BLOG / UPDATE

Update · DEVELOPMENT · 4 min read · Sep 8, 2026

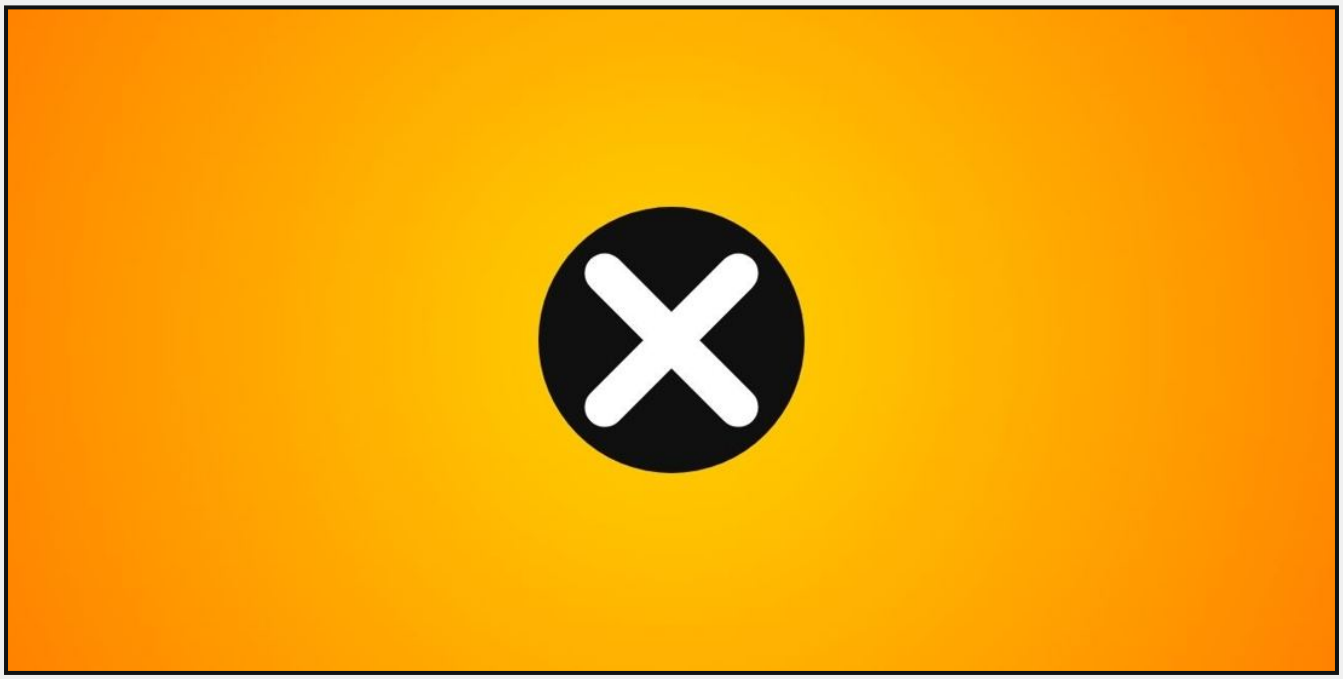
⚠ **This article is over 1 month old.**

We made a mistake

How a simple update led to hours of downtime that didn't even exist.

Updated 4 Nov 2026

Read online: https://traderange.net/blog/sorry-cwqx7sq9/



Some of you may have noticed our extended downtime today. We did too. Having downtime 12 hours before our biggest ad campaign is terrifying and even worse - our first hot fix crashed the site entirely. This is the detailed account of what happened and the measures we have taken to avoid it in future.

Sep 7

Just 2 hours earlier we had been informed our bid for advertising has been accepted with a live start on the 9th of September. We were ecstatic. At 22:14 commit was deployed from our test server to the production one that added the https://trge.link domain to our active production as a short share link domain designed to remove the insanely long links the main site has. The deploy failed to pass the test and the workflow flagged it. At this point the live site was unaffected.

An issue was identified where the https://trge.link domain was missing a private key for the certificate. This error was expected as the workspace never stores the private keys and only the live server does. The issue was overridden and a deploy was forced. at 22:27 the update went live. Slowly the global cache was updated to generate links using the new domains and no bugs were noticed. All tests passed.

Sep 8

On the morning of September the 8th at 06:22 was the first time we noticed something was wrong. The home page lacked any sections below the top banner. This was an abnormal issue not visibly related to any deployments the day prior. We searched for the bug and located an escaped string. This seemed like an easy fix, which was implemented within minutes of being found. At 8:39 AM the site went live again. At 8:44 the home page returned a 404 error.

A combination of minor issues had built up. The major cause of the bug was found and rectified. Our internal description of the issue is below.

> *Every snapshot write is chmod'd 0644 before rename, and start.sh chmods leftover 0600 pages before gunicorn binds.*

This was believed to have been fixed weeks earlier by a dev who was unavailable when we were working on these fixes. At 09:09 another fix was attempted and it appeared to have rectified some of the issues. Around 86% of the site was operational, but specific tests still failed and operation was disrupted.

By 09:16 the entire dev team was online working on a fix. An issue was noticed in how pages ending in .html were being routed to the non-.html version regardless of that page's existence. This transition was effected days earlier and never experienced such issues.

This was fixed at 09:31 and the site went live. With hours left before our campaign launches. We expect that not all the errors have been caught. We are now performing tests and will gradually restore to 100% functionality by 22:00 today.

> *We thank you for your support during this time and apologise for any disruptions. This is not the standard we work towards, but we are proud of our team's fast and professional response.*

Futureproofing

This is hardly a standard situation. We have now limited which developers have the ability to override test fails and also introduced new tests to reduce the fail surface. Things like this do happen, but they should NEVER reach our production server. We apologise.

TradeRange https://traderange.net/blog/sorry-cwqx7sq9/