Incident 829: Major incident affecting Webflow Designer, Dashboard, Marketplace, and other features

Incident Start: July 28, 2025 at 1:27 PM UTC
Incident End: July 31, 2025 at 4:00 PM UTC
Incident Duration: 3 days, 2 hours, and 33 minutes

Summary

Between July 28 at 1:27 PM UTC and July 31 at 4:00 PM UTC, Webflow experienced a major incident that affected the availability of the Designer, Dashboard, Marketplace, and user sign ups across four distinct phases. The first two phases were caused by targeted malicious attacks on API endpoints, resulting in elevated latency and service outages. These attacks were mitigated through firewall protections, IP blocking, and infrastructure investigations. The third and most disruptive phase was triggered by continued attack traffic, combined with performance degradation that occurred after the scaling of a critical backend database cluster to provide operational headroom. The newly scaled cluster introduced severe write latency and replication lag, which persisted until the cluster was scaled down in accordance with vendor guidance. The fourth phase began with another malicious attack, which again caused the critical database cluster to experience severe write latency and replication lag. This phase was mitigated through two separate configuration changes, first by our database vendor and then by scaling the database cluster up using a new hardware architecture. Full stability was restored after these changes, and no further disruptions have been observed since. Throughout all four phases of the incident, all Webflow-hosted websites maintained 100 percent availability.

Webflow takes the needs of its customers very seriously, and recognizes the impact of this incident on your business. Our goal is to provide our customers with an excellent experience, and we sincerely regret that this incident occurred.

Details

This incident occurred over four distinct phases and lasted approximately 3 days, 2 hours, and 33 minutes. During this time Webflow Designer and Dashboard were intermittently down. The root cause for the first two phases appears to be from a malicious attacker producing sustained load on our systems and the root cause for the third and fourth phases appear to be issues with our third-party database provider.

Four phases of the July 28–30 incident. Yellow shows write latency, red shows read latency. Lower is better.

The first phase began at 1:27 PM UTC on July 28, 2025 when we began to receive internal and external reports of increased latency in loading the Webflow Designer and Dashboard. Some of the issues customers may have experienced include long wait times to publish sites and errors when attempting to load the dashboard or designer. We immediately investigated these issues and found some of our API endpoints were being hit by a malicious attack. We put in Web Application Firewall protections, blocked suspect IP address ranges, investigated backend load on our third-party databases (including opening the highest severity case with our third-party vendor), and took actions to improve database efficiency. These attacks were mitigated by 4:55 PM UTC. At this time, Webflow Designer and Dashboard returned to being fully responsive and stable.

The second phase started at 9:03 AM UTC on July 29, 2025 when we saw a new set of attacks to a similar set of API endpoints and we saw increased latency in loading the Webflow Designer and Dashboard. We took immediate additional actions with our Web Application Firewall, blocked more suspect IP address ranges, continued our backend database investigations, and took other corrective actions. These attacks were mitigated by 10:59 AM UTC. At this time, again, Webflow Designer and Dashboard returned to being fully responsive and stable.

The third, longest, and most painful phase for our customers, began at 12:13 PM UTC on July 29, 2025. The malicious attacks and subsequent Designer and Dashboard latencies started again. As before, we blocked suspect IP address ranges, applied operational fixes, examined backend database load in detail, and the issues appeared to be headed toward mitigation by approximately 2:30 PM UTC. At this time, however, we started seeing normal weekday traffic load and our infrastructure began its standard scaling operations. But, because of the additional load from the attacks and the fact that they were not fully mitigated, with guidance from our vendor we decided to scale up a critical backend database cluster using our third-party vendor's automation to give us more operational headroom. This operational scaling of the cluster was completed at 2:50 PM UTC. Shortly thereafter, Webflow Designer and Dashboard latencies increased significantly. For the next 8 hours the availability of the Designer and Dashboard was intermittent, at best, and often down.

The critical backend database cluster that we scaled up immediately increased to over 300 times the normal write latency and replication lag across the database cluster increased to over

500 times the normal replication lag. We quickly escalated to our third-party vendor and took several concurrent actions during this time period to reduce load on the backend database. These actions included turning off data pipelines to reduce load, tuning off several newly launched features, blocking new user sign ups, disabling SCIM, and other mitigating actions. All other engineering and operational work was halted while every available resource worked on this issue. Each of our corrective actions had limited positive effect on the Webflow Designer and Dashboard but the backend database cluster write latency and replication lag were persistent issues. Disabling newly launched features and new user sign ups, for example, did not help mitigate the incident.

P90 (purple) and P50 (light blue) Designer latencies in phase three. Lower is better.

At approximately 8:00 PM UTC, our third-party database vendor gave us guidance to scale down the critical backend database cluster to a smaller cluster size. Their recommendation was based on their assessment that their larger database cluster was using a dual-socket CPU architecture and the smaller instance size was using a single-socket CPU architecture. While the larger database cluster size was an available upgrade and shown to be higher performance, we later learned that the single-socket CPU architecture performs better for our services. We began the scale-down of the critical database cluster to the single-socket CPU architecture and this operation completed at 10:09 PM UTC. Nearly immediately afterwards, the Webflow Designer and Dashboard returned to being fully responsive and stable.

The fourth phase of this incident began at 9:32 AM UTC on July 30, 2025, when a malicious actor initiated an attack against the Webflow Marketplace services, triggering high CPU usage alerts on our critical database cluster. Although Webflow remained initially stable due to caching, error handling, and auto-scaling mechanisms implemented during earlier phases, certain session-related API calls began to introduce elevated write latency. This caused performance degradation in both the Webflow Designer and Dashboard.

We promptly escalated the issue to our third-party database provider and engaged with their team on a live video call for the next 14 hours. To reduce database load, we temporarily removed the attack surface by taking the Marketplace offline, which helped alleviate memory contention on the cluster. As a precautionary measure against a recurrence of earlier database issues, we also disabled new user sign ups. In addition, we optimized specific read operations

and partially redirected database queries to secondary replicas, which helped decrease pressure on the primary nodes. Based on feedback from our vendor, we performed a database cluster failover, which provided immediate relief at 10:18 AM UTC. Webflow Designer and Dashboard returned to being fully responsive and stable.

During this part of the investigation, our database vendor informed us of a known bug in their software that may be a contributing factor to the sudden increase in database cluster write latency seen throughout this incident. This known bug involves controlling the number of sessions that a database cluster can handle at a given time, called the session count. Their engineering team is working on a fix but does not have a solution for us to deploy at this time. We were instructed to monitor the health of our systems and if a subsequent event occurred, to work with them on possible mitigations.

The next part of the incident began at 12:06 PM UTC, without any changes initiated by the Webflow team. During this time, we observed a significant increase in traffic and database operations, which led to high write latencies on the critical database cluster. While monitoring continued, we remained in active discussions with our third-party database vendor over the ongoing video call. On their recommendation, we agreed to proceed with a configuration change to reduce logging of slow queries on the cluster, aiming to alleviate performance pressure. This adjustment by the database vendor led to a quick recovery, and by 1:40 PM UTC, the Webflow Designer and Dashboard returned to being fully responsive and stable and we kept Marketplace offline and user sign ups disabled.

At 4:26 PM UTC we again observed elevated write latency on the critical database cluster. Despite the increased latency, we did not observe any degradation in the performance of the Webflow Designer or Dashboard. We believe this was because of previous mitigations and optimizations deployed during this incident. We continued our live video investigation with our database vendor to assess the issue. Based on their analysis, the vendor recommended reverting a previous operating system configuration change they had earlier advised. Specifically, they disabled aggressive memory decommit, a setting that had previously been enabled to improve memory management. This change was applied at 5:58 PM UTC, and shortly afterward, database write latency returned to expected levels.

Lastly, to give ourselves operational headroom going forward, we upgraded the critical database cluster to a larger capacity that uses the single-socket CPU architecture. This operation completed at 5:59 PM UTC successfully and we continued to see improvements in Webflow and database performance.

The fourth phase. Yellow is database write latency, red is read latency. Lower is better.

We have since taken measured and deliberate steps to re-enable many of the features that were disabled during the incident. No further performance degradation of the Webflow Designer and Dashboard have been observed, although we remain on constant alert and vigilant for a repeat of these issues.

Out of an abundance of caution given the long duration and multiple phases of this incident, we continued our continuous monitoring, video calls with our database vendor, and remained on high alert until 4:00 PM UTC on July 31. We carefully monitored site traffic, database write latencies, and replication lag and found no abnormalities. Webflow Dashboard, Designer, and other features remain performing as expected.

Latency after 5:59 PM UTC on July 30. Yellow is database write latency, red is read latency. Lower is better.

Action Items

As a result of this incident, Webflow completed the following actions in order to improve our systems and ability to prevent analogous situations from occurring in the future, including:

- Added an index to a collection in the impacted database cluster to increase query efficiency and reduce database load
- Implemented rate-limiting for our user sign up systems to better mitigate spikes in traffic

- Increased rate-limit protections in our Web Application Firewall to block malicious traffic in targeted areas
- Introduced circuit breakers for key traffic flows that contribute significant load to the affected database cluster
- Enhanced monitoring to ensure more reliable alerting when database latency issues arise
- Upgraded our critical database cluster to a higher-capacity, single-socket CPU architecture

To be completed
- Replay missed form submissions for all customer forms (by August 4)
- Tune heartbeat configurations to improve the health of database connection pools (by August 1)
- Adjust backup and snapshot schedules to avoid load during peak usage hours (by August 4)
- Evaluate and potentially move additional read-only queries to a dedicated replica (by August 4)
- Evaluate the use of a queuing system for non-critical write requests to allow for eventual consistency (by August 4)
- Complete a detailed root cause analysis with our database vendor for performance recommendations (by August 1)
- Finalize internal root cause analysis with additional follow-up actions (by August 4)
- Upgrade database clusters to the latest software version (by August 15)
- Deploy the fix for the session count bug identified by our database vendor when it becomes available