lack a semantic <nav> element or navigation landmark, potentially difficult to predict for neurodivergent users.
SPI	No Clear Exit or Pause Options	200 of 200 pages do not appear to offer distinct exit, cancel, or pause functionality for interactive components.
SPI	No Time Expectations Communicated	Time-related cues (e.g., estimated read time) were not identified on any of the 200 scanned pages.

Recommended next steps (from the preliminary scan): conduct a deep-dive manual accessibility and cognitive-impact audit focused on the identified navigation and instruction gaps — particularly within dynamic sections such as the main feed — and commission a full Sociological Impact Assessment. Both recommendations are addressed in this report (Sections 3–4).

3. C.A.S. Scoring — AEA / AIS / SPI

Assessor	delante@asmcas.com
Date	September 14, 2026
Highest Burden Group	Racially marginalized, low-resource, and digitally constrained users — especially users relying on mobile access

Attention Equity Assessment (AEA) — 16.5 / 100 · High Inequity

Any AEA score at or below 40 falls into High Inequity, and X.com scores well within that range. This indicates very limited evidence that the platform's design distributes attentional demand in a balanced, user-supportive, or cognitively equitable manner — the architecture places a disproportionate share of pacing, filtering, and cognitive-burden management on the user rather than embedding those protections into the system. This is consistent with a platform environment dominated by high-volume, rapidly updating streams; socially comparative visibility indicators; recommendation-driven discovery; and multi-surface engagement loops spanning feed, trends, search, notifications, communities, media, and live interaction. Promoted content and monetization-linked visibility, quote-post/repost mechanics, and real-time event amplification further concentrate this inequity, and the preliminary scan's absence of positive AEA signals reinforces that no strong pacing controls, simplified task structures, or default interruption safeguards were evident.

Attention Inequity Score (AIS) — 67.9 / 100 · Moderate Inequity (near-significant)

67.9 places X.com in the Moderate Inequity band, only 2.1 points below the Significant Inequity threshold — a substantial concentration of attention-related burden, navigational inconsistency, and uneven participation conditions across devices, workflows, and user cohorts. The preliminary scan's Mobile vs. Desktop Differences signal indicates cross-platform parity is not fully consistent, which in a feature-dense system produces uneven discoverability of controls and differential cognitive load depending on device. The SIA record substantiates this: dense, multi-path, recall-dependent navigation (AIO-O1); abstract, non-persistent workflows (DD-2A, DD-2C); and resource-dependent success pathways tied to premium features and paid visibility (DD-3B, DD-3C) collectively distribute effort unevenly — users with stronger digital fluency, better connectivity, more time, or the ability to pay for visibility are likely to navigate and benefit from X.com's systems more effectively than users without those advantages.

Safeguard Pressure Index (SPI) — 75.8 / 100 · High Coercive Pressure

Scores above 70 represent High Coercive Pressure: the platform's design exerts strong pressure toward re-entry, continuation, reaction, or prolonged session persistence. This is the clearest indicator of acute risk in the assessment. The preliminary scan's four SPI signals — missing instructions, unpredictable navigation, absent exit/pause controls, and no time expectations — map directly onto this result: pressure is produced not only by notifications or feed ranking, but by users being unable to easily understand where they are, how long a process will take, how to pause, or how to safely disengage without missing context or incurring social disadvantage. On a real-time social platform this can turn ordinary participation into a sustained vigilance model — reinforced by the SIA's AIO-A1 finding (urgency-framed notifications coupled with socially visible engagement signals, producing anxiety and compulsive engagement) and AIO-Y1 finding (persistent engagement expectation producing burnout and stress normalization).

Overall Domain Interpretation

X.com's C.A.S. scoring profile shows an unfavorable attention-governance pattern: weak attention-equity protections, substantial unevenness in how attention burden and participation conditions are distributed, and strong engagement-extractive conditions with insufficient default pacing or disengagement protection. The overall combination reflects a low-equity, high-pressure platform environment with meaningful cognitive-accessibility concerns.

4. Sociological Impact Assessment (SIA)

SIA Mode	Mode A — Include SIA
Status	Complete (finalized Sep 14, 2026)
Assessor	delante@asmcas.com
AER Exemptions	None recorded

SIA findings do not alter certification tier, but provide critical contextual evidence regarding who is most affected by platform design and how harms may be distributed.

Age Impact Triggers (AIO)

AIO-A1 — Urgency & Social Visibility: Urgency-framed notifications combined with socially visible engagement signals and ranking-style visibility mechanics. *Harm:* Anxiety, compulsive engagement, punitive exposure.

AIO-O1 — Interface Density: Dense interface with multi-path navigation and recall-dependent interaction across feed, search, trends, communities, and account features. *Harm:* Errors, exclusion, administrative failure.

AIO-Y1 — Persistent Engagement Expectation: Continuous alerts, real-time events, and ongoing notification-driven re-entry. *Harm:* Burnout, stress normalization.

Collectively, these findings indicate that the platform's pace, signaling architecture, and interface density may create disproportionate burden for age-sensitive cohorts and users vulnerable to social-pressure design.

Digital Divide Flags (DDIO)

DD-1A — Digital-Only Access: Mandatory digital-only workflows for core platform functions. *Harm:* Inability to access essential functions; forced exclusion.

DD-1B — Bandwidth Dependency: High-speed-internet assumptions for video, livestream, media, and real-time participation. *Harm:* Persistent access failure; degraded participation.

DD-2A — Workflow Complexity: Abstract interaction logic without guided support. *Harm:* Errors, task failure, process abandonment.

DD-2C — Non-Persistent Workflows: Multi-step, memory-dependent tasks with no state persistence. *Harm:* Incomplete processes, administrative errors.

DD-3B — Resource-Gated Participation: Premium features and paid visibility create resource-dependent participation advantages. *Harm:* Unequal benefit distribution despite equal participation.

DD-3C — Cascading Disadvantage: Failures or penalties compound across connected recommendation, monetization, and participation systems. *Harm:* Cumulative exclusion and long-term disadvantage.

Together, these findings show the platform's burdens are not evenly experienced and may intensify for users with lower bandwidth, lower digital fluency, fewer financial resources, or less tolerance for continuous process management.

Racialized Structural Risk (RSIO)

RSIO-R2 — Surveillance-First Logic: Pervasive monitoring, visibility metrics, and audit-like engagement/compliance systems. *Harm:* Elevated sanction risk, mistrust.

RSIO-R3 — Digital Monoculture: Centralized digital enforcement with no offline or human intermediary access path. *Harm:* Systematic exclusion.

These findings suggest that monitoring-heavy and centralized enforcement structures may deepen mistrust and inequitable outcomes, especially where assistance pathways are less visible than control or sanction pathways.

Required Institutional Fixes (TRE)

Eleven remediation tracks were generated across three levels of intervention:

Finding	L1 — Design	L2 — Structural	L3 — Governance
AIO-A1	Remove/reduce urgency-framed notification language and de-emphasize socially comparative prompts in adolescent-accessible surfaces	Redesign notification and ranking architecture to decouple visibility metrics from re-engagement demands	Adopt platform-wide standards limiting urgency-based social pressure for adolescent-accessible features
AIO-O1	Introduce clearer in-flow instructions, persistent cues, and simplified navigation labels	Simplify information hierarchy; add persistent state-saving across multi-step workflows	Require interface-density and memory-burden review for major navigation/settings releases
AIO-Y1	Batch notifications by default; reduce after-hours alert intensity	Create asynchronous participation pathways that do not penalize delayed response	Establish notification-governance standards prohibiting continuous-attention assumptions
DD-1A	Provide lightweight web-access modes for core information and participation	Develop alternative lower-connectivity access pathways for core functions	Adopt access-equity requirements for non-exclusionary core-feature availability
DD-1B	Add low-bandwidth media defaults, reduced-data playback, static alternatives	Architect media/live/real-time systems for degraded-network participation	Institute bandwidth-equity design review for all new media/live-feature launches
DD-2A	Add visible help, instructions, and guided-support cues across core workflows	Redesign complex interaction flows to reduce abstraction and step burden	Require guided-support and usability-burden review before deployment
DD-2C	Add state persistence, progress indicators, resumable multi-step workflows	Rebuild key workflows to minimize cross-session memory dependence	Mandate persistence and memory-burden checks in design assurance review
DD-3B	Clearly separate paid amplification from organic relevance indicators	Redesign benefit pathways so participation advantages don't depend on surplus resources	Implement equity review for monetization/premium feature launches
DD-3C	Add warnings/recovery prompts before compounding visibility or participation losses	Introduce recovery windows; decouple penalties from single-point failures	Create no-fault attention-failure and escalation-override policies
RSIO-R 2	Reduce prominence of monitoring-style cues; increase visible assistance pathways	Rebalance enforcement so assistance precedes punitive/visibility-restricting action	Establish oversight requiring review of surveillance-weighted features
RSIO-R 3	Provide lower-friction, lower-data participation pathways and clearer human-readable help	Reduce centralized digital monoculture via alternative support/access channels	Adopt access-governance standards requiring mitigation of clustered digital exclusion risk

The SIA supports a broad conclusion that meaningful cognitive-accessibility improvement at X.com will require both interface-level changes and higher-level product-governance controls.

5. ASM Scoring — Attention Safeguard Model

Source: ASM Scoring Dashboard — platform pressure inputs and simulation scores linked to this assessment. Three linked simulation runs were recorded:

Target	ASM Score	API Pressure	Safeguard Freq.	Load Volatility	Recovery	Behaviors
Website — x.com	40.5 / 100	41.9 / 100	4.3 / 10	0.2 / 10	8.2 / 10	Selective exposure, muting
Platform (X) — run 1	67.3 / 100	69.1 / 100	7.3 / 10	0.5 / 10	6.0 / 10	Selective exposure, muting, avoidance
Platform (X) — run 2 (duplicate)	67.3 / 100	69.1 / 100	7.3 / 10	0.5 / 10	6.0 / 10	Selective exposure, muting, avoidance

All runs are C.A.S. Cycle 5, scored September 14, 2026 (website run 9:28 AM EDT; platform runs 9:27 AM EDT). Linked social-pressure inputs: X (“X”) at API pressure 86.0 (assessment-import), and the x.com website at API pressure 35.9.

Historical runs: two terminated runs from May 2026 (RUN-1780048045024, RUN-1779997056280) recorded a Final ASM Score of 15/100 (“Meets ASM Standards”) but remain in review-pending status and are not aligned with the current completed simulation set. They should be treated cautiously and not weighed against the completed September 2026 evidence below.

ASM Narrative

The website-targeted run (40.5/100, Recovery 8.2) indicates limited recovery capacity within the tested simulation context: some self-protective user behaviors are available, but they do not fully offset the platform's broader attention and exposure dynamics. The two platform-input runs (67.3/100, Recovery 6.0) suggest a more impaired outcome profile at the platform level — a higher final ASM score paired with lower recovery resilience. The recurring annotations (selective exposure, muting, avoidance) indicate users may rely on defensive adaptation strategies rather than receiving sufficient baseline protection from platform design. The gap between the website and platform-input results also suggests inconsistent safeguard performance across surfaces or test conditions, warranting separate quarterly testing for each.

Algorithm Detection Report — 62 / 100 Overall Risk (Medium)

An independent automated structural scan of 15 pages found X.com to exhibit a medium-to-high risk profile. Detected patterns included autoplay media (6 pages), forced-continuity mechanisms (7 pages), variable-reward systems (4 pages), attention-capture mechanisms (4 pages), and push-notification prompts (4 pages). An SIA-linked sociological impact score of 58/100 corroborates a meaningful manipulation risk. X.com's core feed algorithm, notification architecture, and content-delivery mechanisms align with established research on social-media persuasive design.

Attention Burden Score: 65 / 100 (substantial)

The combination of autoplay media, algorithmically unpredictable content, and persistent notifications creates a high-frequency interrupt environment. The variable-reward structure artificially elevates the cognitive cost of disengagement — users cannot easily determine when it is “safe” to stop without potentially missing valued content — and the forced-continuity pattern (7 of 15 pages) suggests the platform actively resists natural session-termination points.

- **Autoplay Media Content** (High) — video and media autoplay without explicit initiation, hijacking attention and extending session duration beyond user intent.
- **Forced Continuity Engagement Architecture** (High) — persistent session states, algorithmic feed refresh, and content structures that resist natural stopping points.
- **Variable Reward Feed Algorithm** (High) — the algorithmically curated feed delivers unpredictable, variable-quality rewards mirroring slot-machine reinforcement schedules.

- **Attention Capture Mechanisms** (High) — notification badges, trending topics, and engagement metrics engineered to trigger curiosity and FOMO responses.
- **Push Notification Coercion** (Medium) — prompts establish persistent re-engagement channels, often before users have sufficient context for informed consent.
- **Countdown Timers** (Medium) — introduce artificial time pressure that reduces deliberative thinking and promotes impulsive action.
- **Social Proof Display** (Medium) — likes, reposts, and follower counts are prominently displayed to create social-validation loops.
- **JavaScript Dependency Gate** (High) — core functionality and legal documents (e.g., sitemap) are blocked without JavaScript, limiting autonomy and privacy-tool usage; messaging subtly blames privacy extensions.
- **Broad Consent Bundling at Onboarding** (Medium) — Terms of Service, Privacy Policy, and Cookie Use are bundled into a single consent action with no granular opt-out.

Immediate: remove or neutralize anti-privacy-extension messaging; require affirmative opt-in for autoplay; decouple ToS/Privacy/Cookie consent into separate actions. **Short term:** audit countdown timers for genuine vs. artificial scarcity; add real session-boundary features (usage dashboards, “end of feed” indicators); offer a persistent non-algorithmic chronological feed default. **Long term:** commission an independent, published algorithmic-impact assessment of the “For You” reinforcement schedule, and redesign the notification consent flow to occur after meaningful engagement with clear frequency controls.

6. Economic Impact Analysis

Generated September 14, 2026, 9:26 AM EDT · Scores used — AEA 17 · AIS 68 · SPI 76 · CCS 24

Traffic & Revenue Inputs		Revenue at Risk	
Monthly Visits: 1,500,000,000 Avg. Session Duration: 963s Bounce Rate: 35% Conversion Rate: 2% Avg. Revenue / Conversion: \$3.50 Avg. User Hourly Wage: \$30.00		Daily: \$154,067,586.25 Weekly: \$1,078,473,103.75 Monthly: \$4,622,027,587.50 Quarterly: \$14,020,150,348.75 Annual: \$56,234,668,981.25	
Effective Bounce Rate (AIS-amplified)	44.51%		
Daily Suppressed Conversions (SPI)	189,375.00		
Dominant Loss Vector	Session Depth Loss		

7. DSA Risk Detection — EU Digital Services Act Module

Run Date	September 14, 2026, 9:32 AM EDT
Methodology Version	DSA-RISK-V1.0
Input Mode	Crawl
Assessed Target	X — formerly Twitter · https://x.com
Run By	delante@asmcas.com

TOTAL SYSTEMIC RISK

71 / 100

High Risk

Total = $IRS \times 0.25 + MRS \times 0.25 + UARS \times 0.20 + SIRS \times 0.20 + AS \times 0.10$. Classification bands: Low Risk 0–30, Moderate Risk 31–60, High Risk 61–80, Severe Systemic Risk 81–100.

Systemic Risk Profile — Category Breakdown

Category	Score	Weight	Composition
IRS — Information Risk	75 / 100	25%	Virality rate (35%), amplification bias (40%), exposure skew (25%). Measures how strongly the system distorts what information reaches a user.
MRS — Manipulation Risk	73 / 100	25%	Algorithmic weight (40%), engagement-loop strength (35%), recommendation dependency (25%). Measures how much the system steers behavior rather than serving intent.

Category	Score	Weight	Composition
UARS — User Autonomy Risk	57 / 100	20%	Default bias (30%), control access (35%), opt-out barriers (35%). Measures how far defaults and controls constrain a user's real choice.
SIRS — Societal Impact Risk	77 / 100	20%	Content clustering (30%), echo-chamber score (35%), exposure divergence (35%). Measures population-level fragmentation and polarization pressure.
AS — Addictiveness	72 / 100	10%	Unweighted mean of scroll trigger, autoplay, infinite feed, and session duration. The most direct attention-harm signal.

Attention Harm Findings (Primary C.A.S. Outcome)

- **Compulsion-forming design pressure (AS) — 72/100 (elevated):** from scroll trigger 78, autoplay 60, infinite feed 80, and session duration 68. Higher values mean the interface itself sustains attention beyond user intent.
- **Algorithmic steering of attention (MRS) — 73/100 (severe):** attention is allocated by the recommendation system rather than by the user, so exposure is a product of ranking, not choice.
- **Loss of attentional autonomy (UARS) — 57/100 (elevated):** defaults, control access, and opt-out friction determine how much attentional control the user actually retains.

DSA Compliance Flags

- High algorithmic amplification detected
- High engagement reinforcement patterns
- Elevated societal polarization risk
- Platform-level systemic risk alert
- Severe multi-domain risk — immediate review required

Methodology: the DSA module scores 16 structural risk variables (virality rate 75, amplification bias 78, exposure skew 70, algorithmic weight 76, engagement-loop strength 72, recommendation dependency 68, default bias 65, control access 50, opt-out barriers 58, content clustering 78, echo-chamber score 82, exposure divergence 72, scroll trigger 78, autoplay 60, infinite feed 80, session duration 68) into the five systemic-risk categories above and a weighted total, using this assessment as its source context. DSA output is a governance-facing view of the same attention-harm evidence; it does not alter C.A.S., ASM, or SIA scoring and does not constitute legal advice.

8. C.A.S. Improvement Roadmap

X.com's current research assessment shows a high-risk cognitive-attention profile driven by low attention equity (AEA 16.5), elevated attention inequity (AIS 67.9), and high coercive pressure (SPI 75.8). The most urgent priorities are to reduce re-entry pressure from notifications and real-time ranking systems, simplify dense and recall-dependent navigation across feed/search/trends/settings pathways, and strengthen user control over pace, visibility, and exposure. The findings point to an architecture that currently rewards continuous vigilance, rapid reaction, and high-frequency engagement rather than informed, self-paced participation.

AEA — Attention Exposure

- 1 Replace the default “For You” landing experience with an explicit first-run choice between algorithmic, following-only, and chronological feeds.
- 2 Introduce session-level exposure-balancing controls (“show less of this theme today,” “increase source variety”).
- 3 Create a persistent “attention settings” hub for feed pace, autoplay defaults, trend visibility, media density, and ranking intensity.
- 4 Add transparent labels explaining why posts, trends, and promoted items are shown (engagement velocity, paid promotion, history, or network proximity).
- 5 Reduce prominence of socially comparative engagement counts by default, especially in discovery and quote-post surfaces.
- 6 Offer a low-stimulation mode: reduced motion, reduced visual competition, limited live-update injections, hidden trend modules, finite feed segments.
- 7 Rebalance recommendation objectives to include diversity, comprehension support, and user-declared intent completion — not primarily session duration and re-entry probability.

AIS — Attention Inequity

- 1 Standardize cognitive-access parity between desktop and mobile for feed, notification, safety, monetization, search, and trend controls.
- 2 Simplify multi-path navigation across feed, search, trends, communities, messages, Spaces, and account settings with persistent wayfinding cues.
- 3 Add state persistence and resumable flows for settings, reporting, filtering, account recovery, monetization setup, and premium management.
- 4 Separate paid amplification and premium visibility from organic ranking with unmistakable visual and textual distinction.
- 5 Provide lightweight, lower-bandwidth variants: text-first feed loading, static trend summaries, click-to-load media, low-data Spaces/video.
- 6 Introduce guided-support overlays and plain-language instructions for posting, muting, notifications, feed switching, reporting, and visibility controls.
- 7 Create recovery pathways — explanations, remediation steps, review options — before reach or monetization loss compounds across systems.
- 8 Offer human-readable help and alternative support access for account restrictions, safety settings, or participation barriers.

SPI — Safeguard & Protection

- 1 Batch non-essential notifications by default; add quiet-hour protection for likes, reposts, follows, and trend alerts.

- 2 Remove urgency-framed notification language implying social loss, visibility decline, or missed-event penalties.
- 3 Introduce clear pause, stop, and exit controls in infinite feed scroll, autoplay video, live events, Spaces, and notification streams.
- 4 Replace endless-feed behavior with optional finite session units (“you’re caught up” states, feed caps, session summaries).
- 5 Display time and interruption expectations before entering live, long-thread, or high-volume surfaces.
- 6 Make chronological/following-based consumption sticky across sessions rather than repeatedly defaulting back to algorithmic ranking.
- 7 Add a “slow mode” that reduces real-time injections and sudden ranking reshuffles during active reading.
- 8 Require every new retention mechanic, notification type, or ranking experiment to pass a coercive-pressure review before launch.

SIA — Sociological Impact

- 1 Remove urgency cues and comparative social prompts from notification surfaces; decouple ranking visibility from delayed participation (AIO-A1).
- 2 Reduce interface density by consolidating duplicate pathways and adding persistent orientation cues (AIO-O1).
- 3 Establish asynchronous participation norms so delayed replies or infrequent posting don’t reduce discoverability or social standing (AIO-Y1).
- 4 Provide lightweight web-access modes for core reading, posting, following, and safety actions (DD-1A).
- 5 Add low-bandwidth defaults — text-first loading, disabled autoplay, static substitutes for livestream/event modules (DD-1B).
- 6 Embed visible help and guided support directly into account, ad, premium, and reporting workflows (DD-2A).
- 7 Add progress indicators, autosave, and resumable multi-step workflows for settings, safety, verification, and monetization flows (DD-2C).
- 8 Ensure premium features and paid visibility don’t materially signal superior relevance or trust in discovery systems (DD-3B).
- 9 Introduce warnings, grace periods, and recovery windows before penalties propagate across connected systems (DD-3C).
- 10 Reduce surveillance-first cues; replace with explanation, assistance, and review pathways before punitive action (RSIO-R2).
- 11 Create alternative, human-intermediated support channels for users facing access barriers or platform confusion (RSIO-R3).

ASM — Attention Safeguard Model

- 1 Create explicit recovery tools after overload events: “reset recommendations,” “reduce topic intensity,” “pause suggested accounts,” “rebuild feed from follows only.”
- 2 Expand muting into a more powerful, accessible control suite (one-tap durations, topic-cluster muting, quote-post muting, trend suppression).
- 3 Add guided diversity controls (“more source diversity,” “less repetitive viewpoint clustering”) instead of only binary suppression.
- 4 Make discovery safer and more predictable with recommendation-reason previews and topic-intensity indicators before active threads or live discussions.

- 5 Create a persistent “recovery mode” that disables autoplay, hides trends, limits notifications, and prioritizes followed accounts chronologically.
- 6 Run quarterly ASM scenario testing separately for web and mobile, given the website (40.5) vs. platform (67.3) score gap.
- 7 Investigate the discrepancy between earlier terminated runs (15/100, “Meets Standards”) and current completed runs; standardize scope and publish a single authoritative remediation baseline.

9. C.A.S. v2 Deterministic Scoring

Scored September 14, 2026, 9:32 AM EDT (automated deterministic pass; confidence not yet established).

Domain	Score
AEA	71.5
AIS	80.1
SPI	86.8
CCS (Composite)	75.0

Note: this automated deterministic pass returns notably higher-risk domain scores than the assessor-reviewed scoring in Section 3. Both passes point in the same direction — low equity, high inequity, high pressure — but the deterministic engine has not yet incorporated assessor override and confidence weighting. The Section 3 scores (AEA 16.5, AIS 67.9, SPI 75.8) are the assessor-of-record figures for this report.

Report Status & Composition Scope

Report State	Ready to finalize
Scoring Coverage	Not determined
Aggregate Score Status	Not available
Results Current	Yes
Scope	This report contains only current, non-archived artifacts linked to this assessment.
Artifacts Included	0 requirement results · 0 statutory findings · 0 remediation tasks · 0 citation review items · 0 jurisdiction rollups

Evidence Log

Seven supporting evidence files were logged against this assessment and are retained in the assessment record:

Screenshot — Sep 14, 2026, 8:57 AM (X.com “For You” feed) — Type: Design · Domain: AEA

Demonstrates the platform's engagement-driven content distribution: algorithmic recommendations, trending topics, promoted content, and follower suggestions presented simultaneously. Content observed during the assessment included interpersonal conflict, sexually suggestive discussion, and identity-related commentary positioned alongside ads and news — illustrating how engagement-focused systems can amplify provocative or emotionally charged content that attracts interaction regardless of whether the resulting experience is constructive.

Screenshots (3) — Sep 14, 2026, 8:54–8:55 AM (creator monetization eligibility criteria) — Type: Governance · Domain: AIS

X.com's monetization framework requires account age, identity/email verification, two-factor authentication, good standing, an active Premium subscription, and follower/impression thresholds. Eligibility metrics such as follower counts and impression thresholds inherently reward visibility and engagement, which may incentivize reaction-driven or controversy-generating content since audience growth is directly tied to monetization access. Enforcement and eligibility determination rely heavily on automated systems, which can produce inconsistent outcomes.

Screenshots (3) — Sep 14, 2026, 8:54–8:55 AM (Hateful Conduct enforcement policy) — Type: Policy · Domain: AEA

X.com publishes enforcement options for policy violations (down-ranking, discoverability restriction, engagement-count limits, ad exclusion, post removal, suspension). Enforcement appears to rely heavily on automated detection rather than direct human review, creating risk that harmful content is not reliably identified (false negatives) or that enforcement is perceived as inconsistent.

10. Attached Assessment Comparison

One comparative analysis is attached to this assessment: “X” vs. HBCU Yard Talk (generated September 14, 2026, 9:50 AM EDT).

Platform	AEA	AIS	SPI	Tier
"X" (x.com)	16.5	67.9	75.8	FAIL
HBCU Yard Talk	78.5	17.3	12.4	Gold

“X” and HBCU Yard Talk represent opposite ends of the C.A.S. spectrum among platforms assessed to date. “X” fails to meet minimum performance thresholds across all three metrics, while HBCU Yard Talk achieves Gold Tier status, reflecting strong attention equity and low coercive pressure. HBCU Yard Talk's comparatively higher (though still passing) AIS and SPI scores suggest room for further optimization, but its overall profile should be treated as a positive reference model. “X,” by contrast, requires an immediate strategic overhaul given its failing tier status across every scored dimension. A full comparative report is available as a separate attachment to this assessment.

Validity

This assessment reflects X.com as observed during the documented September 14, 2026 research and scoring period. Because the assessment status is Pending Approval and the platform is not certified, no certification validity term applies at this time. The report should be treated as valid only for the assessed scope, evidence set, and platform conditions present during the assessment window, and should be re-evaluated after any material product, ranking, notification, monetization, or control-surface change. This report evaluates platform architecture and structural design conditions, not the truth or falsehood of individual posts, the legality of specific content or conduct, or the intent, motives, or character of individual users, employees, or stakeholders.

References

Clark, D. (2026). *Cognitive Attentional Standards (C.A.S.): An operational system for the evaluation of digital systems and measurement of attention harm*. Institute for Digital Equity and Attention Science (I.D.E.A.S.).