SECUREFI INTELLIGENCE BRIEF

When Trusted Systems Go Offline

Three major disruptions in five days. None has been attributed to a cyberattack. Together, they show how hidden defects and interconnected systems can turn a technical problem into widespread consequences.

Edited by David Haberland
SecureFi, Inc. | November 17, 2026 | Research current to November 17, 2026

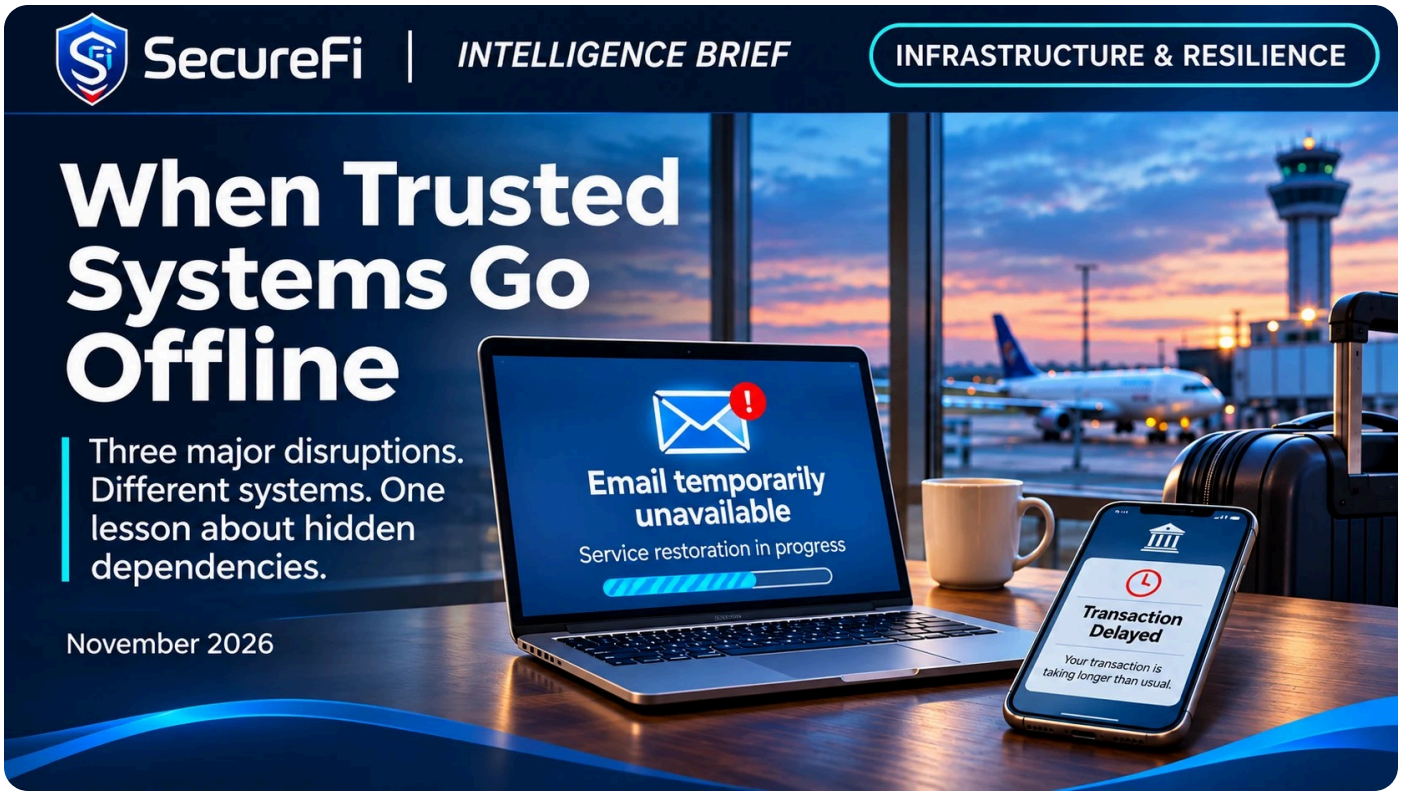

Executive Summary

In five days in November, three systems that millions of people depend on experienced significant disruptions.

On November 12, a Microsoft 365 failure left many customers worldwide unable to reach their email for most of a day. On November 13, the Federal Aviation Administration stopped departures nationwide after network trouble destabilized communications and radar data at 17 of its 22 regional control centers, the FAA says. On November 16, a software problem at JPMorgan Chase disrupted banking, with delays reaching customers of other banks.

None of the three has been attributed to a cyberattack. No evidence connects them.

What they share is a pattern. In each case, a problem that began in one place spread much further than it should have. Microsoft has begun to explain why its own safeguards did not stop that spread. The FAA is still investigating. Chase has said little beyond its first statement.

The question this brief asks is not whether these systems can fail. Every complex system can. It is **how well each was built to keep a failure from spreading, and how well people and operations held up when it did.**

> **Trusting a system means understanding not only what it can do when everything works, but what happens when it doesn't.**

At a glance

	Microsoft 365	FAA air traffic control	JPMorgan Chase
Date	Nov 12 to 13 (monitoring into Nov 14)	Nov 13	Nov 16

	Microsoft 365	FAA air traffic control	JPMorgan Chase
What failed	**Network gateway servers carrying Exchange Online connections**	**FAA telecommunications network feeding communications and radar data**	**Chase banking transactions**
Stated cause	**A latent software defect from a July update, triggered by higher load and reduced capacity (Microsoft, preliminary)**	**Abnormal traffic from routers at two sites during contractor work; root cause unknown (FAA)**	**"A software bug in an on-premise hardware system" (Chase)**
How far it spread	**More than 80% of gateway servers unresponsive at the low point; global (Microsoft)**	**17 of 22 regional centers; nationwide ground stop (FAA)**	**Delays for other banks' customers transacting with Chase (Chase)**
Cyberattack?	**Microsoft: not aware of malicious activity**	**Officials: no evidence of anything nefarious**	**Chase: "not cyber-related"**
Status	**Preliminary report Nov 15; final report due Nov 20**	**Routers under lab testing; investigation ongoing**	**Restored Nov 16; no detailed report**

1. Microsoft 365: When a Safeguard Spreads the Failure

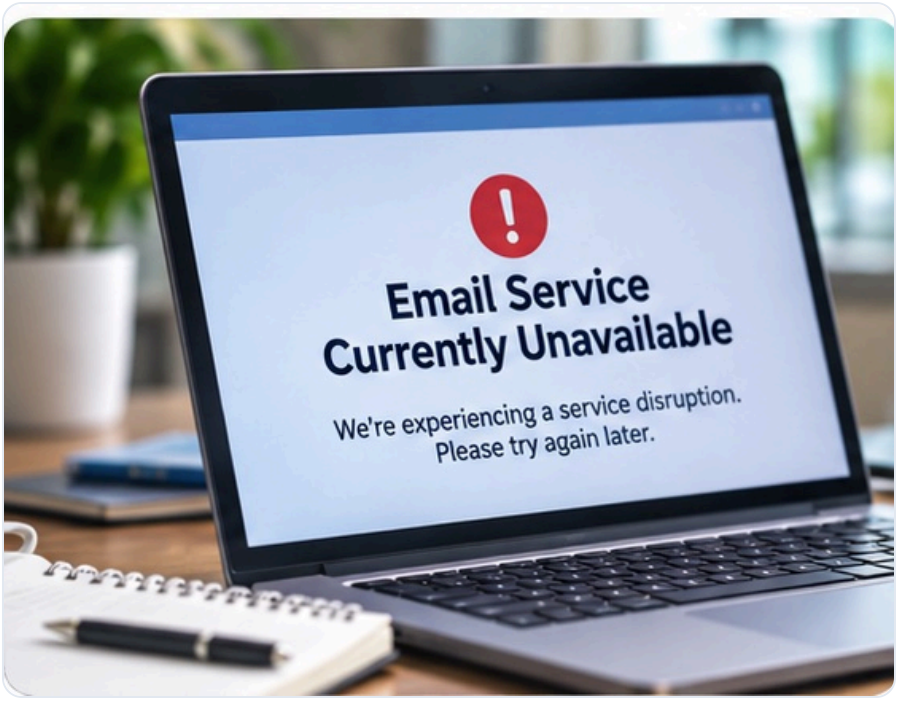

WHAT HAPPENED

According to Microsoft's preliminary report, from about 11:20 a.m. Eastern on November 12 until about 11:10 a.m. on November 13, most affected customers could not reach their mailboxes through Outlook or other email apps. Some features in Copilot, Defender, Purview, Intune, Planner and Power BI that depend on Exchange Online were affected; Microsoft says most of those services otherwise stayed online.

Not everything went down. Teams, OneDrive, SharePoint and the Office apps were able to route around the broken component and recovered sooner, Microsoft says. Teams chats and meetings largely kept working.

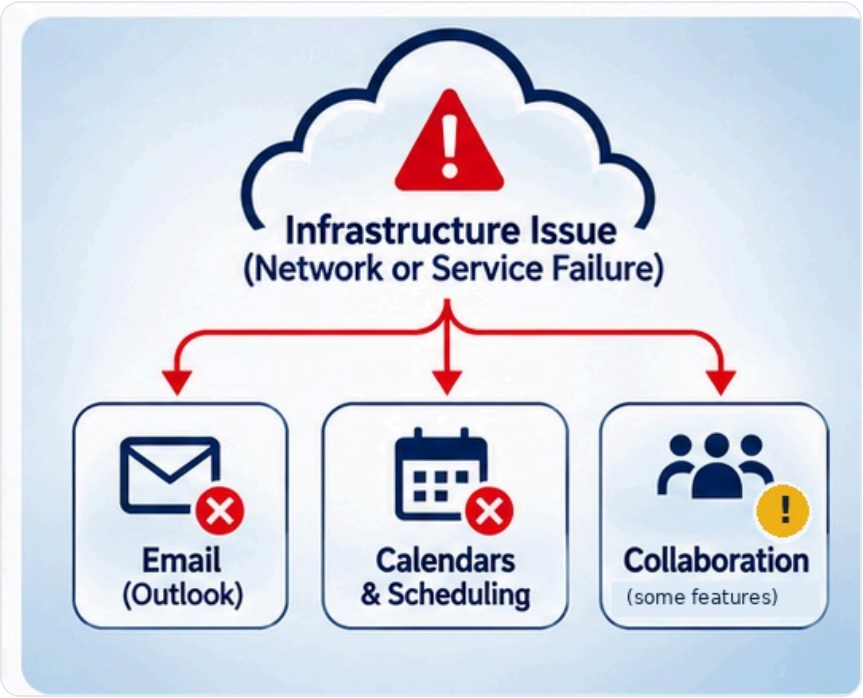

One underlying failure can affect several services at once.

WHAT MICROSOFT'S PRELIMINARY REPORT FOUND

Microsoft's first explanation pointed to a network capacity upgrade. Its November 15 preliminary report gives a more detailed account, built on three conditions that came together:

A hidden defect. A July 2026 update to Microsoft's network gateway servers contained a bug, the report says. Under heavy load, it caused servers to run out of processing power and memory and stop responding. Microsoft's load testing had used machines with fewer processor cores than the newer production servers, where the bug was worse. The change ran for more than three months without visible problems.

A normal failover. About 15 hours before the outage, Teams automatically shifted its traffic onto the gateway that also serves Exchange Online. Microsoft calls this kind of failover "part of normal operations." It increased the load.

Less spare capacity. Nearly 14% of server capacity was already out of service because of an unrelated issue from the previous day, according to the report.

Together, these triggered the defect.

WHY IT SPREAD

This is the most important finding for this brief.

As servers stopped responding, Microsoft's traffic management did what it was designed to do: it sent users to the remaining healthy servers. Those servers then hit the same defect and failed too. In less than two hours, the failure had spread across the fleet. At the lowest point, more than 80% of the servers were unresponsive, Microsoft reports.

Microsoft's systems are divided into separate groups and regions so that a problem in one stays there. According to the report, that separation did not contain this failure because of a separate issue.

> **The mechanism meant to protect availability, redirecting traffic away from failing servers, is what carried the failure across the system.**

WHY RECOVERY TOOK SO LONG

- Servers brought back online were immediately overwhelmed by the backlog of waiting traffic and failed again.
- Microsoft's automated recovery was not built for so many servers failing at once.
- Its usual rollback tools could not quickly undo a change made months earlier.

Microsoft paused all deployments and configuration changes on the afternoon of November 12, built a fix that returned the code to a previously working state, and rolled it out region by region. In its service health updates, Microsoft said it was extending the email delivery window to 48 hours so delayed messages would not be returned as undeliverable.

In those updates Microsoft also said it was "not aware of any malicious activity." Its final report, due November 20, is expected to list further repairs, according to the preliminary report.

SECUREFI TAKEAWAY

A safeguard designed for small failures can make a large one worse.

The question for any critical system is not only whether it fails over, but what happens when everything it fails over to has the same weakness.

2. FAA: When a Problem at Two Sites Reaches 17 Centers

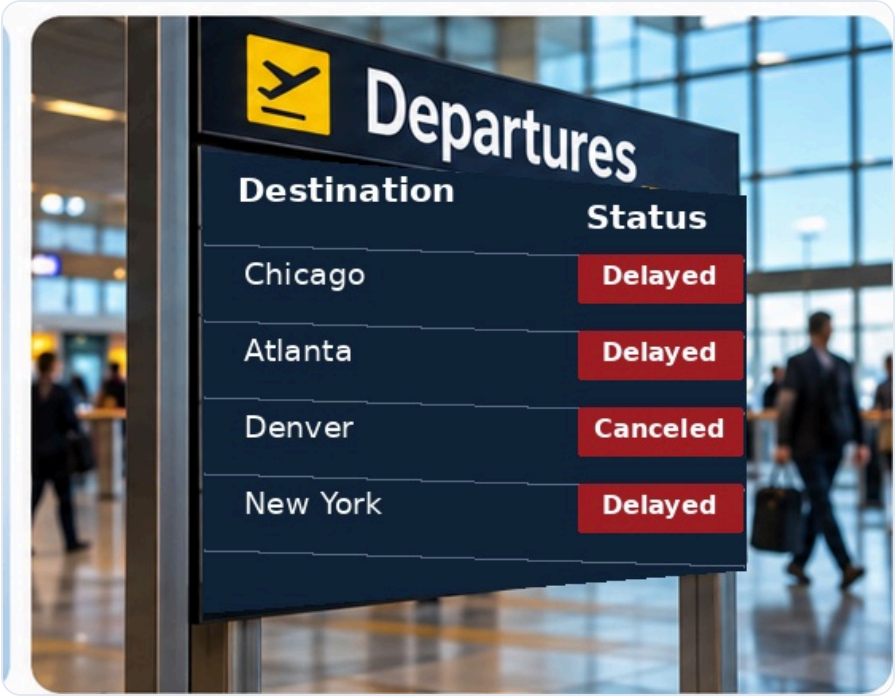

WHAT HAPPENED

At about 3:36 p.m. Eastern on Friday, November 13, air traffic control communications and radar information began to become unstable across much of the FAA's network, according to the Department of Transportation. They did not fail completely, but the FAA stopped departures nationwide. More than 3,500

aircraft were in the air when the ground stop was issued, and controllers brought them in safely using established procedures, DOT says. The FAA reopened the airspace at 8:15 p.m. Eastern, according to the Department of Transportation.

Reopening the airspace did not end the disruption. Aircraft and crews were out of position across the country. The FAA counted about 12,462 cancellations on Friday, November 13, and cautioned that not all of them were directly caused by the ground stop. Reuters reported more than 18,000 cancellations in total over the following days. Officials said airline schedules were expected to fully recover by the end of Sunday, November 15, Scripps News reported.

An airline schedule returning to normal is not the same as every traveler getting home. American Airlines warned that some individual flights might still be canceled as operations were restored, and some airlines extended rebooking waivers to Sunday or Monday, Scripps News reported. Travelers with fewer options, such as those flying from airports with only a few daily departures, can face longer waits than the national numbers suggest.

WHAT INVESTIGATORS HAVE FOUND

The FAA says abnormal traffic, about 1,000 connection attempts every four to five seconds, flooded its telecommunications network. That network is operated under contract by L3Harris. Communications and radar feeds became unstable at 17 of the FAA's 22 regional centers.

The FAA identified two potential originating sites:

El Dorado, Arkansas, a radio communications site where L3Harris was installing a new router.

Amarillo, Texas, a radar site where, according to Secretary of Transportation Sean Duffy, L3Harris had replaced a router about three weeks earlier.

The new routers, made by Juniper Networks, were replacing Cisco equipment dating to 2001, Reuters reported. Blocking the two sites stabilized the network. As a precaution, the FAA says it disabled the same function on more than 60 similar routers.

The FAA says the Arkansas work was routine maintenance of its existing network, not part of its larger modernization program.

WHO IS RESPONSIBLE: NOT YET DECIDED

The root cause is still unknown. FAA Administrator Bryan Bedford has said investigators are examining whether the problem came from the installation, the configuration or the router technology itself, according to news reports.

- Secretary Duffy has said it was either "contractor error or a faulty piece of equipment" and that L3Harris "didn't meet the moment," Al Jazeera reported.
- L3Harris says it is cooperating with the investigation and has completed more than 70% of the planned airspace communications conversion, Al Jazeera reported.
- HPE, which owns Juniper, says its equipment worked as configured, according to Reuters.

Both routers are being tested at an L3Harris laboratory in Florida under FAA oversight, Reuters reported. Duffy has invited the FBI to observe. Officials say they have found no evidence of anything malicious, and the FAA administrator said there is "zero evidence of anything AI-related," addressing questions about a new AI-based FAA planning tool. The FAA and Secretary Duffy say that tool was not connected to the affected systems.

An official's criticism is not a finding. SecureFi will report the investigation's conclusions when they are published.

WHAT WORKED

The technology became unreliable, and the safety system held. Controllers had procedures for exactly this situation, the authority to stop departures, and the training to land thousands of aircraft safely while their tools were degraded.

> **SECUREFI TAKEAWAY**
>
> **The ability to stop safely can be as important as the ability to operate efficiently.**
>
> The unanswered question is why changes at two remote sites could destabilize most of a national network.

3. JPMorgan Chase: A Bank's Problem Reaches Other Banks



Illustration. A problem at one bank can delay payments to and from customers at other banks.

On Monday, November 16, thousands of customers reported trouble reaching online accounts and completing transactions at major U.S. banks.

The problem began at JPMorgan Chase. Spokesman Joe Evangelisti told the Associated Press the issue was related to a software bug "in an on-premise hardware system" and that "it is not cyber-related." "This is not a systemic issue," he said, adding that "other banks may have been impacted by slowness in transactions" only "to the extent that they were a party to a transaction with us."

That is how customers of other banks were affected. U.S. Bank said its systems were not directly impacted, but its customers using Zelle could see delays. Zelle said its network was "operating as it should" and that an issue at the financial institutions might be delaying payments. Chase said transactions were "posting normally" by about 6:30 p.m. Monday, NBC reported.

Chase has not published a detailed explanation, and banks rarely do. What is clear is the pattern: a defect at one bank slowed transactions at other banks that had no reported problem of their own.

> SECUREFI TAKEAWAY
>
> **A problem at a bank you don't use can still delay your money when your transaction involves that bank.**

4. What the Three Incidents Have in Common

These were separate events with different technical causes. No evidence connects them, and treating them as connected would repeat the speculation this brief is meant to answer.

Seen together, they show three things.

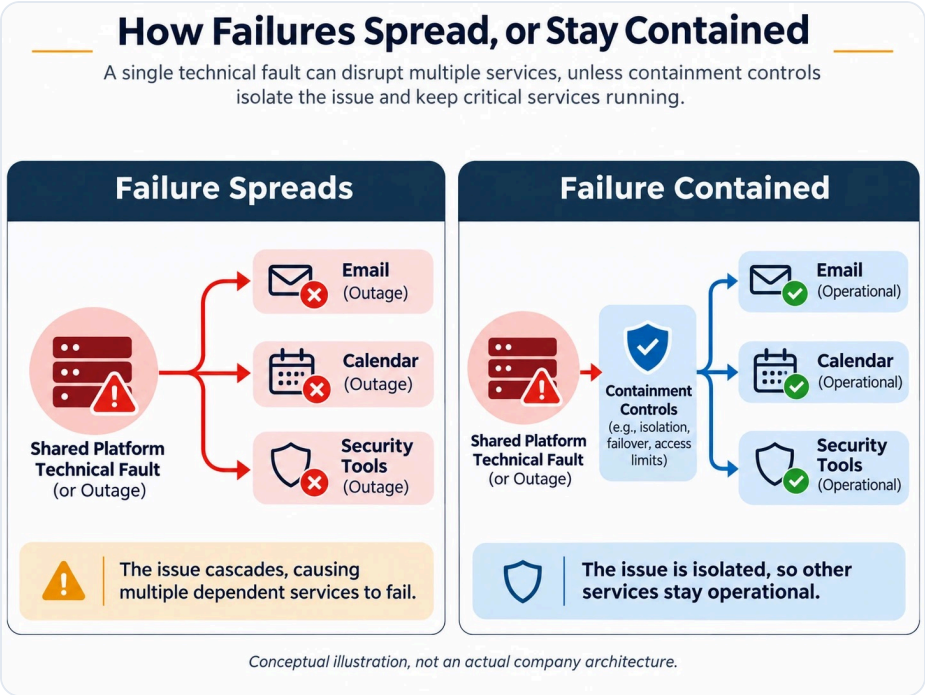

Failures stacked. Microsoft's outage needed an old defect, a normal failover and capacity already offline at the same time, according to its report. The FAA's investigation involves two router sites, one changed weeks earlier. Whether both were needed to produce the disruption is one of the open questions.

Spread mattered more than the trigger. In each case, the starting problem was local. The damage came from how far it traveled: across Microsoft's server fleet, across 17 FAA centers, and across banks through shared payment systems.

People and procedures limited the harm. Air traffic controllers landed every aircraft safely. Microsoft engineers froze changes and restored service region by region. Recovery depended on people as much as on technology.

5. Security, Reliability and Resilience: Three Different Questions

When a major system fails, the first question people ask is whether it was hacked. That is the security question. Two others matter just as much.

Security: can the system resist attack? It covers unauthorized access, fraud, data theft and deliberate disruption. None of these three incidents has been attributed to an attack.

Reliability: does the system work when it is needed? Hidden defects, configuration problems, overloaded equipment and unexpected interactions can stop a system with no attacker involved. All three incidents were reliability failures.

Resilience: can people and organizations keep going, and recover safely, when something fails? It covers containment, fallback procedures and recovery. The FAA's safe landings were resilience working. Microsoft's failed containment was resilience falling short.

> **Security reduces the risk of hostile action. Reliability reduces the chance of failure. Resilience reduces the harm when either happens.**

6. Why It Feels Like Everything Is Breaking

Three major failures in five days left many people asking whether something larger was going on. "It makes you think, like, is this a cyberattack?" one traveler told the Associated Press. Another said, "I don't know what's true because there are so many people saying different things."

That reaction is understandable. When unrelated events cluster, people look for a pattern, and speculation fills the gap faster than official findings arrive. But the evidence so far points to three separate technical failures, each explained, at least in part, by the organization responsible.

The risk in moments like these is acting on what you heard rather than what was confirmed. The useful habits are simple: wait for the responsible organization's

own explanation, separate what is known from what is assumed, and get updates and help directly from the company involved rather than from messages or posts that reach you first.

7. What It Means for You

For individuals and families

Keep a second way to reach people when email or messaging is down.

Save travel and reservation details somewhere you can open without email.

Know your refund rights. Under U.S. Department of Transportation rules, if an airline cancels your flight and you choose not to travel or accept alternative compensation, you are entitled to a refund regardless of the reason. Eligible refunds must generally be issued automatically to the original payment method.

Don't assume a failed transfer failed. During a bank disruption, check your account before resending a payment, to avoid paying twice.

Get status and help from the source. Use the company's own app, website or status page, not messages or posts that reach you first.

For businesses and organizations

Map your hidden dependencies. Which of your vendors rely on the same email platform, network provider or payment system?

Ask vendors how they contain failures, not only how they prevent them.

Test recovery at full scale. Microsoft's recovery tools were built for small failures, not fleet-wide ones.

Decide in advance who can stop operations, and on what signal.

Plan reconciliation. After service returns, how will you find missed messages, duplicate payments and incomplete orders?

8. What We Still Need to Learn

Microsoft: the final report's full list of repairs; what the "separate issue" was that defeated isolation between server groups; and how load testing will change.

FAA: whether the cause was installation, configuration or equipment; how the El Dorado and Amarillo routers interacted; why the traffic reached 17 centers; and what changes the FAA will make to maintenance approval and timing.

Chase: what failed, how many transactions were delayed, and whether any customers were charged fees or paid twice.

These are the questions any responsible review should answer. They are not findings of fault.

The SecureFi Assessment

None of these three failures has been attributed to an attack. All three were reminders that the systems we trust are only as strong as their ability to contain a problem when one appears.

Preventing every failure is impossible. Limiting how far a failure spreads, keeping people safe while it is fixed, and recovering without causing a second failure are what separate a resilient system from a fragile one.

For consumers, that means keeping alternatives and verifying before acting. For organizations, it means understanding dependencies and testing recovery. For providers of critical services, it means designing safeguards that cannot become the path a failure travels.

> **We don't control the systems we depend on. We can understand those dependencies, prepare alternatives, and make better decisions when something fails. That is what risk intelligence means.**

SecureFi will update this brief as Microsoft, the FAA and Chase publish their findings.

Sources and References

1. **Microsoft 365**
   Preliminary Post-Incident Report, EX1550543 / MO1550821 (Microsoft 365 admin center; third-party copy at 365fanatic.com). November 15, 2026.

2. **Microsoft 365 Status on X**
   Announcement of the preliminary post-incident report. November 2026.

3. **Microsoft 365 service health, MO1550821 (mirrored by the University of Pennsylvania)**
   Exchange Online and Microsoft 365 service disruption updates. November 12 to 14, 2026.

4. **The Register**
   Microsoft services restored. November 13, 2026.

5. **Federal Aviation Administration**
   Nationwide Ground Stop Questions and Answers. November 15, 2026.

6. **Federal Aviation Administration**
   FAA Statement on Nationwide Ground Stop. November 13, 2026.

7. **U.S. Department of Transportation**
   Secretary Duffy Announces Lifting of Nationwide Ground Stop. November 13, 2026.

8. **Reuters (via The Business Standard)**
   US probes why a router installation sparked mass flight grounding. November 17, 2026.

9. **Al Jazeera**
   What forced the US to ground all planes? November 15, 2026.

10. **U.S. Department of Transportation**
    Refunds (aviation consumer protection).

11. **Associated Press (via WMUR)**
Online banking delays for major banks triggered by technical issue at JPMorgan Chase. November 16, 2026.

12. **NBC**
Chase says transactions back to normal after outage online and at some branches. November 16, 2026.

13. **Scripps News**
FAA says air travel is safe as airlines clear backlog from nationwide ground stop. November 15, 2026.

14. **Associated Press (via KPTV)**
Grounded jets, Microsoft outage, AI chaos: Could everything be broken? November 15, 2026.

About This Intelligence Brief

SecureFi Intelligence Briefs examine significant technology, cybersecurity and personal-risk developments using publicly available information from authoritative sources.

This brief distinguishes official findings, preliminary explanations, attributed statements and SecureFi analysis. The three incidents are treated as separate events.

It is an educational publication, not an independent forensic investigation.

Edited by David Haberland · Published by SecureFi, Inc. · November 17, 2026