



Reason For Outage (RfO) Report Incident Details

On December 8th, 2022 from 6:23am to 1:03pm EST Zola Suite experienced an interruption of services which prevented users from logging in. The cause of the outage was a series of events which led to a primary database going offline. The storage device became maxed out from a job intended to improve performance, but one that didn't follow our internal process for review.

This incident was human error and a breakdown of supporting monitors, alerts and other protections in place to prevent its occurrence. This incident **was not** a breach in security and no customer data was lost. A more detailed explanation is below with a timeline and resolution summary of what we are improving to prevent this from happening again.

We understand the negative impact to your business. Rest assured, our teams are taking all necessary steps to harden our infrastructure for reliability and implement enterprise level technologies with a focus on performance and stability.

Incident Timelines

All times in EST Eastern Standard Time

December 8th, 2022

6:23 am	Monitoring issued an alert to Network Operations Center (NOC) that a primary database was not responding to requests.
6:25 am	Escalation process began – team members alerted, and diagnostics began.
7:40 am	Identified that the transaction log from the database expanded in size to fill the storage device.
8:45 am	To shrink the log and recover disk space, several running jobs were placed on hold and database was placed into recovery mode – a standard operation for this procedure.
9:00 am	Job hold completed and waited for running tasks to complete.
9:15 am	Performed a database log shrink after removing the database from Availability Group Cluster (AGC). Placed the database in Simple Recovery Mode to free up space.
9:25 am	Executed process to bring primary database node online.
11:58 am	Informed Zola Quality Assurance team primary node was coming back online and prepare testing prior to placing into production mode.
1:03 pm	Quality testing completed. Database online accepting transactions. Notification to clients sent that services are restored.



Corrective Actions

- Created stronger alert correlation from drive space warning and database availability
- Increased staff training on our change control process
- Improved escalation process to accelerate communications internally and to customers/partners
- Changed primary database to a high availability cluster with a minimum of three servers vs mirror pair

Resolution Summary

The root cause of this failure was a decision to run a database index job to improve performance within the application. The job created an unexpected number of logs which filled the primary storage device.

We have a process in place to review all activity prior to deploying into production. In this case, due to human error, this process wasn't followed and with the additional challenges with monitoring (turned off during a recent upgrade but not re-enabled prior to running this job), this outage was unplanned and unexpected.

Incident Categorization

Incident Reference: INC-60

Incident Date: 12/8/2022

Severity Level: Severity 1 – Service Failure

Communication Date: 12/16/2022

End of RFO