Post Incident Report

For

Microsoft 365

Report Date: November 20, 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

Incident ID	**EX1550543/MO1550821**
Incident Title	EX1550543 – Users may have been unable to connect to their mailboxes in Exchange Online through multiple connection methods. Users are unable to send or receive email. MO1550821 – Users may be unable to access or perform several actions in Microsoft 365 services.
Service(s) Impacted	Microsoft 365 services, including Exchange Online, Microsoft 365 Copilot, Microsoft Teams, SharePoint Online, OneDrive for Business, Microsoft Defender, Microsoft Purview, Microsoft Fabric, Microsoft Intune, Microsoft Planner, Power Apps, Microsoft To Do, Power BI, and Universal Print.

Executive Summary

On November 12, 2026, a latent Windows defect caused widespread failures in Microsoft 365 network gateway servers used by Exchange Online, disrupting Exchange Online and dependent features and services. The defect was specific to a configuration used by Exchange Online and was not present in the standard Windows deployment. Under increased load, the defect caused affected servers to become unresponsive.

The defect was triggered by the combination of three operational conditions: (1) reduced production capacity due to planned maintenance, (2) a shift of Teams traffic to the Exchange Online gateways, and (3) peak North American traffic. Normally, these conditions in combination are within normal operating parameters and would not have caused a disruption without the underlying defect. However, with the presence of the latent defect, their combined effect exposed the issue which caused servers to become unresponsive. At this point, when capacity became unavailable, the service automatically redistributed traffic to remaining healthy servers across the worldwide fleet, expanding and accelerating the impact.

Recovery was extended by three primary factors: (1) identifying the underlying defect, which had been introduced months earlier, (2) the need to restart affected servers before applying the fix and restart them again to activate the update, and (3) the time required to deploy the update across the large number of affected servers.

No mail was lost, and all queued messages were successfully delivered. The incident did not impact customer data at rest.

This incident was confirmed not to be caused by a cyberattack, and we found no evidence of malicious activity. This change was not caused by AI tooling.

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 most 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.

For impacted users, below is the detailed information on scenarios that were likely disrupted:

Exchange Online:

- Users were unable to connect to, or experienced significantly reduced access to, their mailboxes through Outlook, Outlook on the web, Outlook for iOS and Android, Exchange ActiveSync, and Messaging and REST APIs.
- Search experience for Exchange Online in Outlook, Outlook on the web, Outlook for iOS and Android and other applications was degraded.
- Business applications and integrations that connect to Exchange Online experienced failures.
- Sending or receiving email messages was severely degraded. To reduce the chance that delayed messages would be returned to senders, we extended the period during which Exchange Online keeps retrying delivery to 48 hours before it generates a non-delivery report. All queued enterprise messages have since been delivered, and we have found no enterprise messages that were permanently lost.
- Some delayed messages destined for consumer mailboxes that had been inactive for more than 30 days expired, and senders received non-delivery reports.

Microsoft Teams:

- Teams services moved their traffic off the affected networking gateway and were partially restored earlier than Exchange Online. Features that depend on Exchange Online data, such as search and people-related features, may have continued to be affected.
- Users may have experienced issues with calendar functionality.
- Users may have experienced issues with search and users' presence may have been stale and not updated automatically or through manual adjustment. User location adjustment may have failed.
- Users may have experienced issues with Microsoft Teams Room (MTR) devices and desktop phones displayed offline in the Teams admin center.
- Users may have experienced issues creating, deleting, and editing Teams channels.

Microsoft 365 Copilot:

- Most users were unable to launch the M365 Copilot desktop apps or load the web apps. Most prompts would fail or time out. Prompts sent to Copilot Cowork may have responded, but would have failed to access mail, file or calendar data in Exchange Online.

Microsoft Graph:

- Users may have experienced failures for scenarios that rely on access to Exchange Online data.

Microsoft 365 Admin Center:

- Admins may have experienced issues accessing the admin center.

Microsoft Purview:

- Users may have been unable to access the Microsoft Purview portal.
- Policy changes may have been delayed.
- Policy matches, alerts, and processing activities may also have been delayed.
- Users may have encountered issues opening or creating protected content.
- Impacted solutions included Data Loss Prevention, Information Protection, Communication Compliance, Data Lifecycle Management, eDiscovery, Insider Risk Management, Audit searches, Data Classification, and On-Demand Scans.

Microsoft Defender XDR:

- Users may have experienced intermittent failures across solutions in Defender.

SharePoint Online and OneDrive for Business:

- Users experienced failures or slowness early in the incident. These services moved their traffic off the affected networking gateway and were partially restored earlier than Exchange Online.
- Features that depend on Exchange Online data, such as search and people-related features, may have continued to be affected.

Microsoft Defender for Office:

- Security Admins/SOC may have had trouble releasing or deleting quarantined messages, managing tenant allow/block lists (TABL), submitting items, or taking remediation actions on email. If the operation didn't complete, please try it again after service was restored.
- Users may have had trouble opening links protected by Safe Links. Depending on your organization's policy, users may also have seen a "Skip verification" option when Safe Links checked a URL at the time it was clicked.

- Email post-delivery protection (ZAP) was delayed because updated threat intelligence wasn't available. As a result, ZAP may not have acted on affected messages as quickly as usual.

Microsoft Fabric, Microsoft Intune, Microsoft Planner, Power Apps, Microsoft To Do, Power BI:

- Features that depend on Exchange Online data or connectivity may have failed or been delayed.

Universal Print:

- Users may have been unable to create, upload, or retrieve print jobs.

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 11:00 AM UTC – Recovery actions were completed across portions of the affected infrastructure in the Asia Pacific region.
- Friday, November 13, at 11:35 AM UTC – Availability in Europe returned to healthy levels.
- Friday, November 13, at 2:17 PM UTC – Global Outlook on the web availability returned to expected levels.
- Friday, November 13, at 3:54 PM UTC – Most major access methods had recovered to normal traffic levels.
- Friday, November 13, at 4:10 PM UTC – Customers were no longer experiencing active impact, and we moved into extended monitoring and mail-backlog recovery.
- Saturday, November 14, at 4:15 AM UTC – After extended monitoring confirmed service stability, the incident was declared resolved.

Root Cause and Mitigation Strategy

Detailed root cause and mitigation details

The incident on Thursday, November 12th, 2026, was caused by four main factors that together caused global impact for the networking gateway that fronts many Microsoft 365 services but primarily impacting Exchange Online.

Factor 1: Connection-processing defect
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 an update that was enabled on Microsoft 365 networking gateway servers; neither the underlying defect nor the update applies to broader Windows deployments.
Moreover, this impacted our newer SKU hardware, but not the older generation SKUs in the service. The newer SKUs have more cores, which allows for more concurrent requests which led to higher contention which triggered this latent defect.

The change that introduced the defect 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 issue because the traffic conditions were different from what was observed during the incident and SKUs used in our labs reduced the likelihood of encountering this defect. As noted previously, the contention issue is worse on machines with more cores and newer hardware; whereas our load testing used older SKU hardware with fewer cores.

Factor 2: Reduced server capacity
Nearly 14% of server capacity for the networking gateway was out of service as part of routine operations when we deploy the latest version of the service on these servers. The service keeps enough spare capacity that this reduction in capacity alone would not have caused impact.

Factor 3: Teams failover
Within Microsoft 365, we maintain two independent and heterogeneous network gateway systems designed for isolation and high availability. This allows many Microsoft 365 services to fail to/from them during outages or disruptions as part of the standard operating procedure for business continuity and disaster recovery (BCDR) and normally creates no impact.

At 1:07 AM UTC on November 12, there was a shift of Teams traffic due to a routing issue on its networking gateway in the US East region which caused Teams connection timeouts and triggered an automatic failover to the redundant network gateway that also supports Exchange Online.

These types of traffic shifts and failovers are routine and have not historically impacted Exchange Online. On November 12, after Teams traffic shifted, all services continued to remain healthy for the next ~16 hours. The gateway servers did not become unresponsive

immediately since these gateway servers were not yet serving daily peak traffic for those regions.

Factor 4: Peak Load Conditions

The first failures in North America occurred near the start of the business day on the US East Coast. Broad failures in North America began at approximately 4:00 PM UTC on November 12, as US traffic approached its daily peak.

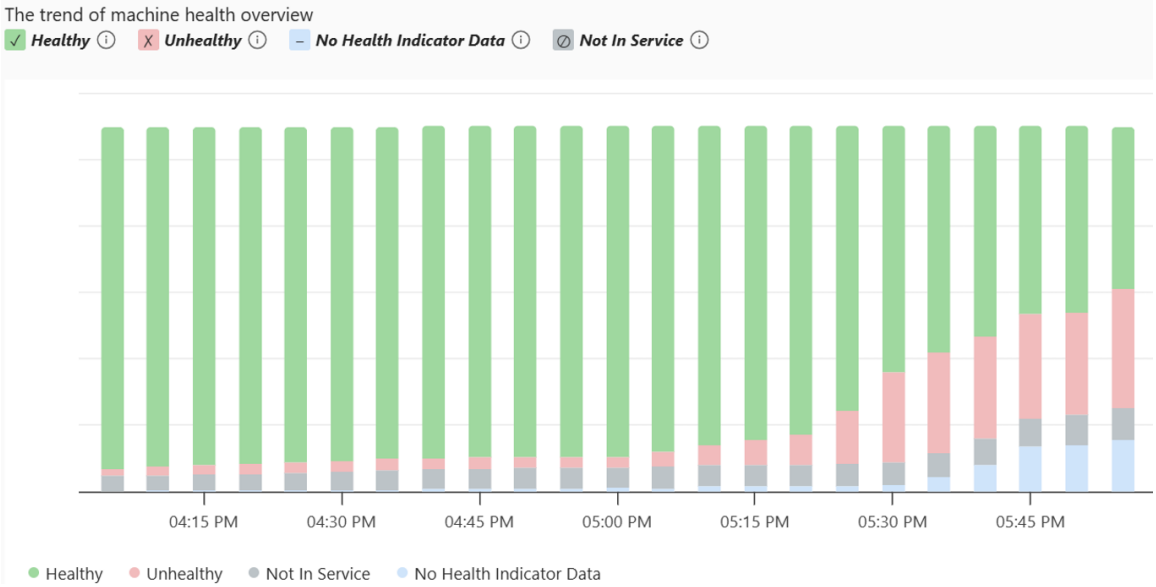


Figure 1: This chart shows the network gateway servers in North America becoming either unhealthy or unresponsive (shown as "no health indicator data") as the daily peak started (all times in UTC)

Taken together, the peak traffic, shift in Teams traffic 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.

How the disruption spread

The network gateway is composed of thousands of 'capacity units' which represent shards of capacity, each shard is composed of compute servers and networking infrastructure. When gateway servers become unresponsive, our traffic management systems automatically route traffic to other healthy 'capacity units' in the same region. This happens often and is critical to avoiding any interruption due to hardware or software defects and is normally what makes the service resilient.

In North America, the Windows performance regression, increased load from redirected Teams traffic, and reduced server capacity in some capacity units combined to overwhelm the remaining servers during peak traffic, causing those capacity units to become unhealthy. As traffic was redirected away from the affected capacity units, the remaining capacity units absorbed additional load. This increased the likelihood of triggering the same

Windows performance regression on other gateway servers and making them unhealthy as well.

The result was a cascading failure across North America: as servers became unresponsive, the same traffic load was distributed across fewer healthy servers, placing additional pressure on them and causing the issue to spread. In less than two hours, this previously unknown defect was triggered across a significant portion of the North American server fleet. At the lowest point, more than 80% of servers were unresponsive.

When a region starts to exhaust its healthy capacity, the traffic management system detects this condition and redirects the traffic to other regions. In this case, this redistribution further spread the impact to other regions (which had the same performance regression), across our network gateway infrastructure globally. Within about two hours of the broad failures that began in North America, servers in every region were affected globally.

Mitigation of incident

The first step in recovery was to quickly move Teams, OneDrive, SharePoint, and Office apps traffic to the redundant network gateway alleviating load and reducing impact to those services. At this point recovery efforts were focused on identifying root cause and determining any method to restore service health regardless of root cause.

Why recovery took longer

Recovery took longer than usual because as gateway servers returned to service, they were immediately overwhelmed by queuing client traffic and became unresponsive again. Our automated recovery processes were not equipped to restore service reliably under this combination of widespread gateway server failures and sustained reconnection demand. Because the defect came from a change deployed several months earlier, our change auditing and diagnostics processes took longer to identify the Windows change as the root cause. Once we understood the root cause, our usual safeguards for recent changes (stopping a config rollout, rolling back a deployment, or switching off a newly enabled feature) could not quickly revert it. As noted earlier, given the defect was already deployed in the worldwide fleet, we had to run recovery for each capacity unit across the globe.

Applying the fix on a gateway server required five steps: (1) the server had to be taken out of service so it would not take load when restarted. (2) the server had to be restarted from its failed, unresponsive state, (3) the patch was then applied, (4) the server was restarted a second time for the patch to be active and (5) the server was put back into service.

Once the server recovery approach was finalized, we created a recovery strategy to apply it globally at scale in a timely manner. We moved large groups of ‘capacity units’ out of traffic rotation, recovered all the unresponsive servers in those capacity units concurrently, and

then returned the group of ‘capacity units’ to service. We then expanded this recovery approach across all regions.

To give delayed email more time to be delivered, we extended the delivery retry window from 24 to 48 hours before a non-delivery report is generated.

Actions Taken (All times UTC)

Time (UTC)	Event
Nov 12, 12:25 AM	Teams users in the US East region began experiencing timeouts when connecting through the alternative networking gateway.
Nov 12, 1:07 AM	The automated failover system moved Teams traffic onto the network gateway front end servers shared with Exchange Online.
Nov 12, 1:07 AM – 4:00 PM	Load on the front-end servers increased gradually. Isolated server failures began in some regions during this period, after which automated mitigations applied.
Nov 12, 4:03 PM	Our connectivity anomaly detection monitor identified an increase in end-user connectivity issues.
Nov 12, 4:17 PM	Due to the increase in connectivity issues observed, automation began to redirect traffic to alternate healthy endpoints.
Nov 12, 4:20 PM	On-call engineers received a high-priority alert indicating impact to Exchange Online services, which triggered immediate communication to the Service Health Dashboard under EX1550543.
Nov 12, 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.
Nov 12, 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. User reports began to accumulate on 3rd party outage reporting sites such as DownDetector.
Nov 12, 5:03 PM	Our connectivity anomaly detection monitors began to see a further increase in connectivity issues.

Time (UTC)	Event
Nov 12, 5:32 PM	Automated recovery actions continued to redirect traffic to healthy endpoints; however, further service degradation was detected across multiple regions.
Nov 12, 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.
Nov 12, 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.
Nov 12, 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.
Nov 12, 6:33 PM	We triggered failovers to ensure traffic was distributed evenly across our infrastructure and components that manage user connections.
Nov 12, 7:00 PM	We expanded our response to accelerate mitigation efforts. As a precaution, we paused all ongoing and planned deployments and configuration changes.
Nov 12, 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.
Nov 12, 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.
Nov 12, 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.

Time (UTC)	Event
Nov 12, 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.
Nov 12, 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.
Nov 12, 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.
Nov 12, 11:55 PM	We determined that the servers we'd returned to service initially remained healthy but then became unresponsive again upon receiving traffic.
Nov 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.
Nov 13, 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.
Nov 13, 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.

Time (UTC)	Event
Nov 13, 1:53 AM	We disabled some automated traffic redirection mechanisms performed by our connectivity anomaly detection monitors.
Nov 13, 2:10 AM	We applied infrastructure capacity improvements where possible in an effort to help support traffic in regions coming online.
Nov 13, 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.
Nov 13, 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.
Nov 13, 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.
Nov 13, 7:20 AM	We identified a mitigation and applied the mitigation: an update that reverted the defective change to a portion of the capacity.
Nov 13, 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.
Nov 13, 10:00 AM	The service began to stabilize, though capacity had not returned to normal levels.
Nov 13, 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.
Nov 13, 11:35 AM	Availability in Europe returned to healthy levels, indicating that services in the region were stabilizing.

Time (UTC)	Event
Nov 13, 1:00 PM	Recovery progress was sustained across all regions and queues of delayed email messages began to reduce.
Nov 13, 2:17 PM	Global webmail availability returned to expected levels, and by 2:44 PM the improvement had remained stable under continued monitoring.
Nov 13, 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.
Nov 13, 3:54 PM	Most major access methods had recovered to their normal traffic levels; however, monitoring for full recovery was still in progress.
Nov 13, 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.
Nov 13, 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.
Nov 14, 4:15 AM	After extended monitoring, which confirmed service stability, we issued our final customer updates to EX1550543 and MO1550821.

Completed actions

Action	Status
Developed an update that reverted the connection handling change that introduced the performance regression	Completed
We permanently updated the Windows image to include the fix which will get applied on all new and existing server deployments	Completed
We extended the email delivery window to ensure delayed messages weren't returned as undeliverable before the incident was mitigated	Completed

Action	Status
Developed an update that reverted the connection handling change that introduced the performance regression	Completed
We permanently updated the Windows image to include the fix which will get applied on all new and existing server deployments	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 temporary guardrails to avoid automatic traffic distribution and avoid cascading impact.	Completed
Return servers from maintenance	Completed

Next Steps

The actions described below consolidate engineering efforts to restore the environment, reduce issues in the future, and enhance resiliency of Microsoft 365 services.

Findings	Action	Completion Date
Traffic distribution design can cascade failures to other regions.	[Blast radius] Fix traffic distribution to eliminate such cascaded failures to other regions.	CY26 Q4
Regional failure zones are too large	[Blast radius] Design traffic distribution for smaller fault zones	TBD
Teams failover contributed to exposing the defect impacting mail flow and Outlook connectivity	[Blast radius] Partition the network gateway capacity into pools that serve different workloads to isolate failures	CY27Q1
A routing issue in one region moved all Teams traffic worldwide to the primary network gateway.	[Failover design] Teams traffic fails over only for the affected region and not worldwide	CY26Q4
Failover testing did not identify the defect in isolation	[Failover drills, Blast radius] Scoped failover testing to cover peak conditions and other changes that rollout as part of safe deployment	TBD

Change auditing and diagnostics processes took longer to identify the Windows change as the root cause	[Diagnosis] Update diagnostics collection and change auditing for Windows changes	CY27Q1
Distributing and validating the corrective configuration across the global fleet took too long for timely recovery	[Recovery] Improve recovery procedures for fleet-wide failures	CY27Q1
Servers were re-overwhelmed on return to service.	[Recovery] Add gradual traffic ramp-up and admission control when servers rejoin. Build controls to reduce traffic by protocol, workload, or region during recovery. Improve and implement additional retry and back-off behavior within the service.	TBD
Service retries may spread the issues to other capacity units	[Blast radius] Improved implementation of retry and back-off behavior with other services.	TBD
Our testing did not adequately represent the combination of hardware and traffic mix encountered during this incident.	[Testing] We're expanding load testing to our newest, highest-core hardware and with production-representative traffic, including running regular tests at peak-load limits.	CY27Q1
The network gateways did not successfully manage a defect in combination with other factors creating impact at peak traffic conditions	[Capacity Design] Add additional capacity for more headroom	CY27Q1
Some capacity units had reduced server capacity (part of routine maintenance), can be risky in combination with other factors	[Capacity Design] Detection for extended capacity loss that risk running out of headroom	CY26Q4
First customer notification came about 20 minutes after broad impact	[Customer communications] Improve detection and automated notification to reduce time to communicate	CY26Q4