

Summary

On April 14, 2026, beginning at 14:06 UTC, Webflow customers experienced 5xx errors when accessing hosted websites, the Dashboard, the Designer, webflow.com, form submissions, and API endpoints.

This incident was not caused by a security vulnerability, a malicious attack, or any of the recent public supply-chain security incidents, nor was it related to a Webflow feature release. Rather, the root cause was a previously unknown and undocumented storage limit on a database cluster operated by our cloud service provider. One cluster within our content management system (CMS) infrastructure reached capacity, did not have any available storage, and could not be restarted.

As a result, our application servers were unable to connect to the affected database cluster, rendering Webflow platform features and functionality unavailable during this period. Webflow-hosted sites cached by our CDN remained accessible; however, uncached or recently published sites may have returned 5xx errors.

For a large majority of our sites (approximately 96%), the incident was resolved without data loss by 16:21 UTC on April 14. The remaining 4% of our sites were resolved by 04:01 UTC on April 15.

Webflow takes the reliability of our platform seriously and recognizes the impact this incident had on your business. Our goal is to provide an excellent experience, and we sincerely regret that this disruption occurred.

Details

14:06 UTC, April 14, 2026. Webflow became aware of elevated error rates and database connectivity failures. Our team quickly determined that one CMS database cluster had failed and that we were unable to restart it through standard automation and operational procedures. We observed that storage capacity on the cluster had dropped from approximately 1.74 TiB to 0 TiB, effectively leaving no CMS data available to serve on sites for a subset of our customers. We immediately took action to bring the database cluster and its storage back online but were unsuccessful. We did not observe any CMS data loss but the database cluster was not able to access any storage capacity.

14:15 UTC. The time from incident creation to drafting the status page was 5 minutes, and it went live 9 minutes after the incident was created.

We quickly established a crisis video call with our cloud provider, who told us that Webflow was not the only customer affected and that they were escalating multiple issues on their side to resolve the cluster storage issue.



After several hours of escalations and discussions, we learned that the cluster could not be restarted because the affected shard had reached the 128 TiB logical storage allocation limit of its database engine version. This logical allocation limit was unexpected and went undetected by our monitoring because the cloud provider's dashboard metrics reported only approximately 1.74 TiB of used storage on the cluster, representing roughly 1.35% utilization.

We later learned that the database engine that we use for CMS reserves approximately 2 MiB of logical space per database file, and these reservations persist even after data is deleted, only being reclaimed through a database engine version upgrade. Over time, based on our CMS usage, approximately 66 million total database files accumulated on the cluster. These persistent reservations silently consumed the full 128 TiB allocation cap while reported data usage showed a peak usage of 1.74 TiB.

When the database instance attempted to restart, the database engine could not allocate any additional logical storage against the exhausted allocation cap. As a result, the instance failed to initialize and entered a crash-restart loop.

After learning about the silent consumption of the full 128 TiB logical allocation cap being the root cause for this event, we were informed by our cloud provider that this issue was isolated to Webflow and did not affect other customers at the same time.

Because the Webflow application layer treats any unresponsive database cluster as a fatal startup condition, application pods that were starting or restarting during this window could not come online, regardless of which cluster their customers' data resided on, expanding the impact well beyond the affected cluster alone.

The team pursued two parallel recovery workstreams: removing the affected cluster from the active cluster group to restore service for the majority of customers, and initiating a point-in-time restore to ensure the cluster's data remained stable.

16:21 UTC. The change to remove the failing cluster was manually deployed. With the failing cluster removed, the remaining four clusters resumed serving traffic normally. The application pods returned to normal operation and were operational. At this point in time, a large majority of Webflow customers were able to access hosted websites, the Dashboard, the Designer, webflow.com, form submissions, and API endpoints.

16:49 UTC. The status page was updated confirming service restoration for the majority of customers.

16:55 UTC. A point-in-time restore of the affected shard was initiated. As a precautionary measure during recovery, Webflow Canvas automated backup and restore operations were temporarily disabled to reduce additional load on the database clusters.

22:46 UTC. Our cloud provider manually provisioned an additional 160 GiB of storage for the database cluster to provide the capacity needed to perform restoration operations.



02:00 UTC, April 15, 2026. The restoration process remained ongoing, with a small number of sites (4%) still affected at this time. We remained engaged with our cloud provider on a video conference call to work through the remainder of the incident.

04:01 UTC. We have restored the failing database cluster and our cloud provider has provided doubling of the logical allocation capacity to 256 TiB. We confirmed restoration of service to 100% of all Webflow customers. Incident resolved.

Action Items

As a result of this incident, Webflow has taken or plans to take the following actions to improve our systems and ability to prevent analogous situations from occurring in the future:

- **Restored Webflow automated backup and restore operations.** We re-enabled the automated backup and restore operations that were temporarily disabled to reduce additional load on the database clusters.
- **Upgraded all CMS database clusters to a newer engine version.** The upgraded version raises the logical storage allocation limit from 128 TiB to 256 TiB, directly resolving the logical allocation constraint that caused this incident.
- **Added monitoring and alerting for database logical storage allocation.** We have added monitoring and alerting for database logical storage allocation. We now track the total number of database files per cluster as a proxy for logical space consumption, with alerting thresholds that trigger well before clusters approach engine limits. We have also deployed monitoring that tracks actual logical storage consumption on every cluster, rather than relying on incorrectly reported data usage.
- **Add additional shards to the database clusters.** This will reduce the number of files per database cluster and provide horizontal scaling for the database clusters. Target date: April 30, 2026
- **Isolate database shard failures from application availability.** Modify application startup behavior so unavailable database shards do not prevent the application from starting for customers whose data resides on healthy shards. Target date: April 30, 2026
- **Build rapid site-publishing circuit breakers.** Introduce operational controls to pause or limit high-write operations such as site publishing during incidents, reducing load on degraded clusters and helping maintain availability of cached content. Target date: TBD