



Reason For Outage Report (RFO)


Incident: Connectivity Issues on US-EAST-2 on 20th of February 2024 at 20:40 UTC

London, March 29, 2024

This is the Reason For Outage Report regarding the connectivity issues on US-East-2 datacenter that occurred on February 20th, 2024. This report details the incident timeline, the emergency response actions taken to mitigate the issue, a comprehensive post-mortem with root-cause analysis, recommendations and measures that we have committed to implement in order to avoid similar occurrences from happening again in the future.

Incident Timeline and Actions Taken

- **[2024-02-20 at 19:51 UTC]: Incident Identified** - A continuous increase in latencies across nodes of multiple clusters within the Ashburn Datacenter was identified. Immediate inspection revealed extreme slowness affecting every system in the region. Examination began while maintaining stability and searching for any possible clues about the cause of the issue.
- **[2024-02-20 at 22:31 UTC]: Incident Isolation** - No apparent configuration issues were found with our web infrastructure. After isolating the incident to the Ashburn region and observing high network latency, bandwidth, and packet loss, we concluded that the issue was related to the region's upstream network provider.
- **[2024-02-21 at 01:40 UTC]: Escalation** - Discussions with our upstream network provider continued, and our team still observed major issues and reported no change in severity (high-priority emergency). All network components indicated that everything was working properly. Our efforts shifted towards examining every aspect of our web stack and other variables.
- **[2024-02-22 at 06:27 UTC]: Further Investigation** - We had indications that the cause of the issue could be in our storage implementation. Investigations shifted to this direction.



- **[2024-02-22 at 10:54 UTC]: Mitigation** - After conducting repeated checks, we were unable to identify any issues with our storage systems. Since no other avenue was left to explore, we decided to replace network components that may affect our storage systems. **Subsequent on-site examination revealed that the root cause of the network issues, including high latencies and connectivity disruptions, was attributed to a faulty network component that was misreporting its status as Healthy.**
- **[2024-02-22 at 15:55 UTC]: Post Mitigation Actions** - After the faulty hardware was replaced by the Data Center team, network latency returned to normal levels. Our focus shifted to reverting any temporary mitigations we had in place and monitoring the situation.

Incident Background

On February 20th, our Ashburn data center experienced severe performance degradation, resulting in significant latency increases across multiple clusters. The primary contributing factor was a malfunction in a critical network component which failed to maintain its expected bandwidth capacity.

This component falsely indicated that it was operating at full capacity and was healthy, leading to a delay in identifying the root cause of the performance degradation.

Comprehensive Post-Mortem and Root-Cause Analysis

A series of concerning issues emerged during a performance degradation, including:

Increased IO latency: Across the network storage, affecting all servers attached to the specific storage in the region.


Bandwidth Reduction: The malfunctioning network component caused a significant reduction in available bandwidth, impacting the speed and reliability of network communication.

Latency and Packet Loss: The decreased bandwidth resulted in increased latency and packet loss, affecting the responsiveness and stability of network-dependent services.

Upon detecting anomalies in storage performance and elsewhere, an immediate investigation was initiated to identify the root cause of the degradation. Given that all network components indicated proper functionality, our attention shifted to the storage systems.

This false indication derailed our troubleshooting efforts for more than a day during which the hosted websites were underperforming. After conducting repeated thorough checks on the storage systems and finding nothing, we decided to shift our troubleshooting focus back to network components.

Despite the absence of any indicators suggesting a problem, we opted to replace certain network



components which led us to resolve the issue. In the meantime, we implemented temporary measures such as rerouting traffic or reallocating resources, to mitigate the immediate impact of the degraded performance.

Lessons Learned and Future Measures

Unfortunately, our initial investigation led us down an incorrect path. While storage performance anomalies were present, they were a secondary symptom, and not directly related to the root cause. This resulted in us spending significant time troubleshooting storage sub-systems without finding the real issue. Thus, we made the decision to revisit the network components despite the lack of initial red flags.

This move ultimately led to the identification of the faulty component and a swift restoration of full functionality.

Moving forward, we have taken measures to prevent such incidents from occurring again. These include, the consideration of multiple factors and system events so that our team can be alerted for uncommon behaviors, **even in the absence of indications of hardware failure**. In addition to this:

Proactive Hardware replacement - The malfunctioning network component was replaced and the full bandwidth capacity was restored.

Monitoring and alerting - We improved our network monitoring and alerting systems to proactively detect and notify us of anomalies in network performance.

These measures reflect our ongoing commitment to providing reliable, secure, and high-performance hosting solutions.

We apologize for any interruptions caused by this incident and recognize the potential impact it may have caused to our customers' operations. We appreciate your patience and understanding.

At Pressidium, we are dedicated to delivering to all of our customers and partners the most reliable, dependable and productive experience in the industry.

Thank you for your continued trust and commitment.

Team Pressidium

For assistance or further information, please open a ticket from the [Pressidium Dashboard](#) or contact us at support-at-pressidium.com. Follow us on [X](#) & [Facebook](#).