FINAL ASSESSMENT

# Inauthentic Activity in the Run-up to Election Day 2026

Five U.S. midterm races on X, Facebook and Instagram

Ohio and Georgia Governor · Michigan, Maine and Kansas Senate

**Released Monday, 9 November 2026**

Data period: 31 August – 3 November 2026 · Data pulled 6 November 2026

Cyabra Narrative Intelligence

# EXECUTIVE SUMMARY

**We witnessed a swarm of inauthentic accounts and fake results on Election Day, concentrated on X.**

Cyabra analyzed 379,025 public posts, replies and comments about five competitive 2026 midterm races on X, Facebook and Instagram, comparing the final stretch of the campaign (20 October – 3 November) with the period before it (31 August – 19 October). This report is Cyabra's final assessment of inauthentic activity in the run-up to Election Day. Comparisons cover the three platforms measured; TikTok could not be assessed for the final stretch.

**On Election Day, 18.7% of X posts about the five races came from inauthentic accounts,** up from 13.3% in August to October. Facebook stayed at 3.9%. Across all three platforms the share was 8.0% (8.8% before), because Facebook, by far the largest platform by volume, stayed low.

## Five key takeaways

| # | Takeaway | Key figure | Confidence |
|---|---|---|---|
| 1 | Both parties were targeted in candidates' replies and comments | Share rose for 5 of 8 candidates; 31.4% for Sen. Roger Marshall on Election Day | High (Moderate for Acton, Ramaswamy) |
| 2 | Fake results spread before polls closed | 3.4 million views of fake Michigan and Ohio "results," all on X | High |
| 3 | Sleeper inauthentic accounts woke up at the end | 84% of 8,344 inauthentic accounts in the final stretch were new; 76 created after 1 October | Moderate |
| 4 | On X, the accounts were overwhelmingly US-based, and they swarmed rather than organized | 84.9% of located accounts US-based; 299 reply swarms, one repeat team | Location: Moderate · Coordination: Low |
| 5 | Most of the conversation is real | 99.8% of views went to real accounts | High (Moderate for Acton) |

## What it means

The target was trust in the result, not any single race. Both parties were hit, the activity on X peaked on Election Day, and the most-viewed fake content was fabricated results posted before polls closed. In 2028, Cyabra expects the same playbook, faster and closer to the count: fresh sleeper accounts each cycle, AI chatbots invoked as a credibility stamp, and the same accounts working many races at once.

Most fake posts fizzle. The risk lies in the few that break through: those that land in the reply sections leaders read, and those real people pass on.

# SCOPE & DATA

| Item | Detail |
|---|---|
| Races | Ohio Governor (Amy Acton, D; Vivek Ramaswamy, R); Georgia Governor (Keisha Lance Bottoms, D; Rick Jackson, R); Michigan Senate (Abdul El-Sayed, D; Mike Rogers, R); Maine Senate (Troy Jackson, D; Sen. Susan Collins, R); Kansas Senate (Adam Hamilton, D; Sen. Roger Marshall, R) |
| Platforms | X, Facebook and Instagram |
| Before period | 31 August – 19 October 2026: 196,619 posts, replies and comments |
| Final stretch | 20 October – 3 November 2026 (to 6:59 pm ET): 182,406 posts, replies and comments |
| Election Day | 3 November 2026, 12:00 am – 6:59 pm ET: 51,520 posts, replies and comments |
| Data pulled | 6 November 2026 |
| Not included | TikTok (see Methodology); the Nevada races covered in Cyabra's pre-election survey |

## Election Day by platform

| Platform | Election Day posts | Inauthentic share, Election Day | Inauthentic share, 31 Aug – 19 Oct |
|---|---|---|---|
| X | 12,078 | 18.7% | 13.3% |
| Facebook | 36,315 | 3.9% | — |
| Instagram | 3,127 | 17.6%* | — |
| All three | 51,520 | 8.0% | 8.8% |

** Many Instagram accounts were not yet classified; treat the Instagram share as indicative only. Share = inauthentic ÷ (inauthentic + authentic). Platform-level baselines for Facebook and Instagram are not reported separately.*

# KEY TAKEAWAY 1: BOTH PARTIES WERE TARGETED IN THE REPLIES

**Confidence:** High; Moderate for Acton and Ramaswamy

> **The inauthentic share of replies and comments on candidates' own pages rose for 5 of the 8 candidates with enough data, in both parties.**

Reply sections are where candidates, reporters and campaigns read the public mood. Flooding them distorts what voters appear to think. The sharpest rise was on Sen. Roger Marshall's pages, from 13.1% to 23.3%; on Election Day, 31.4% of the classified replies and comments he received came from inauthentic accounts, nearly 1 in 3.

| Candidate | Race | Before | Final stretch | Change |
|---|---|---|---|---|
| Amy Acton | Ohio Governor · D | 11.1% | 14.1% | +3.0 |
| Vivek Ramaswamy | Ohio Governor · R | 10.8% | 12.5% | +1.7 |
| Susan Collins | Maine Senate · R | 9.4% | 9.8% | +0.4 |
| Troy Jackson | Maine Senate · D | 12.1% | 14.4% | +2.3 |
| Mike Rogers | Michigan Senate · R | 9.8% | 8.6% | −1.2 |
| Abdul El-Sayed | Michigan Senate · D | 8.2% | 5.2% | −3.0 |
| Roger Marshall | Kansas Senate · R | 13.1% | 23.3% | +10.2 |
| Rick Jackson | Georgia Governor · R | 13.1% | 8.9% | −4.2 |

*Replies and comments on each candidate's own X, Facebook and Instagram pages. Before = 31 Aug – 19 Oct 2026; final stretch = 20 Oct – 3 Nov 2026. Unclassified accounts excluded. Ten candidates in five races; Adam Hamilton (KS) and Keisha Lance Bottoms (GA) are not compared because their pages lacked comparable final-stretch data on X, the baseline platform.*

## Election Day replies and comments

| Candidate | Inauthentic share, 3 November |
|---|---|
| Roger Marshall | 31.4% |
| Troy Jackson | 17.7% |
| Amy Acton | 11.2% |
| Vivek Ramaswamy | 10.7% |
| Rick Jackson | 8.8% |
| Mike Rogers | 6.0% |

*Marshall: 76 of 242 classified replies and comments. 3 Nov 2026, 12:00 am – 6:59 pm ET.*

## Implications

*Cyabra assessment, based on the findings above.*

### WHY IT MATTERS

- **Reply sections shape perception.** Candidates, campaign staff, reporters and pollsters read the replies under a candidate's posts as a quick gauge of public mood. When up to 1 in 3 of those replies are inauthentic, as on Sen. Marshall's pages on Election Day, that gauge is distorted.
- **A small share can set the tone.** Replies appear directly under the original post and are often the first reactions a reader sees. A coordinated burst of hostile replies can make a fringe view look like a groundswell, even when most of the wider conversation is authentic.
- **It erodes trust in real voices.** Once audiences learn that reply sections are partly fake, they tend to discount genuine criticism and support as well. The damage extends beyond the inauthentic posts themselves.

### WHAT IT MEANT FOR THE 2026 ELECTION

- **No party was spared.** The share rose for Republicans (Marshall, Ramaswamy, Collins) and Democrats (Acton, Troy Jackson) alike. This was not a one-sided effort to help a single candidate.
- **Pressure peaked when it counted.** The highest Election Day shares landed on candidates in competitive races; at the moment campaigns were making final pushes and reporters were reading the room.
- **Candidates were exposed on their own pages.** Because the activity targeted candidates' own accounts, campaigns had the most direct view of it, and the most reason to misread it as genuine voter sentiment.

### WHAT IT MEANS FOR FUTURE ELECTIONS

- **Expect reply sections to be a primary battleground.** They are low-cost to target and high visibility. In 2028, with more candidates, more debates and national attention, the volume of targets will multiply.
- **Campaigns and newsrooms need a check before they react.** Treating reply sentiment as public opinion without verifying who is posting will become a routine vulnerability.
- **Bipartisan targeting points to a broader aim.** Activity that hits both sides suggests the goal is to inflame and distort, not to pick a winner. Defenses should be nonpartisan by design.

# KEY TAKEAWAY 2: FAKE RESULTS SPREAD BEFORE POLLS CLOSED

**Confidence:** High

> **Fake Michigan and Ohio "results" drew 3.4 million views on X before polls closed.**

At 1:35 pm ET on Election Day, an inauthentic account with 592 followers posted fabricated Michigan Senate "results" (Abdul El-Sayed 64%, Mike Rogers 36%) and asked X's AI chatbot, Grok, to "confirm this." The post drew 566,927 views. Real accounts carried the rest: about 2.8 million of the 3.4 million views came from authentic accounts repeating fake numbers.

| Time (ET), 3 Nov | Post | Account | Views |
|---|---|---|---|
| 1:35 pm | Michigan "results" 64% – 36%, asking Grok to confirm | Inauthentic | 566,927 |
| 3:16 pm | Ohio "results": Ramaswamy 53.67%, Acton 46.33% | Authentic | 1,258,296 |
| 6:25 pm | Same 64% – 36% Michigan figures | Authentic | 760,932 |
| Afternoon | Michigan "projection" with exact vote totals (4 posts) | Authentic | 772,225 |

*Public X posts, 3 Nov 2026; view counts as collected. Michigan polls closed at 8 pm ET. No fabricated results posts were found about Georgia, Kansas or Maine. Those found on Facebook drew under 40 interactions each; Facebook does not report views. Prediction-market odds excluded. The inauthentic account is not named.*

The pattern matters more than any single post: a small fake seed, passed on by real people, reached millions before a single official number existed, and an AI chatbot was invoked to lend it credibility.

## Implications

*Cyabra assessment, based on the findings above.*

### WHY IT MATTERS

- **It filled the information gap before official numbers existed.** In the hours before polls close there are no results, and that vacuum is easy to fill. Fake numbers posted in that window face no official figure to contradict them.
- **Real people did most of the spreading.** One inauthentic account seeded the Michigan numbers; authentic accounts carried about 2.8 million of the 3.4 million views. Once a fake claim crosses into real networks, removing the original source does little to stop it.
- **AI was invoked as a referee.** Asking a chatbot to "confirm this" borrows the perceived neutrality of AI to lend a fake claim credibility, whatever the chatbot answers.

### WHAT IT MEANT FOR THE 2026 ELECTION

- **Expectations were set before the count.** Voters who saw a confident "64–36" or "53.67% vs 46.33%" before polls closed had a reference point that real results could later be measured against and called suspect if they differed.
- **It targeted the closest-watched races.** The fake results concerned Michigan and Ohio, both competitive, high-profile contests. No comparable fake results were found for Georgia, Kansas or Maine.

- **The window of risk was hours, not days.** All the fake results appeared on Election Day afternoon and evening, leaving almost no time for correction.

## WHAT IT MEANS FOR FUTURE ELECTIONS

- **Fake results will be a standard Election Day tactic.** The technique costs almost nothing and reached millions. In a presidential year, with a single national result at stake and staggered poll closings across time zones, the window for fake "early results" is longer.
- **AI verification requests will grow.** As more voters turn to chatbots for answers, expect fake claims packaged with "AI, confirm this," and expect bad actors to quote chatbot replies as validation.
- **Speed of official communication becomes a defense.** Election officials who state clearly and repeatedly that no results exist before polls close, and who publish results through a single trusted channel, reduce the space fake numbers can occupy.

# KEY TAKEAWAY 3: SLEEPER INAUTHENTIC ACCOUNTS WOKE UP AT THE END

**Confidence:** Moderate

> **6,992 of the 8,344 inauthentic accounts active in the final stretch (84%) were sleepers: they had not appeared in the August–October period.**

Of the 4,402 sleeper accounts with a known creation date, only 76 were created after 1 October. These were existing accounts switched on for the finish, not freshly made ones. A list of known bad accounts built in September would have missed most of them; only behavior-based detection keeps up.

| Measure | Value |
|---|---|
| Inauthentic accounts active, final stretch | 8,344 |
| Sleeper accounts (not seen 31 Aug – 19 Oct) | 6,992 (83.8%) |
| Sleeper accounts with a known creation date | 4,402 |
| Of those, created after 1 October | 76 |
| Inauthentic accounts posting about 3+ races | 103 (64 before) |

*Accounts classified as inauthentic that posted about the five races on X, Facebook and Instagram. Creation dates as reported by the platform. Moderate confidence: collection caps limit how many accounts each scan captures.*

## Implications

*Cyabra assessment, based on the findings above.*

### WHY IT MATTERS

- **Static lists fail.** 84% of the inauthentic accounts in the final stretch were not active in August to October. Any defense built on lists of known bad accounts would have missed most of the activity that mattered.
- **Aged accounts look more credible.** Most sleeper accounts with a known creation date were created well before October. Older accounts with posting history are harder to spot and more convincing to real users than brand-new ones.
- **The effort was held back for the finish.** Activating accounts late concentrates impact in the final days, when there is least time to detect and respond.

### WHAT IT MEANT FOR THE 2026 ELECTION

- **The final stretch was a different landscape.** Monitoring that was calibrated in September would have underestimated the scale of inauthentic activity in the last two weeks.
- **Late activation coincided with the Election Day peak.** The influx of sleeper accounts lined up with the rise in inauthentic share on X, from 13.3% to 18.7% on Election Day.
- **Accounts worked across races.** 103 inauthentic accounts posted about three or more races in the final stretch, up from 64 before, showing that the same accounts were spread across contests.

## WHAT IT MEANS FOR FUTURE ELECTIONS

- **Detection must be behavioral and continuous.** Identifying accounts by how they behave, not by whether they were seen before, is the only approach that keeps pace with accounts that are rotated each cycle.
- **Expect a reserve of dormant accounts.** The 2026 pattern suggests accounts are built or acquired in advance and held until needed. Ahead of 2028, monitoring should look for dormant accounts that suddenly turn active.
- **Surge capacity matters.** Teams monitoring 2028 will need capacity to scale up in the final weeks, not just steady-state coverage.

# KEY TAKEAWAY 4: ON X, THEY WERE US-BASED, AND THEY SWARMED RATHER THAN ORGANIZED

**Confidence:** Location: Moderate · Coordination: Low

> **84.9% of inauthentic X accounts with a known location were US-based (3,922 of 4,620).**

Cyabra found 299 reply swarms on X: three or more inauthentic accounts replying to the same post within 10 minutes. The swarms drew in 969 accounts. Only one team of three accounts, swarmed together more than once; the rest never repeated with the same partners. The activity looks opportunistic, not orchestrated.

Foreign-based accounts were a small minority: 96 of the 969 swarming accounts, never more than four in any swarm. This does not rule out foreign involvement, including on other platforms. Cyabra is not attributing the activity to any country or actor.

| Race | Swarming accounts | US-based | Foreign-based | Unknown |
|---|---|---|---|---|
| Ohio | 442 | 388 | 41 | 13 |
| Maine | 271 | 244 | 23 | 4 |
| Kansas | 162 | 138 | 20 | 4 |
| Michigan | 70 | 53 | 12 | 5 |
| Georgia | 24 | 24 | 0 | 0 |
| Total | 969 | 847 | 96 | 26 |

| Region of foreign-based swarming accounts | Accounts |
|---|---|
| Canada | 33 |
| Europe | 26 |
| Asia | 18 |
| Australia / New Zealand | 9 |
| Middle East / Africa | 6 |
| Latin America | 4 |

*Inauthentic X accounts, 20 Oct – 3 Nov 2026. Accounts are assigned to the race they posted about most. Location = where an account appears to be based (Cyabra assessment), not who operates it. Location is X-only: Facebook and Instagram accounts carry no location data.*

## Implications

*Cyabra assessment, based on the findings above.*

### WHY IT MATTERS

- **The problem is not only foreign.** 84.9% of inauthentic X accounts with a known location were US-based. Framing election manipulation solely as a foreign threat misses most of what Cyabra observed in these races.

- **Swarming is cheap and effective.** 299 reply swarms required no ongoing organization: only one team of three accounts swarmed together more than once. Loose, opportunistic pile-ons can still create the appearance of a crowd.
- **Location is not attribution.** Where an account appears to be based does not reveal who operates it. These findings describe behavior, not responsibility.

## WHAT IT MEANT FOR THE 2026 ELECTION

- **The activity looked homegrown and reactive.** Most swarms came from US-based accounts piling onto whatever a candidate had just posted, especially in Ohio, Maine and Kansas.
- **Foreign-based accounts were on the margins.** 96 of 969 swarming accounts were foreign-based, never more than four in a swarm. The swarm data showed no sign of a joint domestic–foreign operation in these races on X.
- **It did not require a central plan to have an effect.** Even without organized teams, the cumulative result was the rising inauthentic share in candidates' replies described in Key takeaway 1.

## WHAT IT MEANS FOR FUTURE ELECTIONS

- **Defenses must cover domestic activity.** Programs that focus only on foreign influence will miss much of the activity seen in 2026. Monitoring should be source-agnostic.
- **Watch for organization emerging.** 2026 showed swarming without stable teams. If repeat teams and scheduled coordination appear in 2028, that would mark an escalation worth flagging early.
- **Foreign involvement cannot be ruled out.** This analysis covers X only for location and does not cover TikTok. Foreign activity may be present on other platforms or behind US-based accounts.

# KEY TAKEAWAY 5: MOST OF THE CONVERSATION IS REAL

**Confidence:** High; Moderate for the Acton finding

> **Real accounts drew 99.8% of views in the final stretch (99.9% before).**

Most inauthentic posts went unseen. Even the biggest attack of the final stretch, on Amy Acton, was carried mostly by real people.

| Personal attack on Acton | Before (7 weeks) | Final stretch (15 days) |
|---|---|---|
| Posts | 488 | 1,198 |
| Inauthentic share | 8.3% | 17.1% |
| Posts from authentic accounts | — | 948 (82.9% of classified) |

*Posts about the Ohio race carrying the personal attack on Acton (keyword match); Cyabra does not repeat the claim. View share = views on authentic accounts' posts ÷ views on all classified posts, where reported.*

The lesson is proportion: inauthentic accounts are a minority whose impact depends on where they concentrate.

## Implications

*Cyabra assessment, based on the findings above.*

### WHY IT MATTERS

- **Proportion prevents overreaction.** Overstating the reach of inauthentic accounts can itself undermine trust in elections and in genuine debate.
- **Impact comes from where fakes land, not how many there are.** The posts that mattered were in reply sections leaders read, and seeds real people spread.
- **Real people carry most narratives.** Even in the Acton attack, where the inauthentic share doubled, about 4 in 5 posts were authentic. Inauthentic accounts amplify; they rarely dominate.

### WHAT IT MEANT FOR THE 2026 ELECTION

- **The 2026 conversation was mostly authentic.** Voters online in these races were overwhelmingly engaging with real people.
- **Inauthentic accounts acted as accelerants.** They seeded, inflamed and amplified, most visibly in the fake results and the Acton attack.
- **Cyabra makes no claim about outcomes.** This analysis measures online activity, not votes.

### WHAT IT MEANS FOR FUTURE ELECTIONS

- **Focus resources on breakout points.** Watch for the moments when inauthentic content crosses into authentic networks, rather than chasing every fake account.
- **Communicate with precision.** Describe inauthentic activity accurately, neither dismissing nor exaggerating it, so warnings stay credible in 2028.
- **Support the public as the last line of defense.** Because real people do most of the spreading, habits such as checking a profile before sharing are among the most effective protections.

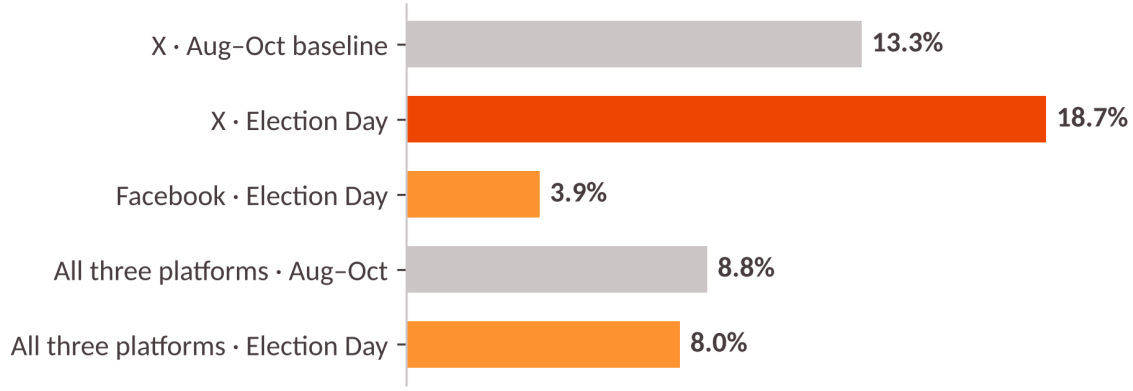
# WHAT IT MEANT FOR THE 2026 MIDTERMS

**Confidence:** Cyabra assessment

**The target was trust in the result, not one race.**

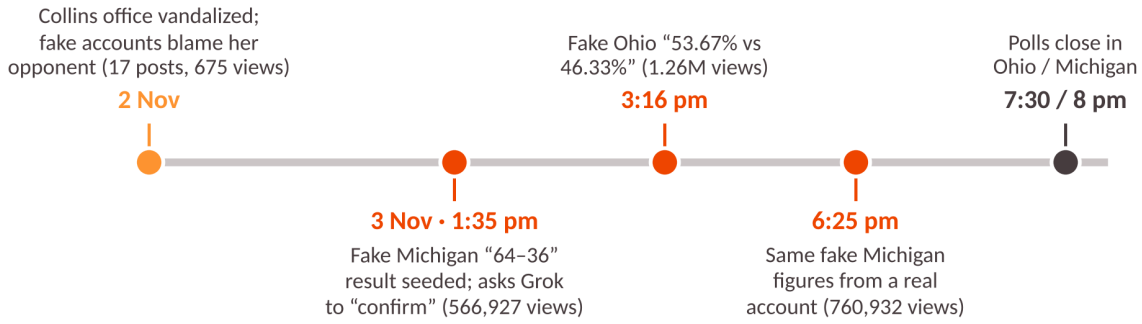
- **Both sides were hit.** Republican and Democratic candidates alike saw rising inauthentic shares in their replies and comments. The sharpest Election Day rise was on X.
- **Results were faked before they existed.** Fabricated Michigan and Ohio results drew 3.4 million views before polls closed.
- **Fake accounts react within hours.** After Sen. Collins's Biddeford office was vandalized on 2 November, inauthentic accounts blamed her opponent's side. Their 17 posts, out of 204 about the incident, drew 675 views in total.
- **The count is the next target.** Expect claims about "dumps," "flips" and recounts in close races, aimed at election officials.

## Inauthentic Share of Posts: The Election Day Rise Was on X




Source: Cyabra. Share = inauthentic ÷ (inauthentic + authentic). Election Day = 3 Nov 2026, to 6:59 pm ET.
Instagram (17.6%) not shown: many Instagram accounts were not yet classified.

## Election Day on X: Fake Results Filled the Hours Before Any Count



## Cyabra's assessment of intent

**Confidence:** Low–Moderate. Intent is inferred from observed behavior; it is a judgment, not a finding of fact.

| Apparent aim | How it worked | Evidence in 2026 |
|---|---|---|
| Fake a crowd | Flood candidates' reply sections so a fringe view looks like public mood | Inauthentic share rose for 5 of 8 candidates in both parties; 31.4% of Sen. Marshall's Election Day replies |
| Discredit the result before it exists | Seed fake tallies before polls close so any outcome can be called suspect | 3.4 million views of fake Michigan and Ohio results; an AI chatbot asked to "confirm" |
| Inflame and assign blame | Turn real events and personal attacks into anger at the other side | Blame posted within hours of the Collins office vandalism; inauthentic share of the Acton attack rose from 8.3% to 17.1% |
| Stay hard to catch | Rotate accounts, spread across races, avoid stable teams | 84% sleeper accounts; 103 accounts in 3+ races; 299 swarms with one repeat team |

Both sides were targeted. Cyabra judges the aim to be eroding trust in the electoral process, not tipping a particular race.

## What did not happen

An accurate assessment also records what Cyabra did not find. These limits matter as much as the headline findings, because overstating inauthentic activity can itself undermine trust.

- **No widespread cross-platform surge.** Facebook, the largest platform by volume, stayed at 3.9% on Election Day, and the combined share across X, Facebook and Instagram fell slightly, from 8.8% to 8.0%.
- **No fake results with real reach outside Michigan and Ohio.** None were found for Georgia, Kansas or Maine, and those found on Facebook drew under 40 interactions each.
- **No sign of organized teams.** Only one team of three accounts swarmed together more than once across 299 swarms.
- **No foreign-led swarms.** Foreign-based accounts were never more than four in a swarm, and 84.9% of located inauthentic X accounts were US-based.
- **No takeover of the conversation.** Real accounts drew 99.8% of views.

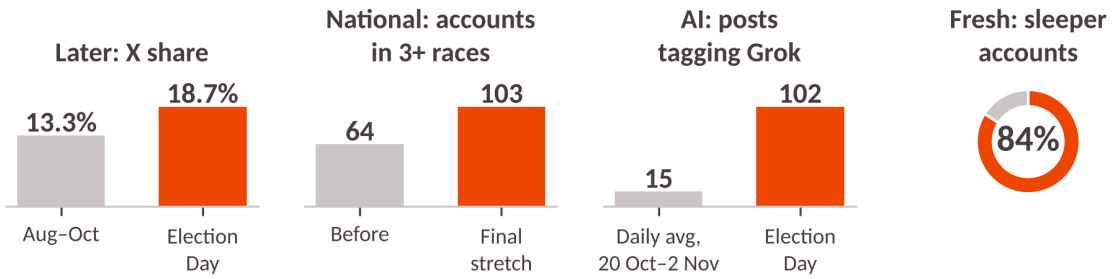
# INSIGHTS FOR THE 2028 PRESIDENTIAL CYCLE

**Confidence:** Forward-looking Cyabra assessment

**2026 was a preview. Expect the same playbook, faster.**

| Signal in 2026 | Evidence | What to expect in 2028 |
|---|---|---|
| Later, closer to the count | On X, 18.7% on Election Day vs 13.3% Aug–Oct | The surge will land when there is least time to respond |
| Fresh accounts every cycle | 84% of final-stretch inauthentic accounts were new | Blocklists won't keep up; behavior-based detection will |
| AI as a credibility stamp | 310 X posts tagged an AI chatbot in the final stretch, 102 on Election Day | Fake claims will arrive with "AI, confirm this" |
| National, not local | 103 accounts worked three or more races, up from 64 | The same accounts across many swing states in a presidential year |

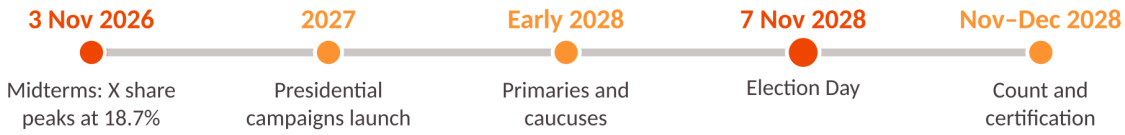
### Four 2026 Signals That Point to 2028




Source: Cyabra. X share on Election Day vs 31 Aug–19 Oct. Grok: 310 X posts in the final stretch, 102 on Election Day (other days averaged ~15). Sleeper accounts = share of 8,344 final-stretch inauthentic accounts not seen before.

The road to 2028 has several targets: presidential campaign launches in 2027, primaries and caucuses in early 2028, Election Day on 7 November 2028, and the count and certification that follow.

### The Road to 2028: Every Stage Is a Target



## Early-warning indicators for 2028

Cyabra recommends tracking these indicators from the start of the 2028 cycle. The 2026 baseline gives a reference point; the escalation signal is what would show the playbook growing.

| Indicator | 2026 baseline | Escalation signal to watch |
|---|---|---|
| Inauthentic share near voting | 18.7% on X on Election Day, up from 13.3% | A rise that starts earlier, during early voting, or spreads to more platforms |
| Sleeper activation | 84% of final-stretch inauthentic accounts new; most with known dates created before October | Waves of dormant accounts switching on before debates, primaries or early voting |
| Cross-race accounts | 103 accounts in three or more races, up from 64 | The same accounts spanning many swing states and presidential and down-ballot races |
| Fake results | 3.4 million views, Michigan and Ohio only, all on X | Fake results across several states and time zones, or presented as screenshots of official sites |
| AI as validator | 310 X posts tagging an AI chatbot; 102 on Election Day | Fake claims routinely paired with chatbot “confirmations” or AI-generated evidence |
| Organization | 299 swarms; one repeat team of three | Repeat teams, scheduled posting or coordination across platforms |
| Foreign-based participation | 96 of 969 swarming accounts; never more than four in a swarm | A larger foreign share, or swarms led by foreign-based accounts |

# RECOMMENDATIONS: PREPARING FOR 2028

> **See coordinated activity before it goes viral.**

## Federal: see it all

- Cross-state, cross-platform monitoring: the same accounts worked several races at once.
- Shared standards for calling out coordinated inauthentic activity.
- Fast channels with platforms and AI providers to flag fabricated results.

## State: own the count

- Live monitoring through Election Night, the count and certification.
- Say early and often: no results exist until polls close.
- Prebunk the known claims: ballot rules, machines, late counts.

## Local: respond first

- Early alerts on false claims about polling places and incidents.
- Verify a pile-on before reacting: it may be fake accounts, not voters.
- One trusted source for local updates on election night.

## Voters and the public

- Check the profile before you share.
- A reply section isn’t a poll.
- Results come from election officials, not from posts or chatbots.

# METHODOLOGY

## Data collection

Cyabra collected public posts, replies and comments on X, Facebook and Instagram about the ten candidates in five races, using searches for candidate names and campaign hashtags and scans of each candidate's own accounts. The analysis compares two periods: 31 August – 19 October 2026 (the window of Cyabra's pre-election survey) and 20 October – 3 November 2026, through 6:59 pm ET on Election Day. Data was pulled on 6 November 2026.

## Account classification

Cyabra classifies each account as authentic or inauthentic using behavioral signals, not the views it expresses. "Inauthentic" means an account is not what it presents itself to be; it may be automated or operated by a person, so it is broader than "bot." No single trait determines the label. Accounts Cyabra could not classify are excluded from every share and reported separately.

## Key definitions

| Term | Definition |
|---|---|
| Inauthentic share | Inauthentic ÷ (inauthentic + authentic). Unclassified accounts excluded. |
| Replies and comments | Responses posted on a candidate's own X, Facebook or Instagram pages. |
| Fake results | Posts presenting vote shares or outcomes as results before polls closed. Prediction-market odds excluded. |
| Sleeper account | An inauthentic account active in the final stretch that did not appear in the before period. |
| Reply swarm | Three or more inauthentic accounts replying to the same X post within 10 minutes. |
| Location | Where an account appears to be based (Cyabra assessment). It does not identify who operates the account. |
| View share | Views on authentic accounts' posts ÷ views on all classified posts, where the platform reports views. |

## Confidence levels

Each finding carries a Cyabra analyst confidence level, based on sample size, the share of unclassified accounts, collection limits and method.

| Level | Meaning | Applies to |
|---|---|---|
| High | Direct counts from large samples with little unclassified data | Fake results; most reply shares; Election Day share on X; view share |
| Moderate | A solid pattern with known limits | Acton and Ramaswamy reply shares; the Acton attack (keyword match); sleeper accounts (collection caps); locations (may be masked) |
| Low | Limited evidence or a judgment | Coordination beyond one-off swarms; who operated accounts; intent |

## Limitations

- **TikTok.** Cyabra's scans hold TikTok comments only up to 19 October, so the final stretch could not be compared. TikTok is excluded from this report.
- **Instagram.** Many Instagram accounts were not yet classified at the time of the pull; Instagram shares are indicative only.
- **Location.** Available for X accounts only and may be masked by VPNs or false profile information.
- **Collection limits.** Scans capture a capped number of items per search, which can understate account counts.
- **Candidate comparisons.** Eight of ten candidates are compared; Adam Hamilton and Keisha Lance Bottoms lacked comparable final-stretch data on X. The Nevada races in Cyabra's pre-election survey are not covered.
- **Coordination.** Swarm analysis identifies accounts acting at the same moment on the same post. It is not Cyabra's scored coordination model, and it does not establish common control.
- **Outcomes.** Cyabra measures online activity, not votes. This report makes no claim about whether inauthentic activity changed any election result.
- **Privacy.** Inauthentic accounts are not named. Example posts are anonymized.