OpusFlow

Service Incident Report:

November 18, 2025

Date: November 18, 2025

Incident Duration: 12:48 – 15:41 (CET)

Status: Resolved

Visit Now >

opusflow.io

OpusFlow

Executive Summary

On November 18, 2025, our platform experienced a service disruption impacting PDF generation, Chart generation, and user authentication for customers in specific regions. This incident was directly caused by a major global outage affecting Cloudflare, a critical infrastructure provider used by a significant portion of the internet, including our services.

While the root cause was external, our engineering team successfully restored full functionality by migrating affected services to an alternative cluster, bypassing the impacted infrastructure. This report details the incident, its impact on your experience, and the measures we are implementing to ensure greater resilience in the future.

Impact Analysis

During the incident window, customers may have experienced the following issues:

- **PDF & Chart Generation:**
  - **Phase 1 (12:48 – 13:19 CET):** Intermittent failures. Users attempting to generate reports or visualize data may have seen occasional error messages or timeouts.
  - **Phase 2 (13:19 – 15:30 CET):** As the upstream provider issues escalated, these services experienced a near 100% failure rate, rendering document and chart creation unavailable.
- **Regional Login Issues (Belgium):**
  - A subset of users, specifically those accessing the platform from Belgium, experienced difficulties logging in. These authentication pathways were routed through the affected Cloudflare infrastructure and were restored immediately once the upstream provider resolved their internal configuration errors.

OpusFlow

Timeline of Events (Amsterdam Time / CET)

- **12:48:** Cloudflare officially acknowledges an internal service degradation and begins investigation. Simultaneously, our monitoring systems detect the first elevated error rates in the PDF and Chart generation microservices.
- **13:00:** Engineering teams identify the correlation between our service degradation and the global Cloudflare outage. Intermittent failures begin to escalate.
- **14:09:** Cloudflare identifies the issue and begins implementing a fix.
- **14:30:** Despite upstream progress, failure rates for our chart and PDF generation reach critical levels (100%) due to continued instability in the edge network.
- **15:15:** A mitigation strategy is finalized. The decision is made to migrate the Chart and PDF generation workloads to a backup cluster that does not rely on the impacted Cloudflare ingress points.
- **15:30:** Mitigation Implemented. Migration of our internal services commences. Traffic is rerouted to the backup cluster.
- **15:41:** Service Restored. Migration completed. Success rates for PDF and Chart generation return to 100%.
- **15:42:** Cloudflare reports a fix has been implemented and they are monitoring the results. Our services were already stable at this point due to the internal migration.
- **20:28:** Cloudflare marks the global incident as fully resolved.

Root Cause

The primary trigger for this incident was a configuration error within Cloudflare's global network. According to their status reports (Incident 8gmgl950y3h7), an internal service degradation disrupted traffic for thousands of online services, including ours. This resulted in our application traffic being dropped or mishandled before it could reach our servers.

OpusFlow

Resolution and Mitigation

Our immediate priority was restoring service availability regardless of the status of our upstream provider.

Once it became clear that the upstream issues would not be resolved quickly, our infrastructure team initiated an emergency failover. We moved the affected PDF and Chart generation services to a secondary cluster configured with a direct traffic ingress that bypassed the failing Cloudflare nodes. This action immediately restored service health independent of the ongoing global outage.

Future Prevention Measures

We are committed to learning from every incident to improve our platform's reliability. While we cannot control the uptime of external providers, we can control how our system reacts to their failure.

To prevent a recurrence of this specific failure mode, we are implementing the following architectural changes:

1. **Multi-Channel Traffic Routing:** We are formally operationalizing the failover strategy used today. We are establishing permanent redundant traffic channels.
2. **Automated Failover:** We are developing automated health checks that will detect specific upstream provider failures and automatically re-route traffic to backup clusters. This will reduce the time to recovery from hours to minutes, often before users even notice a degradation.
3. **Regional Redundancy:** We are auditing our login flows to ensure that regional access points (like those in Belgium) have diverse routing options available to bypass localized edge network failures.

We apologize for the disruption this caused to your workday and thank you for your patience as we navigated this global internet event.