Preliminary Post Incident Report
For
Microsoft 365

Report Date: November 15, 2026

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

The descriptions of other companies' products in this document, if any, are provided only as a convenience to you. Any such references should not be considered an endorsement or support by Microsoft. Microsoft cannot guarantee their accuracy, and the products may change over time. Also, the descriptions are intended as brief highlights to aid understanding, rather than as thorough coverage. For authoritative descriptions of these products, please consult their respective manufacturers.

© 2026 Microsoft Corporation. All rights reserved. Any use or distribution of these materials without express authorization of Microsoft Corp. is strictly prohibited.

Microsoft and Windows are either registered trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Incident Information

Important Note	This is a preliminary Post Incident Report (PIR) that is being delivered to provide early insight into details of the issue. The information in this PIR is preliminary and subject to change. A final PIR will be provided within five (5) business days from full event resolution and will supersede this document upon publication.
Incident ID	EX1550543/MO1550821
Incident Title	Users may be unable to access or perform several actions in Microsoft 365 services Users may have been unable to connect to their mailboxes in Exchange Online through multiple connection methods.
Service(s) Impacted	Microsoft 365 services, including Exchange, Copilot, Teams, SharePoint, OneDrive, Defender, Purview, Fabric, Intune, Planner, Power Apps, To Do, Power BI, and Universal Print.

User Impact

Primary impact scenario: The most significant impact affected Exchange Online, with most users being unable to access their mailbox through Outlook or other mail clients. Additionally, email delivery was delayed for many users.

Secondary impact scenarios: Other services that utilize capabilities of Exchange Online were impacted, though most of these services largely remained online and accessible. This included, but was not limited to, elements of Copilot, Teams, SharePoint, OneDrive, Defender, and Purview that interact with Exchange Online.

Scope of Impact

This event affected most users when they accessed or used features that rely on Exchange Online. The severity of impact varied by region, connection method, and time over the course of the incident. Partial or intermittent access may have occurred while the service was degraded.

Incident Start Date and Time

Thursday, November 12, 2026, at 4:20 PM UTC

Incident End Date and Time

Friday, November 13, 2026, at 4:10 PM UTC

Root Cause

The incident on Thursday, November 12th, 2026, was caused by three main factors that together caused a global outage for the networking gateway that fronts many Microsoft 365 services.

First, a latent defect caused Windows servers to become overwhelmed and unresponsive under certain load conditions. This was due to lock contention inside an internal Windows process which caused the machine's resources (CPU & memory) to be exhausted. This defect was the result of a fix that was prepared specifically for Microsoft 365 networking gateway servers; neither the underlying defect nor the fix applies to broader Windows deployments.

The initial change in behavior first started deployment in July of this year. The update went through safe deployment and saturated globally over several weeks, during which time no issues were observed. For more than three months, the change was deployed worldwide without impact.

Our normal load testing did not detect the latent issue because the traffic mix and SKUs used in our labs. The contention issue is worsened on machines with more cores and newer hardware; whereas our load testing tended to focus on weaker hardware with fewer cores.

Second, a surge in load increased the likelihood for the networking gateways to encounter the contention. Within Microsoft 365, we maintain two independent and heterogenous network gateway systems designed for isolation and high availability. This allows many Microsoft 365 services to fail to/from them during outages or disruptions. Approximately fifteen hours before the outage, Teams failed over from the alternative networking gateway to the 'primary' gateway that also supports Exchange Online. This increased the load on the networking gateway and changed its traffic patterns. This type of failover is part of normal operations and is integral to standard Microsoft resiliency efforts.

Third, at the same time, nearly 14% of server capacity for the networking gateway was out of service due to an unrelated issue from the previous day. This server capacity was on track to be returned to active service through routine maintenance and health processes. The service had sufficient capacity headroom, so the reduction in server capacity had no impact on its own.

Taken together, the increased load and the reduced server capacity, caused the network gateway to encounter the latent bug from July. This caused requests to build up until servers ultimately stopped responding. This occurred during the peak window of traffic in the United States.

When servers lose connectivity, our traffic management systems automatically route users to other healthy endpoints. This happens often and is critical to avoiding any interruption due to hardware or software defects and is normally what makes the service resilient. However, the redirected traffic increased the likelihood for the same problem to occur on the remaining endpoints. This caused a cascading failure from server to server as the same amount of load was spread across fewer servers. In less than two hours, the previously unknown defect was encountered across our fleet and rendered them unresponsive. At its lowest, more than 80% of servers were unresponsive.

Normally, the network gateway is composed of thousands of 'capacity units' which represent shards of capacity, hardware, and networking infrastructure. These 'capacity units' are then organized into regions. In this case, due to a separate issue, the normal bulkheads between these

different partitions did not stop the cascading failure. This is what allowed the impact to spread across our network gateway infrastructure globally.

Those services that were able – like Teams, OneDrive, SharePoint, and Office apps – failed out of the faulty network gateway infrastructure. That is why those services were partially restored earlier, with Teams chats and meetings functioning without issue for most of the outage. But not all services, like Exchange Online, can fail out of this network gateway for scale reasons.

Recovery took longer than usual because as servers returned to service, they were immediately overwhelmed by waiting traffic and became unresponsive again. Our normally automated recovery processes were not designed to restore service from this broad unresponsive state. Additionally, since the main code defect was a latent bug from several months previous, our normal rollback, flighting, and feature flags were unable to quickly revert the change.

After engineers confirmed the defect as the cause, we developed an update that returned the code path to the previously known good state in Windows. To effectively apply this fix, and to slowly bring the service online again, we took groups of servers out of service, applied the change, and restarted them before returning them to service. This process was carefully applied to large quantities of infrastructure at a time (e.g. 50%+ of a region).

The final Post-Incident Review (PIR) will provide more details.

Actions Taken (All times UTC)

Thursday, November 12

01:07AM - The automated failover system moved Teams traffic onto the network gateway front end servers shared with Exchange Online.
1:07 AM – 4:00 PM – Load on the front-end servers increased gradually. Isolated server failures began in some regions during this period, automated mitigations applied.
4:03 PM – Our connectivity anomaly detection monitor identified an increase in end-user connectivity issues.
4:17 PM – Due to the increase in connectivity issues observed, automation began to redirect traffic to alternate healthy endpoints.
4:20 PM – We received a high-priority alert indicating impact to Exchange Online services, which triggered immediate communication to the Service Health Dashboard under EX1550543.
4:25 PM – We confirmed that customers were experiencing issues across multiple Exchange connection methods, including Outlook, Outlook on the web, Exchange ActiveSync, and Messaging and REST APIs. We began reviewing internal telemetry.
4:33 PM – We received the first support case associated with the broader incident, confirming that customers were reporting issues alongside those already detected through our monitoring. Users reports began to accumulate on 3rd party outage reporting sites such as DownDetector.
5:03 PM – Our connectivity anomaly detection monitors began to see a further increase in connectivity issues.
5:32 PM – Automated recovery actions continued to redirect traffic to healthy endpoints; however, further service degradation was detected across multiple regions.
5:45 PM - We identified connectivity errors occurring on a portion of infrastructure which receives and processes users' access requests. We collected network logs to further investigate the source of the errors.
6:00 PM - We progressively restarted and attempted to repair impacted portions of our infrastructure and worked to redirect traffic to alternate infrastructure to provide relief.

6:27 PM – We published additional communication under MO1550821 to the Service Health Dashboard, describing impact to additional services affected by this event, including Microsoft Teams, Copilot, and Defender.
6:33 PM – We triggered failovers to ensure traffic was distributed evenly across our infrastructure and components that manage user connections.
7:00 PM – We expanded our response to accelerate mitigation efforts. As a precaution, we paused all ongoing and planned deployments and configuration changes.
7:20 PM – We determined that the elevated load we were observing was not driven by an overall increase in traffic, but by many servers being unexpectedly unavailable, which was shifting load onto and overwhelming the remaining infrastructure.
8:00 PM – We deployed a targeted change to bypass components we found to be triggering failed health checks on affected servers while we also investigated why those servers were showing elevated processor usage while handling requests.
8:30 PM – While attempting to validate the effectiveness of our targeted change, we found that many of the affected servers had become unresponsive. We began work to restart those servers.
9:00 PM – We prepared an alternative health check method on a subset of affected infrastructure, which allowed the servers in that subset to successfully return to service, and we began applying that change more broadly. In parallel, we prepared an automated recovery process to continually detect and restart unresponsive servers, to help keep them in service while the change progressed.
10:00 PM - Our telemetry and signals confirmed that the actions we'd taken to restore service had not been effective, and we broadened the scope of our investigation and response accordingly. We began analyzing traffic patterns across our entire service infrastructure and evaluating whether isolating traffic within each region would allow affected regions to recover independently.
11:00 PM - Work was underway to return a portion of the affected servers to service and shift traffic and connections to alternative pathways wherever possible. We observed limited recovery in some regions. To further support recovery, we implemented protective traffic management measures in a specific location to reduce load on the affected infrastructure, with plans to expand these measures to additional locations if we observed positive results.
11:55 PM - We determined that the servers we'd returned to service initially remained healthy but then become unresponsive again upon receiving traffic.

Friday, November 13
12:30 AM – We identified the connection negotiation as the primary area responsible for the widespread access failures and determined that competing processes were preventing connections from being handled successfully.
1:00 AM - Our investigation found that a process on the affected servers responsible for establishing connections was becoming overwhelmed, causing the servers to stop processing requests.
1:15 AM – We introduced measures to reduce pressure on the affected systems, including adjustments to monitoring activity and service load. These measures were partially successful and helped stabilize the service while we continued working toward broader recovery.
1:53 AM – We disabled some automated traffic redirection mechanisms performed by our connectivity anomaly detection monitors.
2:10 AM – We applied infrastructure capacity improvements where possible in an effort to help support traffic in regions coming online.
3:00 AM – While monitoring the progression of the changes applied to improve service health, we determined those changes were not applying as expected and established new workstreams in response.
4:00 AM – We worked to improve service stability through traffic protocol adjustments, traffic redistribution, and strategic traffic blocking measures. We observed partial recovery in regions we applied adjustments to, but these actions were not resulting in sustained recovery.

5:00 AM – We encountered persistent challenges with machines being stuck in bad states, preventing our changes from syncing as expected. We scaled up virtual machines to increase capacity and reduce service load.
7:20 AM – We identified a mitigation and applied the mitigation - an update that reverted the defective change to a portion of the capacity
8:00 AM - We confirmed server stability across the portions of infrastructure to which we'd applied the update and initiated its expansion systematically to focus on infrastructure supporting regions that were currently within or entering peak business hours.
10:00 AM – The service began to stabilize, though capacity had not returned to normal levels.
11:00 AM - Recovery actions were completed across portions of the affected infrastructure in the Asia Pacific region. Traffic continued to be gradually restored, and we observed improvements in key customer scenarios including mailbox access and email processing.
11:35 AM – Availability in Europe returned to healthy levels, indicating that services in the region were stabilizing.
1:00 PM – Recovery progress was sustained across all regions and queues of delayed emails messages began to reduce
2:17 PM – Global webmail availability returned to expected levels, and by 2:44 PM the improvement had remained stable under continued monitoring.
3:05 PM – We completed the move of eligible customer traffic to alternative connection paths. This further improved sign-in and service access across affected regions.
3:54 PM – Most major access methods had recovered to their normal traffic levels, however monitoring for full recovery was still in progress.
4:10 PM – We confirmed that customers were no longer experiencing active impact from the incident and moved into an extended monitoring and backlog recovery phase.
>4.10 PM – We continued to monitor service for sustained recovery, ensuring that all queued mail had been processed, and as load continued to increase, we confirmed the service remained healthy

Saturday, November 14
4:15 AM – After extended monitoring, which confirmed service stability, we issued our final customer updates to EX1550543 and MO1550821.

Mitigation / Repair Items

The repair items below focus on the actions that restored service and have been applied as protections since the incident. Many longer-term changes from our full root cause analysis will be in the final Post Incident Report later this week.

Action	Status
A comprehensive list of repair items will be delivered as part of the final PIR	11/20/2026
Developed an update that returned the connection handling code path to the previously known good state	Completed
Extended the email delivery window to ensure delayed messages weren't returned as undeliverable and all messages that were queued for delivery during the event were processed.	Completed
Added additional capacity as a proactive measure to help with recovery.	Completed
Additional monitoring for traffic-redirection patterns and connection-handling health.	Completed
Implemented guardrails for traffic distribution to avoid cascading impact.	Completed