

[Resolved] ClubRunner Service Interruption

2021-04-01 - Zach W. - Service Issues

[APR 01 2021 5:30 PM]

Starting at approximately 5:15 PM [NAEDT](#) our team noticed that there were issues accessing some parts of ClubRunner. We can confirm that ClubRunner is experiencing an outage due to a disruption on Microsoft Azure's end.

We are monitoring the status on <https://status2.azure.com/> and will post updates as we learn more.

We apologize for the inconvenience.

[APR 05 2021 9:00 AM]

As our team continued to monitor on Thursday evening, we observed that all systems had recovered and were fully operational by 8:30 PM NAEDT on Apr 01 2021.

Microsoft was able to resolve the issue. Here is the summary of events from their status page:

RCA - DNS issue impacting multiple Microsoft services (Tracking ID GVV5-TZZ)

Summary of Impact: Between 21:21 UTC and 22:00 UTC on 1 Apr 2021, Azure DNS experienced a service availability issue. This resulted in customers being unable to resolve domain names for services they use, which resulted in intermittent failures accessing or managing Azure and Microsoft services. Due to the nature of DNS, the impact of the issue was observed across multiple regions. Recovery time varied by service, but the majority of services recovered by 22:30 UTC.

Root Cause: Azure DNS servers experienced an anomalous surge in DNS queries from across the globe targeting a set of domains hosted on Azure. Normally, Azure's layers of caches and traffic shaping would mitigate this surge. In this incident, one specific sequence of events exposed a code defect in our DNS service that reduced the efficiency of our DNS Edge caches. As our DNS service became overloaded, DNS clients began frequent retries of their requests which added workload to the DNS service. Since client retries are considered legitimate DNS traffic, this traffic was not dropped by our volumetric spike mitigation systems. This increase in traffic led to decreased

availability of our DNS service.

Mitigation: The decrease in service availability triggered our monitoring systems and engaged our engineers. Our DNS services automatically recovered themselves by 22:00 UTC. This recovery time exceeded our design goal, and our engineers prepared additional serving capacity and the ability to answer DNS queries from the volumetric spike mitigation system in case further mitigation steps were needed. The majority of services were fully recovered by 22:30 UTC. Immediately after the incident, we updated the logic on the volumetric spike mitigation system to protect the DNS service from excessive retries.

Next Steps: We apologize for the impact to affected customers. We are continuously taking steps to improve the Microsoft Azure Platform and our processes to help ensure such incidents do not occur in the future. In this case, this includes (but is not limited to):

- Repair the code defect so that all requests can be efficiently handled in cache.
- Improve the automatic detection and mitigation of anomalous traffic patterns.