

Title:	PIR – Multiple services impacted after a power/cooling event – Australia East
Tracking ID:	VVTQ-J98
Event Type:	Service Issue
Status:	Resolved
Service(s):	Azure Relay, Azure Purview, Azure Health Data Services, Virtual Machines, Azure Container Apps, Backup, IoT Central, Azure Cosmos DB, Azure Data Explorer, Event Hubs, Azure Arc enabled Kubernetes, Azure Database for MySQL flexible servers, ExpressRoute \ ExpressRoute Circuits, Stream Analytics, Azure Chaos Studio, API Management, Log Analytics, Application Insights, ExpressRoute \ ExpressRoute Gateways, Azure NetApp Files, Azure Databricks, SQL Database, Service Fabric, Azure IoT Hub, Azure Digital Twins, Notification Hubs, Azure Kubernetes Service (AKS), Logic Apps, Data Factory, Network Infrastructure, Service Bus, App Service, Azure Search, Azure API for FHIR, Batch, Storage, Container Registry, Azure Database for PostgreSQL flexible servers, HDInsight, Redis Cache, Device Update for IoT Hub
Region(s):	Australia East, Australia Southeast
Start time:	2023-08-30T09:00:00.373Z
Resolve time:	2023-08-30T15:10:00.000Z
Last update time:	2023-09-14T21:00:25.017Z
Impacted subscriptions:	fe4a62bd-9f59-4d84-903b-0ff05610f075, 46995d20-05a1-49c1-a8b3-83753d6ec980

Last update:

Post Incident Review (PIR) – Multiple services impacted after a power/cooling event – Australia East

This incident will be discussed in an upcoming Azure Incident Retrospective customer livestream webcast event.

Join either session to hear from our engineering leaders - on what happened, how we responded, and what we learned.

The event includes a live Q&A experience, to ask our engineering experts any questions you have about the incident.

Session option #1 - 20 September 2023 @ 01:00 UTC - Register at: <https://www.aka.ms/air6/reg/ash/1of2>

Session option #2 - 20 September 2023 @ 15:00 UTC - Register at: <https://www.aka.ms/air6/reg/ash/2of2>

What happened?

Starting at approximately 10:30 UTC on 30 August 2023, customers may have experienced issues accessing or using Azure, Microsoft 365 and Power Platform services. This incident was triggered by a utility power sag at 08:41 UTC on 30 August 2023, which impacted one of the three Availability Zones of the Australia East region. This power sag tripped a subset of the cooling system chiller units offline and, while working to restore cooling, temperatures in the datacenter increased to

levels above operational thresholds. We powered down a small subset of selected compute and storage scale units, both to lower temperatures and to prevent damage to hardware. Although the vast majority of services recovered by 22:40 UTC on 30 August 2023, full mitigation was not until 20:00 UTC on 3 September 2023 – as some services experienced prolonged impact, predominantly as a result of dependencies on recovering subsets of Storage, SQL Database, and/or Cosmos DB services.

Multiple Azure services were impacted by this incident – including Azure Active Directory (AAD), Azure Active Directory B2C, Azure Active Directory Conditional Access, Azure Active Directory Connect Health, Azure Active Directory MyApps, Azure Activity Logs & Alerts, Azure API Management, Azure App Service, Azure Application Insights, Azure Arc enabled Kubernetes, Azure API for FHIR, Azure Backup, Azure Batch, Azure Chaos Studio, Azure Container Apps, Azure Container Registry, Azure Cosmos DB, Azure Databricks, Azure Data Explorer, Azure Data Factory, Azure Database for MySQL flexible servers, Azure Database for PostgreSQL flexible servers, Azure Digital Twins, Azure Device Update for IoT Hub, Azure Event Hubs, Azure ExpressRoute, Azure Health Data Services, Azure HDInsight, Azure IoT Central, Azure IoT Hub, Azure Kubernetes Service (AKS), Azure Logic Apps, Azure Log Analytics, Azure Log Search Alerts, Azure NetApp Files, Azure Notification Hubs, Azure Redis Cache, Azure Relay, Azure Resource Manager (ARM), Azure Role Based Access Control (RBAC), Azure Search, Azure Service Bus, Azure Service Fabric, Azure SQL Database, Azure Storage, Azure Stream Analytics, Azure Virtual Machines, Microsoft Purview, and Microsoft Sentinel.

What went wrong and why?

Starting at approximately 08:41 UTC on 30 August 2023, a utility voltage sag was caused by a lightning strike on electrical infrastructure approximately 18 miles from the impacted Availability Zone of the Australia East region. The voltage sag caused cooling system chillers for multiple datacenters to shut down. While some chillers automatically restarted, 13 failed to restart and required manual intervention. To do so, the onsite team accessed the datacenter rooftop facilities, where the chillers are located, and proceeded to sequentially restart chillers moving from one datacenter to the next. By the time the team reached the final five chillers requiring a manual restart, the water inside the pump system for these chillers (chilled water loop) had reached temperatures that were too high to allow them to be restarted. In this scenario, the restart is inhibited by a self-protection mechanism that acts to prevent damage to the chiller that would occur by processing water at the elevated temperatures. The five chillers that could not be restarted supported cooling for the two adjacent data halls which were impacted in this incident.

The two impacted data halls require at least four chillers to be operational. Before the voltage sag, our cooling capacity for these halls consisted of seven chillers, with five chillers in operation and two chillers in standby. At 10:30 UTC some networking, compute, and storage infrastructure began to shutdown automatically as data hall temperatures increased, impacting service availability. At 11:34 UTC, as temperatures continued to increase, our onsite datacenter team began a remote shutdown of remaining networking, compute, and storage infrastructure in the impacted data halls to protect data durability, infrastructure health, and address the thermal runaway. This shutdown allowed the chilled water loop to return to a safe temperature which allowed us to restart the chillers. This shutdown of infrastructure resulted in a further reduction of service availability for this Availability Zone.

The chillers were successfully brought back online at 12:12 UTC, and data hall temperatures returned to operational thresholds by 13:30 UTC. Power was then restored to the affected infrastructure and a phased process to bring the infrastructure back online commenced. All power to infrastructure was restored by 15:10 UTC. Once all networking and storage infrastructure had power restored, dependent compute scale units were then also returned to operation. As the underlying compute and storage scale units came online, dependent Azure services started to recover, but some services experienced issues coming back online.

From a storage perspective, seven storage scale units were impacted – five standard storage scale units, and two premium storage scale units. Availability impact to affected storage accounts began at 10:30 UTC as hardware shut down in response to elevated data hall temperatures. This was most impactful to storage accounts configured with the default local redundant storage (LRS), which is not resilient to a zonal failure. Accounts configured as zonally redundant (ZRS) remained

100% available, and accounts configured as geographically redundant (GRS) were eligible for customer-managed account failover. After power restoration, storage nodes started coming back online from 15:25 UTC. Four scale units required engineer intervention to check and reset some fault detection logic – the combination of investigation and management tooling performance problems delayed restoration of availability for these scale units. We also identified that some automation was incorrectly marking some already recovered nodes as unhealthy, which slowed storage recovery efforts. By 20:00 UTC, 99% of storage accounts had recovered. Restoring availability for the remaining <1% of storage accounts took more time due to hardware troubleshooting and replacement required on a small number of storage nodes in a single scale unit. Even identifying problematic hardware in the storage nodes took an extended period of time, as the nodes were offline and therefore not able to provide diagnostics. By 01:00 UTC on 31 August, availability was restored for all except a handful of storage accounts, with complete availability restored by 07:00 UTC on 1 September.

From a SQL Database perspective, database capacity in a region is divided into tenant rings. The Australia East region includes hundreds of rings, and each ring consists of a group of VMs (10-200) hosting a set of databases. Rings are managed by Azure Service Fabric to provide availability in cases of VM, network or storage failures. When infrastructure was powered down, customers using zone-redundant Azure SQL Database did not experience any downtime, except for a small subset of customers using proxy mode connection, due to one connectivity gateway not being configured with zonal-resilience. The fix for this issue was already being rolled out, but had not yet deployed to the Australia East region. As infrastructure was powered back on, all tenant rings except one came back online and databases became available to customers as expected. However, one ring remained impacted even after Azure Compute became available. In this ring, 20 nodes did not come back online as expected, so databases on these nodes continued to experience unavailability. As a result of Service Fabric attempting to move databases to healthy nodes, other databases on this ring experienced intermittent availability issues as a side-effect of the overall replica density and unhealthy nodes. The recovery involved first moving all the databases from unhealthy nodes to healthy nodes. All remote storage (general purpose) databases were successfully recovered by this move, but databases using local storage (business critical) only recovered as their underlying nodes recovered. All databases on unhealthy nodes were recovered by 11:00 UTC on 31 August. Since the health and capacity of this final ring did not completely recover, we decided to move all databases out of the ring, which extended the overall recovery time but did not negatively impact customer availability. During this extended recovery, most customers were not experiencing any issues but it was important to move all databases out of this unhealthy ring to prevent any potential impact. The operation of moving all databases out of this ring was completed at 20:00 UTC on 3 September. During this incident, customers who had 'active geo-replication' setup were able to failover manually to restore availability. For customers who have 'auto-failover groups' enabled, we did not execute automatic failover – our automatic failover policy was not initiated for the region, due to an incorrect initial assessment of the impact severity to SQL Database.

From a Cosmos DB perspective, zone resilient accounts and those with multi-region writes remained operational during the incident, transparently serving requests from a different zone or region, respectively. However, accounts not configured for AZ or multi-region writes experienced full or partial loss of availability, due to the infrastructure that was powered down. Multi-region accounts with single region write eligible for failover (those with Service Managed Failover enabled) were failed over to their alternate regions to restore availability. These were initiated at 12:07 UTC, 33 minutes after decision to power down scale units. The reason for this delay was to identify and failover the Cosmos DB control plane system resources – in retrospect this delay was unnecessary, as the Cosmos DB control plane was already fully zone-resilient. 95% of database accounts were failed over within 35 minutes by 12:42 UTC, and all eligible accounts were failed over by 16:13 UTC on 30 August. Accounts that were not eligible for failover had service restored to partitions only once the dependent storage and compute were restored.

From an Azure Kubernetes Service (AKS) perspective, the service experienced a loss of compute for the AKS control plane for Australia East as well as data access loss to SQL Database. The AKS control plane underlay is deployed across multiple availability zones. AKS uses Azure SQL Database for its operation queue which is used for Create/Read/Update/Delete (CRUD) activities. Although scheduled to be converted, the SQL Database in Australia East was not configured with AZ resiliency selection, leaving it unavailable during the incident period. In addition, AKS services in the Australia Southeast region depended on this same database, causing an AKS incident for CRUD activities in that region also. Existing customer workloads running on AKS clusters in either region should not have been impacted by the downtime, as long as they did not need to access the AKS resource provider for scaling or other CRUD activities. As the SQL Database recovered, service was restored without any other mitigation required.

From an Azure Resource Manager (ARM) perspective, the impact on customers was the result of degradation in Cosmos DB. This degradation impacted ARM between 10:45 UTC and 12:25 UTC and resulted in ARM platform availability for the Australia East region dropping from ~99.999% to (at its lowest) 88%, with a 62% success rate for write operations. For data consistency reasons, write operations are required to be sent to the associated Cosmos DB regional replica for a given resource group. While the migration to our next generation zonally-redundant storage architecture is still ongoing, it has not been completed and as a result this region is not yet leveraging fully zonally redundant storage for ARM. This meant that for the duration of the incident, customers worldwide trying to manage resources whose resource groups were homed in Australia East saw increased error rates (and this manifested in a small impact to global platform availability until 15:00 UTC).

How did we respond?

30 August 2023 @ 08:41 UTC – Voltage sag occurred on utility power line
30 August 2023 @ 08:43 UTC – 13 chillers failed to restart automatically
30 August 2023 @ 08:51 UTC – Remote resets on chillers commenced
30 August 2023 @ 09:09 UTC – Team arrived at first group of chillers for manual restarts
30 August 2023 @ 09:18 UTC – Team arrived at second group of chillers for manual restarts
30 August 2023 @ 09:42 UTC – Team arrived at third group of chillers for manual restarts
30 August 2023 @ 09:45 UTC – Team arrived at the final group of chillers which could not be restarted
30 August 2023 @ 10:30 UTC – Initial impact from automated infrastructure shutdown
30 August 2023 @ 10:47 UTC – Cosmos DB Initial impact detected via monitoring
30 August 2023 @ 10:48 UTC – First automated communications sent to Azure Service Health
30 August 2023 @ 11:30 UTC – Initial communications posted to public Azure Status page
30 August 2023 @ 11:36 UTC – All subscriptions in Australia East sent portal communications
30 August 2023 @ 11:34 UTC – Decision made to shutdown impacted infrastructure
30 August 2023 @ 12:07 UTC – Failover initiated for eligible Cosmos DB accounts
30 August 2023 @ 12:12 UTC – Five chillers manually restarted
30 August 2023 @ 13:30 UTC – Data hall temperature normalized
30 August 2023 @ 14:10 UTC – Safety walkthrough completed for both data halls
30 August 2023 @ 14:25 UTC – Decision made to start powering up hardware in the two affected data halls
30 August 2023 @ 15:10 UTC – Power restored to all hardware
30 August 2023 @ 15:25 UTC – Storage, networking, and compute infrastructure started coming back online after power restoration
30 August 2023 @ 15:30 UTC – Identified three specific storage scale units still experiencing fault codes
30 August 2023 @ 16:00 UTC – 35% of VMs recovered / Began manual recovery efforts for three remaining storage scale units
30 August 2023 @ 16:13 UTC – Account failover completed for all Cosmos DB accounts
30 August 2023 @ 17:00 UTC – All but one SQL Database tenant ring had recovered
30 August 2023 @ 19:00 UTC – 99% of storage accounts were back online
30 August 2023 @ 19:20 UTC – 90% of VMs recovered
30 August 2023 @ 19:29 UTC – Successfully recovered all premium storage scale units
30 August 2023 @ 22:35 UTC – Standard storage scale units were recovered, except for one scale unit
30 August 2023 @ 22:40 UTC – 99% of VMs recovered
31 August 2023 @ 05:00 UTC – 100% of VMs recovered
31 August 2023 @ 04:04 UTC – Restoration of Cosmos DB accounts to Australia East initiated
31 August 2023 @ 04:43 UTC – Final Cosmos DB cluster recovered, restoring all traffic for accounts that were not failed over to alternate regions
1 September 2023 @ 06:40 UTC – Successfully recovered all standard storage scale units
3 September 2023 @ 20:00 UTC – Final SQL Database tenant ring evacuated and all customer databases online

How are we making incidents like this less likely or less impactful?

Incremental load increases over time in the Availability Zone resulted in a chiller operating configuration that was susceptible to a timing defect in the Chiller Management System. We have de-risked this failure to restart due to voltage fluctuations, by implementing a change to the control timing logic on the Chiller Management System (Completed). Emergency Operation Procedure for manual restarts of simultaneous chiller failures has changed from an 'adjacent' sequence to a data hall 'load based' sequence to ensure all impacted data halls have partial cooling, to slow thermal runaway while full cooling is being restored (Completed).

Following this incident, as a temporary mitigation, we increased technician staffing levels at the datacenter to be prepared to execute manual restart procedures of our chillers prior to the change to the Chiller Management System to prevent restart failures. Based on our incident analysis the staffing levels at the time would have been sufficient to prevent impact if a 'load based' chiller restart sequence had been followed, which we have since implemented (Completed).

Datacenter staffing levels published in the Preliminary PIR only accounted for "critical environment" staff onsite. This did not characterize our total datacenter staffing levels accurately. To alleviate this misconception, we made a change to the preliminary public PIR posted on the Status History page.

Our Storage team has identified several optimizations in our large scale recovery process which will help to reduce time to mitigate. This includes augmenting data provided in our incidents to enable quicker decision making, and updates to our troubleshooting guides (TSGs) that enable faster execution (Estimated completion: December 2023).

Our Azure Service Fabric team is working to improve reliability of SQL Database tenant ring recovery. (Estimated completion: December 2023).

Our SQL Database team is reviewing our 'auto-failover group' trigger criteria, to ensure that failovers can happen within the expected timeframe. (Estimated completion: October 2023).

Our SQL Database team is upgrading internal tooling to enable mass migration of databases. (Estimated completion: December 2023).

Our Cosmos DB team is working to optimize Service Managed Failover for single region write accounts to reduce time to mitigate (Estimated completion: November 2023).

Our AKS team is immediately converting all operation queue SQL Database databases to be zone redundant (Estimated completion: September 2023).

Our AKS team is also replacing all cross-region SQL Database queue usage with Service Bus queues that are zone redundant (Estimated completion: September 2023).

Our ARM team will complete its storage layer migration to the next generation, zonally redundant architecture (Estimated completion: December 2023).

Our incident management team is exploring ways to harden our readiness, process, and playbook surrounding power down scenarios (Estimated completion: October 2023)

How can customers make incidents like this less impactful?

Consider using Availability Zones (AZs) to run your services across physically separate locations within an Azure region. To help services be more resilient to datacenter-level failures like this one, each AZ provides independent power, networking, and cooling. Many Azure services support zonal, zone-redundant, and/or always-available configurations: <https://docs.microsoft.com/azure/availability-zones/az-overview>

For mission-critical workloads, customers should consider a multi-region geodiversity strategy to avoid impact from incidents like this one that impacted a single region: <https://learn.microsoft.com/training/modules/design-a-geographically-distributed-application/> and <https://learn.microsoft.com/azure/architecture/patterns/geodes>

Consider which are the right storage redundancy options for your critical applications. Zone redundant storage (ZRS/GZRS) remains available throughout a zone localized failure, like in this incident. Geo-redundant storage (GRS) enables account level failover in case the primary region endpoint becomes unavailable: <https://learn.microsoft.com/azure/storage/common/storage-redundancy>

Consider the relevant guidance for recovering your SQL Database databases during disaster recovery scenarios:

<https://learn.microsoft.com/azure/azure-sql/database/disaster-recovery-guidance>

Consider the relevant guidance for achieving high availability with Azure Cosmos DB: <https://learn.microsoft.com/azure/cosmos-db/high-availability>

More generally, consider evaluating the reliability of your applications using guidance from the Azure Well-Architected Framework and its interactive Well-Architected Review: <https://docs.microsoft.com/azure/architecture/framework/re-siliency>

Finally, consider ensuring that the right people in your organization will be notified about any future service issues – by configuring Azure Service Health alerts. These can trigger emails, SMS, push notifications, webhooks, and more: <https://aka.ms/ash-alerts>

How can we make our incident communications more useful?

You can rate this PIR and provide any feedback using our quick 3-question survey: <https://aka.ms/AzPIR/VVTQ-J98>

Update history:

Fri Sep 15 2023 07:00:25 GMT+1000 (Australian Eastern Standard Time)

Post Incident Review (PIR) – Multiple services impacted after a power/cooling event – Australia East

This incident will be discussed in an upcoming Azure Incident Retrospective customer livestream webcast event.

Join either session to hear from our engineering leaders - on what happened, how we responded, and what we learned.

The event includes a live Q&A experience, to ask our engineering experts any questions you have about the incident.

Session option #1 - 20 September 2023 @ 01:00 UTC - Register at: <https://www.aka.ms/air6/reg/ash/1of2>

Session option #2 - 20 September 2023 @ 15:00 UTC - Register at: <https://www.aka.ms/air6/reg/ash/2of2>

What happened?

Starting at approximately 10:30 UTC on 30 August 2023, customers may have experienced issues accessing or using Azure, Microsoft 365 and Power Platform services. This incident was triggered by a utility power sag at 08:41 UTC on 30 August 2023, which impacted one of the three Availability Zones of the Australia East region. This power sag tripped a subset of the cooling system chiller units offline and, while working to restore cooling, temperatures in the datacenter increased to levels above operational thresholds. We powered down a small subset of selected compute and storage scale units, both to lower temperatures and to prevent damage to hardware. Although the vast majority of services recovered by 22:40 UTC on 30 August 2023, full mitigation was not until 20:00 UTC on 3 September 2023 – as some services experienced prolonged impact, predominantly as a result of dependencies on recovering subsets of Storage, SQL Database, and/or Cosmos DB services.

Multiple Azure services were impacted by this incident – including Azure Active Directory (AAD), Azure Active Directory B2C, Azure Active Directory Conditional Access, Azure Active Directory Connect Health, Azure Active Directory MyApps, Azure Activity Logs & Alerts, Azure API Management, Azure App Service, Azure Application Insights, Azure Arc enabled Kubernetes, Azure API for FHIR, Azure Backup, Azure Batch, Azure Chaos Studio, Azure Container Apps, Azure Container Registry, Azure Cosmos DB, Azure Databricks, Azure Data Explorer, Azure Data Factory, Azure Database for MySQL flexible servers, Azure Database for PostgreSQL flexible servers, Azure Digital Twins, Azure Device Update for IoT Hub, Azure Event Hubs, Azure ExpressRoute, Azure Health Data Services, Azure HDInsight, Azure IoT Central, Azure IoT Hub, Azure Kubernetes Service (AKS), Azure Logic Apps, Azure Log Analytics, Azure Log Search Alerts, Azure NetApp Files, Azure Notification Hubs, Azure Redis Cache, Azure Relay, Azure Resource Manager (ARM), Azure Role Based Access Control (RBAC), Azure Search, Azure Service Bus, Azure Service Fabric, Azure SQL Database, Azure Storage, Azure Stream Analytics, Azure Virtual Machines, Microsoft Purview, and Microsoft Sentinel.

What went wrong and why?

Starting at approximately 08:41 UTC on 30 August 2023, a utility voltage sag was caused by a lightning strike on electrical infrastructure approximately 18 miles from the impacted Availability Zone of the Australia East region. The voltage sag caused cooling system chillers for multiple datacenters to shut down. While some chillers automatically restarted, 13 failed to restart and required manual intervention. To do so, the onsite team accessed the datacenter rooftop facilities,

where the chillers are located, and proceeded to sequentially restart chillers moving from one datacenter to the next. By the time the team reached the final five chillers requiring a manual restart, the water inside the pump system for these chillers (chilled water loop) had reached temperatures that were too high to allow them to be restarted. In this scenario, the restart is inhibited by a self-protection mechanism that acts to prevent damage to the chiller that would occur by processing water at the elevated temperatures. The five chillers that could not be restarted supported cooling for the two adjacent data halls which were impacted in this incident.

The two impacted data halls require at least four chillers to be operational. Before the voltage sag, our cooling capacity for these halls consisted of seven chillers, with five chillers in operation and two chillers in standby. At 10:30 UTC some networking, compute, and storage infrastructure began to shutdown automatically as data hall temperatures increased, impacting service availability. At 11:34 UTC, as temperatures continued to increase, our onsite datacenter team began a remote shutdown of remaining networking, compute, and storage infrastructure in the impacted data halls to protect data durability, infrastructure health, and address the thermal runaway. This shutdown allowed the chilled water loop to return to a safe temperature which allowed us to restart the chillers. This shutdown of infrastructure resulted in a further reduction of service availability for this Availability Zone.

The chillers were successfully brought back online at 12:12 UTC, and data hall temperatures returned to operational thresholds by 13:30 UTC. Power was then restored to the affected infrastructure and a phased process to bring the infrastructure back online commenced. All power to infrastructure was restored by 15:10 UTC. Once all networking and storage infrastructure had power restored, dependent compute scale units were then also returned to operation. As the underlying compute and storage scale units came online, dependent Azure services started to recover, but some services experienced issues coming back online.

From a storage perspective, seven storage scale units were impacted – five standard storage scale units, and two premium storage scale units. Availability impact to affected storage accounts began at 10:30 UTC as hardware shut down in response to elevated data hall temperatures. This was most impactful to storage accounts configured with the default local redundant storage (LRS), which is not resilient to a zonal failure. Accounts configured as zonally redundant (ZRS) remained 100% available, and accounts configured as geographically redundant (GRS) were eligible for customer-managed account failover. After power restoration, storage nodes started coming back online from 15:25 UTC. Four scale units required engineer intervention to check and reset some fault detection logic – the combination of investigation and management tooling performance problems delayed restoration of availability for these scale units. We also identified that some automation was incorrectly marking some already recovered nodes as unhealthy, which slowed storage recovery efforts. By 20:00 UTC, 99% of storage accounts had recovered. Restoring availability for the remaining <1% of storage accounts took more time due to hardware troubleshooting and replacement required on a small number of storage nodes in a single scale unit. Even identifying problematic hardware in the storage nodes took an extended period of time, as the nodes were offline and therefore not able to provide diagnostics. By 01:00 UTC on 31 August, availability was restored for all except a handful of storage accounts, with complete availability restored by 07:00 UTC on 1 September.

From a SQL Database perspective, database capacity in a region is divided into tenant rings. The Australia East region includes hundreds of rings, and each ring consists of a group of VMs (10-200) hosting a set of databases. Rings are managed by Azure Service Fabric to provide availability in cases of VM, network or storage failures. When infrastructure was powered down, customers using zone-redundant Azure SQL Database did not experience any downtime, except for a small subset of customers using proxy mode connection, due to one connectivity gateway not being configured with zonal-resilience. The fix for this issue was already being rolled out, but had not yet deployed to the Australia East region. As infrastructure was powered back on, all tenant rings except one came back online and databases became available to customers as expected. However, one ring remained impacted even after Azure Compute became available. In this ring, 20 nodes did not come back online as expected, so databases on these nodes continued to experience unavailability. As a result of Service Fabric attempting to move databases to healthy nodes, other databases on this ring experienced intermittent availability issues as a side-effect of the overall replica density and unhealthy nodes. The recovery involved first moving all the databases from unhealthy nodes to healthy nodes. All remote storage (general purpose) databases were successfully recovered by this move, but databases using local storage (business critical) only recovered as their underlying nodes recovered. All databases on unhealthy nodes were recovered by 11:00 UTC on 31 August. Since the health and capacity of this final ring did not completely recover, we decided to move all databases out of the ring, which extended the overall recovery time but did not negatively impact customer availability. During this extended recovery,

most customers were not experiencing any issues but it was important to move all databases out of this unhealthy ring to prevent any potential impact. The operation of moving all databases out of this ring was completed at 20:00 UTC on 3 September. During this incident, customers who had 'active geo-replication' setup were able to failover manually to restore availability. For customers who have 'auto-failover groups' enabled, we did not execute automatic failover – our automatic failover policy was not initiated for the region, due to an incorrect initial assessment of the impact severity to SQL Database.

From a Cosmos DB perspective, zone resilient accounts and those with multi-region writes remained operational during the incident, transparently serving requests from a different zone or region, respectively. However, accounts not configured for AZ or multi-region writes experienced full or partial loss of availability, due to the infrastructure that was powered down. Multi-region accounts with single region write eligible for failover (those with Service Managed Failover enabled) were failed over to their alternate regions to restore availability. These were initiated at 12:07 UTC, 33 minutes after decision to power down scale units. The reason for this delay was to identify and failover the Cosmos DB control plane system resources – in retrospect this delay was unnecessary, as the Cosmos DB control plane was already fully zone-resilient. 95% of database accounts were failed over within 35 minutes by 12:42 UTC, and all eligible accounts were failed over by 16:13 UTC on 30 August. Accounts that were not eligible for failover had service restored to partitions only once the dependent storage and compute were restored.

From an Azure Kubernetes Service (AKS) perspective, the service experienced a loss of compute for the AKS control plane for Australia East as well as data access loss to SQL Database. The AKS control plane underlay is deployed across multiple availability zones. AKS uses Azure SQL Database for its operation queue which is used for Create/Read/Update/Delete (CRUD) activities. Although scheduled to be converted, the SQL Database in Australia East was not configured with AZ resiliency selection, leaving it unavailable during the incident period. In addition, AKS services in the Australia Southeast region depended on this same database, causing an AKS incident for CRUD activities in that region also. Existing customer workloads running on AKS clusters in either region should not have been impacted by the downtime, as long as they did not need to access the AKS resource provider for scaling or other CRUD activities. As the SQL Database recovered, service was restored without any other mitigation required.

From an Azure Resource Manager (ARM) perspective, the impact on customers was the result of degradation in Cosmos DB. This degradation impacted ARM between 10:45 UTC and 12:25 UTC and resulted in ARM platform availability for the Australia East region dropping from ~99.999% to (at its lowest) 88%, with a 62% success rate for write operations. For data consistency reasons, write operations are required to be sent to the associated Cosmos DB regional replica for a given resource group. While the migration to our next generation zonally-redundant storage architecture is still ongoing, it has not been completed and as a result this region is not yet leveraging fully zonally redundant storage for ARM. This meant that for the duration of the incident, customers worldwide trying to manage resources whose resource groups were homed in Australia East saw increased error rates (and this manifested in a small impact to global platform availability until 15:00 UTC).

How did we respond?

30 August 2023 @ 08:41 UTC – Voltage sag occurred on utility power line
30 August 2023 @ 08:43 UTC – 13 chillers failed to restart automatically
30 August 2023 @ 08:51 UTC – Remote resets on chillers commenced
30 August 2023 @ 09:09 UTC – Team arrived at first group of chillers for manual restarts
30 August 2023 @ 09:18 UTC – Team arrived at second group of chillers for manual restarts
30 August 2023 @ 09:42 UTC – Team arrived at third group of chillers for manual restarts
30 August 2023 @ 09:45 UTC – Team arrived at the final group of chillers which could not be restarted
30 August 2023 @ 10:30 UTC – Initial impact from automated infrastructure shutdown
30 August 2023 @ 10:47 UTC – Cosmos DB Initial impact detected via monitoring
30 August 2023 @ 10:48 UTC – First automated communications sent to Azure Service Health
30 August 2023 @ 11:30 UTC – Initial communications posted to public Azure Status page

30 August 2023 @ 11:36 UTC – All subscriptions in Australia East sent portal communications
30 August 2023 @ 11:34 UTC – Decision made to shutdown impacted infrastructure
30 August 2023 @ 12:07 UTC – Failover initiated for eligible Cosmos DB accounts
30 August 2023 @ 12:12 UTC – Five chillers manually restarted
30 August 2023 @ 13:30 UTC – Data hall temperature normalized
30 August 2023 @ 14:10 UTC – Safety walkthrough completed for both data halls
30 August 2023 @ 14:25 UTC – Decision made to start powering up hardware in the two affected data halls
30 August 2023 @ 15:10 UTC – Power restored to all hardware
30 August 2023 @ 15:25 UTC – Storage, networking, and compute infrastructure started coming back online after power restoration
30 August 2023 @ 15:30 UTC – Identified three specific storage scale units still experiencing fault codes
30 August 2023 @ 16:00 UTC – 35% of VMs recovered / Began manual recovery efforts for three remaining storage scale units
30 August 2023 @ 16:13 UTC – Account failover completed for all Cosmos DB accounts
30 August 2023 @ 17:00 UTC – All but one SQL Database tenant ring had recovered
30 August 2023 @ 19:00 UTC – 99% of storage accounts were back online
30 August 2023 @ 19:20 UTC – 90% of VMs recovered
30 August 2023 @ 19:29 UTC – Successfully recovered all premium storage scale units
30 August 2023 @ 22:35 UTC – Standard storage scale units were recovered, except for one scale unit
30 August 2023 @ 22:40 UTC – 99% of VMs recovered
31 August 2023 @ 05:00 UTC – 100% of VMs recovered
31 August 2023 @ 04:04 UTC – Restoration of Cosmos DB accounts to Australia East initiated
31 August 2023 @ 04:43 UTC – Final Cosmos DB cluster recovered, restoring all traffic for accounts that were not failed over to alternate regions
1 September 2023 @ 06:40 UTC – Successfully recovered all standard storage scale units
3 September 2023 @ 20:00 UTC – Final SQL Database tenant ring evacuated and all customer databases online

How are we making incidents like this less likely or less impactful?

Incremental load increases over time in the Availability Zone resulted in a chiller operating configuration that was susceptible to a timing defect in the Chiller Management System. We have de-risked this failure to restart due to voltage fluctuations, by implementing a change to the control timing logic on the Chiller Management System (Completed). Emergency Operation Procedure for manual restarts of simultaneous chiller failures has changed from an 'adjacent' sequence to a data hall 'load based' sequence to ensure all impacted data halls have partial cooling, to slow thermal runaway while full cooling is being restored (Completed).

Following this incident, as a temporary mitigation, we increased technician staffing levels at the datacenter to be prepared to execute manual restart procedures of our chillers prior to the change to the Chiller Management System to prevent restart failures. Based on our incident analysis the staffing levels at the time would have been sufficient to prevent impact if a 'load based' chiller restart sequence had been followed, which we have since implemented (Completed).

Datacenter staffing levels published in the Preliminary PIR only accounted for "critical environment" staff onsite. This did not characterize our total datacenter staffing levels accurately. To alleviate this misconception, we made a change to the preliminary public PIR posted on the Status History page.

Our Storage team has identified several optimizations in our large scale recovery process which will help to reduce time to mitigate. This includes augmenting data provided in our incidents to enable quicker decision making, and updates to our troubleshooting guides (TSGs) that enable faster execution (Estimated completion: December 2023).

Our Azure Service Fabric team is working to improve reliability of SQL Database tenant ring recovery. (Estimated completion: December 2023).

Our SQL Database team is reviewing our 'auto-failover group' trigger criteria, to ensure that failovers can happen within the expected timeframe. (Estimated completion: October 2023).

Our SQL Database team is upgrading internal tooling to enable mass migration of databases. (Estimated completion: December 2023).

Our Cosmos DB team is working to optimize Service Managed Failover for single region write accounts to reduce time to mitigate (Estimated completion: November 2023).

Our AKS team is immediately converting all operation queue SQL Database databases to be zone redundant (Estimated completion: September 2023).

Our AKS team is also replacing all cross-region SQL Database queue usage with Service Bus queues that are zone redundant (Estimated completion: September 2023).

Our ARM team will complete its storage layer migration to the next generation, zonally redundant architecture (Estimated completion: December 2023).

Our incident management team is exploring ways to harden our readiness, process, and playbook surrounding power down scenarios (Estimated completion: October 2023)

How can customers make incidents like this less impactful?

Consider using Availability Zones (AZs) to run your services across physically separate locations within an Azure region. To help services be more resilient to datacenter-level failures like this one, each AZ provides independent power, networking, and cooling. Many Azure services support zonal, zone-redundant, and/or always-available configurations:

<https://docs.microsoft.com/azure/availability-zones/az-overview>

For mission-critical workloads, customers should consider a multi-region geodiversity strategy to avoid impact from incidents like this one that impacted a single region: <https://learn.microsoft.com/training/modules/design-a-geographically-distributed-application/> and <https://learn.microsoft.com/azure/architecture/patterns/geodes>

Consider which are the right storage redundancy options for your critical applications. Zone redundant storage (ZRS/GZRS) remains available throughout a zone localized failure, like in this incident. Geo-redundant storage (GRS) enables account level failover in case the primary region endpoint becomes unavailable: <https://learn.microsoft.com/azure/storage/common/storage-redundancy>

Consider the relevant guidance for recovering your SQL Database databases during disaster recovery scenarios: <https://learn.microsoft.com/azure/azure-sql/database/disaster-recovery-guidance>

Consider the relevant guidance for achieving high availability with Azure Cosmos DB: <https://learn.microsoft.com/azure/cosmos-db/high-availability>

More generally, consider evaluating the reliability of your applications using guidance from the Azure Well-Architected Framework and its interactive Well-Architected Review: <https://docs.microsoft.com/azure/architecture/framework/resiliency>

Finally, consider ensuring that the right people in your organization will be notified about any future service issues – by configuring Azure Service Health alerts. These can trigger emails, SMS, push notifications, webhooks, and more: <https://aka.ms/ash-alerts>

How can we make our incident communications more useful?

You can rate this PIR and provide any feedback using our quick 3-question survey: <https://aka.ms/AzPIR/VVTQ-J98>

Sat Sep 02 2023 14:30:43 GMT+1000 (Australian Eastern Standard Time)

This is our "Preliminary" PIR that we endeavor to publish within 3 days of incident mitigation, to share what we know so far. After our internal retrospective is completed (generally within 14 days) we will publish a "Final" PIR with additional details/learnings.

What happened?

Between approximately 08:41 UTC on 30 August 2023 and 06:40 UTC on 1 September 2023 customers may have experienced issues accessing or using Azure, Microsoft 365 and Power Platform services. This event was triggered by a utility power sag in the Australia East region which tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, in an attempt to avoid damage to hardware.

Multiple downstream Azure services with dependencies on this infrastructure were also impacted – including Activity Logs & Alerts, API Management, App Service, Application Insights, Arc enabled Kubernetes, Azure API for FHIR, Backup, Batch, Chaos Studio, Container Apps, Container Registry, Cosmos DB, Databricks, Data Explorer, Data Factory, Database for MySQL flexible servers, Database for PostgreSQL flexible servers, Digital Twins, Device Update for IoT Hub, Event Hubs, ExpressRoute, Health Data Services, HDInsight, IoT Central, IoT Hub, Kubernetes Service (AKS), Logic Apps, Log Analytics, Log Search Alerts, Microsoft Sentinel, NetApp Files, Notification Hubs, Purview, Redis Cache, Relay, Search, Service Bus, Service Fabric, SQL Database, Storage, Stream Analytics, Virtual Machines. A small number of these services experienced prolonged impact, predominantly as a result of dependencies in recovering subsets of Storage, SQL, and/or Cosmos DB services.

What went wrong and why?

Starting at approximately 08:41 UTC on 30 August 2023, a utility power sag in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. We performed our documented Emergency Operational Procedures (EOP) to attempt to bring the chillers back online, but were not successful. The cooling capacity was reduced in two data halls for a prolonged time, so temperatures continued to rise. At 11:34 UTC, infrastructure thermal warnings from components in the affected data halls directed a shutdown of selected compute, network and storage infrastructure – by design, to protect data durability and infrastructure health. This resulted in a loss of service availability for a subset of this Availability Zone.

The cooling capacity for the two affected data halls consisted of seven chillers, with five chillers in operation and two chillers in standby (N+2) before the voltage dip event. When the event occurred, all five chillers in operation faulted and restart because the corresponding pumps did not get the run signal from the chillers. This is important as it is integral to the successful restarting of the chiller units. In addition, faulted chillers could not be restarted manually, as the chilled water loop temperature had exceeded the threshold. We had two chillers that were in standby which attempted to restart automatically – one managed to restart and came back online, the other restarted but was tripped offline again within minutes. Subsequently, thermal loads had to be reduced by shutting down servers. This successfully allowed the chilled water loop temperature to drop below the required threshold and enabled the restoration of the cooling capacity. When the data hall temperature levels were within operational thresholds, we began to restore power to the affected infrastructure and started a phased process to bring the infrastructure back online. Once all networking and storage infrastructure had power restored, dependent compute scale units were then also restored to operation. As the underlying compute and storage scale units came online, dependent Azure services started to recover, but some services experienced issues coming back online.

From a Storage perspective, seven tenants were impacted – five standard storage tenants, and two premium storage tenants. While all storage data is replicated across multiple storage servers, there were cases in which all of the copies were unavailable due to failures on multiple impacted storage servers. After power restoration, storage nodes started coming back online from 15:25 UTC. Generally speaking, there were three main factors that contributed to delays in bringing storage infrastructure back to full functionality. Firstly, the hardware damaged by the data hall temperatures required extensive troubleshooting. Diagnostics were not able to identify the faults, because the storage nodes themselves were not online – as a result, our onsite datacenter team needed to remove components manually, and re-seat them one by one to identify which particular component(s) were preventing each node from booting. Secondly, several components needed to be replaced for successful data recovery and to restore impacted nodes. In order to completely recover data,

some of the original/faulty components were required to be temporarily re-installed in individual servers. Thirdly, we identified that our automation was incorrectly approving stale requests, and marking some healthy nodes as unhealthy, which slowed storage recovery efforts.

From a SQL perspective, once power had fully returned to the impacted halls, service restoration was initially impacted by the slow recovery of dependent services, those being primarily Azure Standard Storage offerings. Many general purpose databases remained unavailable until those premium Azure Storage services had recovered. After storage services were >99% recovered, a single tenant ring that hosted databases failed to recover completely after the incident. This ring, hosting over ~250K databases had a mix of failure modes. Some databases may have been completely unavailable, some would have experienced intermittent connectivity issues, and some databases would have been fully available. This uneven impact profile for databases in the degraded ring, meant that it was difficult to summarize which customers were still impacted, which continued to present a challenge throughout the incident. As we attempted to migrate databases out of the degraded ring, SQL did not have well tested tools on hand that were built to move databases when the source ring was in degraded health scenario. Soon this became our largest impediment to mitigating impact. Our standard DB migration workflows are designed to contain many safety and health checks to make sure that the DB being migrated does not experience downtime during the migration. Each of these contact points from the SQL control plane to the underlying Service Fabric (which hosts the DB's compute and coordinates replicas) for the ring was another opportunity for the operation to fail or get stuck. Because every DB moved required manual mitigation via scripts, it seriously undermined our ability to move fast even once impacted DBs were identified, and DB moves were scheduled.

From a Cosmos DB perspective, the service experienced a loss of compute underpinning three clusters, and a loss of the operating system for 11 clusters due to dependencies on Azure storage. In total this mean that approximately half of Cosmos DB clusters in the Australia East region were either down or heavily degraded. Eligible accounts (multi-region with Service Managed Failover enabled) were failed over to their alternate regions to restore availability. A set of internal accounts were not originally configured for failover, so the Cosmos DB team worked with these internal service teams to configure and then failover those accounts upon request. Accounts that were not eligible had service restored to partitions once the dependent storage and compute were restored and Cosmos DB clusters recovered.

How did we respond?

30 August 2023 @ 08:41 – Voltage sag occurred on utility power line
30 August 2023 @ 08:43 – Five chillers failed to restart
30 August 2023 @ 08:45 – Two standby chillers started automatically
30 August 2023 @ 08:47 – One standby chiller tripped and went offline
30 August 2023 @ 09:41 – Onsite team arrived at the rooftop chiller area
30 August 2023 @ 09:42 – Onsite team manually restarted the five chillers as per the EOP
30 August 2023 @ 10:30 – Storage and SQL alerted by monitors about failure rates
30 August 2023 @ 10:57 – Cosmos DB Initial impact detected via monitoring
30 August 2023 @ 11:15 – Attempts to stabilize the five chillers were unsuccessful after multiple chiller restarts
30 August 2023 @ 11:20 – Chiller OEM support team arrives onsite
30 August 2023 @ 11:34 – Decision was made to shutdown infrastructure in the two affected data halls
30 August 2023 @ 12:07 – Failover initiated for eligible Cosmos DB accounts
30 August 2023 @ 12:12 – Five chillers manually restarted successfully
30 August 2023 @ 13:30 – Data hall temperature normalized
30 August 2023 @ 14:10 – Safety walkthrough completed for both data halls

30 August 2023 @ 14:25 – Decision made to start powering up hardware in the two affected data halls
30 August 2023 @ 15:10 – Power restored to all hardware
30 August 2023 @ 15:25 – Storage infrastructure started coming back online after power restoration
30 August 2023 @ 15:30 – Identified three specific storage tenants still experiencing fault codes
30 August 2023 @ 16:00 – Began manual recovery efforts for these three storage tenants
30 August 2023 @ 16:13 – Account failover completed for all Cosmos DB accounts
30 August 2023 @ 19:29 – Successfully recovered all premium storage tenants
30 August 2023 @ 20:29 – All but two SQL nodes recovered
30 August 2023 @ 22:35 – Standard storage tenants were recovered, except for one scale unit
31 August 2023 @ 04:04 – Restoration of Cosmos DB accounts to Australia East initiated
31 August 2023 @ 04:43 – Final Cosmos DB cluster recovered, restoring all traffic for accounts that were not failed over
31 August 2023 @ 08:45 – All external customer accounts back online and operating from Australia
1 September 2023 @ 06:40 – Successfully recovered all standard storage tenants

How are we making incidents like this less likely or less impactful?

Based on our initial assessment, we have already identified the following learnings from a datacenter power/cooling perspective. The Final PIR will include additional learnings and repairs based on the service-specific extended recovery timelines.

Due to the size of the datacenter campus, the staffing of the team at night was insufficient to restart the chillers in a timely manner. We have temporarily increased the team size from three to seven, until the underlying issues are better understood and appropriate mitigations can be put in place.

The EOP for restarting chillers is slow to execute for an event with such a significant blast radius. We are exploring ways to improve existing automation to be more resilient to various voltage sag event types.

Moving forward, we are evaluating ways to ensure that the load profiles of the various chiller subsets can be prioritized so that chiller restarts will be performed for the highest load profiles first.

Utilizing the playbook in sequencing workload failovers and equipment shutdown could have been prioritized differently with better insights. We are working to improve reporting on chilled water temperature, to enable more timely decisions for failover/shutdown based on thresholds.

The five chillers did not manage to restart because the corresponding pumps did not get the run signal from the chillers. This is important as it is integral to the successful restarting of the chiller units. We are partnering with our OEM vendor to investigate why the chillers did not command their respective pump to start.

One standby chiller did not automatically restart due to an unknown error. Our OEM vendor is running diagnostics to understand what caused this specific issue.

How can we make our incident communications more useful?

You can rate this PIR and provide any feedback using our quick 3-question survey: <https://aka.ms/AzPIR/VVTQ-J98>

Sat Sep 02 2023 01:35:02 GMT+1000 (Australian Eastern Standard Time)

Summary of Impact: Starting at approximately 08:30 UTC on 30 August 2023, a number of Azure customers in Australia East experienced degraded performance and availability issues for resources hosted in this region. Impact was limited to one of the three Availability Zones in a datacenter in this region. As of 16:07 UTC on 31 Aug 2023, most services were recovered.

Preliminary Root Cause: We determined a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted by this issue.

Mitigation: We worked on recovering the failed cooling units and reducing the overall temperature within the impacted area. Once temperature levels were within safe thresholds, we began to restore power to the affected infrastructure and started a phased process to bring this infrastructure back online. Once storage infrastructure was fully restored, dependent compute scale units were then also restored to operation. As the underlying compute and storage scale units became healthy, compute and other dependent Azure services recovered.

While we have broadly recovered, a small subset of services (SQL and Cosmos DB) are still working on post recovery checks, and we are closely monitoring the datacenter metrics for storage and compute resources to ensure they continue to show as healthy. For any residual customers with services still in the recovery process, we will continue communicating directly through the Azure Portal.

Next Steps: Next Steps: We will continue working on recovering a small number of instances still affected. Additionally:

Our services team will be completing an internal retrospective to understand the incident in more detail. We will publish a Preliminary Post Incident Review (PIR) within approximately 72 hours, to share more details on what happened and how we responded. After our internal retrospective is completed, generally within 14 days, we will publish a Final Post Incident Review with any additional details and learnings.

To get notified when that happens, and/or to stay informed about future Azure service issues, make sure that you configure and maintain Azure Service Health alerts – these can trigger emails, SMS, push notifications, webhooks, and more: <https://aka.ms/ash-alerts>.

For more information on Post Incident Reviews, refer to <https://aka.ms/AzurePIRs>. Finally, for broader guidance on preparing for cloud incidents, refer to <https://aka.ms/incidentreadiness>.

Fri Sep 01 2023 06:07:28 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are beginning to see the recovery on the impacted nodes. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with recovery ongoing. Impact is minimal at this time and most dependent services should begin to see recovery.

Power has been restored to all clusters hosting Cosmos DB. Accounts that were not failed over should continue to see recovery in Australia East. Accounts failed over to alternate regions will continue to operate in these alternate regions; the Cosmos DB team has made additional progress in bringing the Australia East region back online.

Over 99% of impacted Azure storage accounts are now fully recovered. Due to thermal damage, we continue working on recovering a small number of resources to bring the remaining impacted storage accounts back online. We have recovered 75% of the remaining impacted resources, and as we continue making progress, additional services are displaying indications of improvement.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided as events warrant.

Thu Aug 31 2023 22:55:12 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are beginning to see the recovery on the impacted nodes. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with recovery ongoing.

Power has been restored to all clusters hosting Cosmos DB. Accounts that were not failed over should see recovery in Australia East. Accounts failed over to alternate regions will continue to operate in these alternate regions; the Cosmos DB team has begun to bring Australia East region online.

Over 99% of impacted Azure storage accounts are now fully recovered. Due to thermal damage, we continue working

on recovering a small number of resources to bring the remaining impacted storage accounts back online. We have recovered 75% of the remaining impacted resources, and as we continue making progress, additional services are displaying indications of improvement.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided 2 hours, or as events warrant.

Thu Aug 31 2023 20:52:21 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are beginning to see the recovery on the impacted nodes. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with recovery ongoing.

Power has been restored to all clusters hosting Cosmos DB. Accounts that were not failed over should see recovery in Australia East. Accounts failed over to alternate regions will continue to operate in these alternate regions; the Cosmos DB team has begun to bring Australia East region online.

Over 99% of impacted Azure storage accounts are now fully recovered. Due to thermal damage, we continue working on replacing or recovering a small number of resources to bring the remaining impacted storage accounts back online. At this point we can confirm that we have recovered 50% of the remaining impacted resources, and as we continue making progress, additional services are displaying indications of improvement.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided 2 hours, or as events warrant.

Thu Aug 31 2023 19:35:06 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are beginning to see the recovery on the impacted nodes. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with recovery ongoing.

Power has been restored to all clusters hosting Cosmos DB. Accounts that were not failed over should see recovery in

Australia East. Accounts failed over to alternate regions will continue to operate in these alternate regions; the Cosmos DB team has begun to bring Australia East region online.

Over 99% of impacted Azure storage accounts are now fully recovered. Due to thermal damage, we continue working on replacing or recovering a small number of resources to bring the remaining impacted storage accounts back online.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 17:53:12 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are actively investigating individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with investigations ongoing.

Power has been restored to all clusters hosting Cosmos DB. Accounts that were not failed over should see recovery in Australia East. Accounts failed over to alternate regions will continue to operate in these alternate regions; the Cosmos DB team has begun onlining the Australia East region for accounts that were failed over earlier.

Over 99% of impacted Azure storage accounts are now fully recovered, a small number resources are currently being replaced or repaired due to thermal damage in order to bring the remaining impacted storage accounts back online.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 16:55:41 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are actively investigating individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with investigations ongoing.

Power has been restored to 100% of Cosmos DB clusters. Accounts who were not failed over should see recovery in Australia East. The Cosmos DB team has begun onlining this region for accounts that were failed over earlier.

Over 99% of impacted Azure storage accounts are now fully recovered, a small number of resources are currently being replaced or repaired due to thermal damage in order to bring the remaining impacted storage accounts back online.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 15:45:14 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are actively investigating individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with investigations ongoing.

Over 90% of Cosmos DB clusters in the Australia East region have been successfully recovered. Cosmos DB Accounts, that were not failed over to alternate read/write regions, should now begin to see signs of recovery. Once all clusters are fully restored, the Cosmos DB team will initiate the process of bringing this region back online for accounts that were previously failed over.

Over 99% of impacted Azure storage accounts are now fully recovered, a small number of resources are currently being replaced or repaired due to thermal damage in order to bring the remaining impacted storage accounts back online.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

We have assessed that achieving complete mitigation will require additional amount of time. Further updates on these remaining recovery efforts will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 14:38:57 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are actively investigating individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with investigations ongoing.

Over 90% of Cosmos DB clusters in the Australia East region have been successfully recovered. Cosmos DB Accounts, that were not failed over to alternate read/write regions, should now begin to see signs of recovery. Once all clusters are fully restored, the Cosmos DB team will initiate the process of bringing this region back online for accounts that were previously failed over.

Over 99% of impacted Azure storage accounts are now fully recovered, a small number of resources are currently being replaced or repaired due to thermal damage in order to bring the remaining impacted storage accounts back online.

Azure SQL is continuing to recover impact to a single backend cluster, impacting a small subset of databases instances. Additionally, we are taking steps to reduce the overall zonal load to facilitate improved resource allocation and recovery.

Further updates on these remaining recovery efforts will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 13:36:16 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: While we have successfully recovered 99% of storage services and impacted Virtual Machines, we continue to actively investigate individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of the three services mentioned below, with investigations ongoing.

Storage - Our Storage team is making progress with the final remaining storage scale unit that is still experiencing isolated issues. We have engaged our onsite datacenter team to support replacing drives as needed.

SQL - Our SQL team is working to mitigate one final cluster that is experiencing a capacity issue, due to several Service Fabric nodes that have not fully recovered. We are rebalancing capacity to mitigate.

Cosmos DB - Our Cosmos DB team continues to investigate why some services have not yet recovered fully while the majority of customers and services are already mitigated.

A further update on the remaining investigation will be provided within 60 minutes, or as events warrant.

Thu Aug 31 2023 12:28:49 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East

region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: While we have successfully recovered 99% of storage services and impacted Virtual Machines, we continue to actively investigate individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of the three services mentioned below, with investigations ongoing.

Storage - Our Storage team is making progress with the final remaining storage scale unit that is still experiencing isolated issues. We have engaged our onsite datacenter team to support replacing drives as needed.

SQL - Our SQL team is working to mitigate one final cluster that is experiencing a capacity issue, due to several Service Fabric nodes that have not fully recovered. We are rebalancing capacity to mitigate.

Cosmos DB - Our Cosmos DB team continues to investigate why some services have not yet recovered fully while the majority of customers and services are already mitigated.

The next update will be provided within 60 minutes, or as events warrant.

Thu Aug 31 2023 11:10:00 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are actively investigating individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with investigations ongoing. Firstly, our Storage team are making progress with the final remaining storage scale unit that is still experiencing isolated issues - we have engaged our onsite datacenter team to support replacing drives as needed. Secondly, our SQL team are working to mitigate one final cluster that is experiencing a capacity issue, due to several Service Fabric nodes that have not fully recovered - we are rebalancing capacity to mitigate. Finally, our Cosmos DB team continue to investigate why some services have not yet recovered fully. While the majority of customers and the majority of services are already mitigated, further updates on these remaining investigations will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 09:49:39 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: Having successfully recovered 99% of storage services and 99% of impacted Virtual Machines, we are actively investigating individual downstream services to confirm their recovery status and mitigate remaining issues. At this stage, we believe most downstream services that are still experiencing impact are the result of dependencies on one of three services with investigations ongoing. Firstly, our Storage team are making progress with the final remaining storage scale unit that is still experiencing isolated issues - we have engaged our onsite datacenter team to support replacing drives as needed. Secondly, our SQL team are working to mitigate one final cluster that is experiencing a capacity issue, due to several Service Fabric nodes that have not fully recovered - we are rebalancing capacity to mitigate. Finally, our Cosmos DB team continue to investigate why some services have not yet recovered fully. While the majority of customers and the majority of services are already mitigated, further updates on these remaining investigations will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 08:48:09 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware. Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health.

Current Status: With 99% of storage services and 99% of impacted Virtual Machines back online and healthy, we are actively investigating remaining issues with individual downstream services to confirm their recovery status. Our Storage team are making progress on one specific storage scale unit that is still experiencing isolated issues. Our SQL team are investigating a potential issue with an underlying Service Fabric dependency. Our Cosmos DB team are investigating why some services have not fully recovered. Despite these remaining investigations, the majority of customers and services should already be recovered. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 07:50:51 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware.

Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health. Impact to services is limited to Australia East, except for Azure Kubernetes Service (AKS) which has impact in both Australia East and Australia Southeast due to a dependency in the former.

Current Status: With 99% of storage services and 99% of impacted Virtual Machines back online and healthy, we are now supporting individual downstream services to confirm their recovery status. We are aware of one specific storage scale unit that is still experiencing isolated issues, but the majority of customers and services should already be recovered. Beyond this known storage issue, we are investigating which services are still not fully mitigated and why. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 06:54:44 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware.

Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health. Impact to services is limited to Australia East, except for Azure Kubernetes Service (AKS) which has impact in both Australia East and Australia Southeast due to a dependency in the former. If your workloads are protected by Azure Site Recovery or Azure Backup, and you need critical services back online before all services in this datacenter are fully recovered, we recommend either to initiate a failover to the recovery region or recover using Cross Region Restore. Note that any new allocation requests for the Australia East region will automatically avoid the impacted scale units.

Current Status: We are in the final phases of restoring core services, and expect that the vast majority of remaining impacted services should be back online in the next hour. After restoring power and stabilizing temperatures, all network infrastructure and 99% of storage services are back online. All premium disk storage has fully recovered, we continue to work towards mitigating the final remaining storage devices. The vast majority of underlying compute services are back online, with more than 99% of Virtual Machines (VMs) that were impacted now back online and healthy.

While many customers and services have already recovered, we are now prioritizing our investigations with the remaining downstream impacted services. We expect that these remaining services should be back online and healthy within the next hour. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 05:55:40 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware.

Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health. Impact to services is limited to Australia East, except for Azure Kubernetes Service (AKS) which has impact in both Australia East and Australia Southeast due to a dependency in the former. If your workloads are protected by Azure Site Recovery or Azure Backup, and you need critical services back online before all services in this datacenter are fully recovered, we recommend either to initiate a failover to the recovery region or recover using Cross Region Restore. Note that any new allocation requests for the Australia East region will automatically avoid the impacted scale units.

Current Status: We have made significant progress in restoring core services, and expect that the vast majority of remaining services should be back online in the next 1-2 hours. After restoring power and stabilizing temperatures, all network infrastructure and 95% of storage services are back online. All premium disk storage has fully recovered, we continue to work towards mitigating the final remaining storage devices. The majority of underlying compute services are back online, with more than 85% of Virtual Machines (VMs) that were impacted now back online and healthy. For the remaining VMs, we are investigating potential issues in connecting to their corresponding storage services.

While many customers have already recovered, we continue to work with downstream impacted services to ensure that they are coming back online in the next 1-2 hours as expected. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 04:59:04 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one datacenter, within one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units, to avoid damage to hardware.

Multiple downstream services were impacted, with targeted communications being distributed via Azure Service Health. Impact to services is limited to Australia East, except for Azure Kubernetes Service (AKS) which has impact in both Australia East and Australia Southeast due to a dependency in the former. If your workloads are protected by Azure Site Recovery or Azure Backup, and you need critical services back online before all services in this datacenter are fully recovered, we recommend either to initiate a failover to the recovery region or recover using Cross Region Restore. Note that any new allocation requests for the Australia East region will automatically avoid the impacted scale units.

Current Status: Mitigation efforts are continuing, we have made significant progress in restoring core services, and we expect that the vast majority of remaining services should be back online in the next 2-3 hours. After restoring power and stabilizing temperatures, all network infrastructure and 95% of storage services are back online. All premium disk storage has fully recovered, we continue to work towards mitigating the final remaining storage devices. The majority of underlying compute services are back online, with more than 85% of Virtual Machines that were impacted now back online and healthy. As a result, many customers of these services have already recovered - but we continue to work with downstream impacted services to ensure that they are coming back online in the next 2-3 hours as expected. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 04:07:32 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units to avoid damage to hardware. All impacted storage and compute scale units were in the same datacenter, within one of the region's three Availability Zones (AZs).

Multiple downstream services have been identified as impacted, with targeted communications being distributed via Azure Service Health. Impact to services is limited to Australia East, except for Azure Kubernetes Service (AKS) which has impact in both Australia East and Australia Southeast due to a dependency in the former. If your workloads are protected by Azure Site Recovery or Azure Backup, and you need critical services back online before this datacenter is fully recovered, we recommend to either initiate a failover to the recovery region or recover using Cross Region Restore. Note that any new allocations for resources in the Australia East region will automatically avoid the impacted scale units.

Current Status: Mitigation efforts are continuing, we have made significant progress in restoring core services but we are not able to provide a mitigation ETA at this time. Power to all hardware has been restored, temperatures in the impacted datacenter have stabilized. All network infrastructure is back online. The majority of storage devices are back online, we

are validating issues with a few remaining storage nodes. The majority of underlying compute services are back online, with more than 75% of Virtual Machines that were impacted back online and healthy. While many customers of these core services have seen signs of recovery, we continue to work with downstream impacted services to ensure that they are coming back online as expected. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 03:06:34 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one of the Availability Zones. While working to restore cooling, temperatures in the datacenter increased so we proactively powered down a small subset of selected compute and storage scale units to avoid damage to hardware. All impacted storage and compute scale units are in the same datacenter, within one of the region's three Availability Zones (AZs). Multiple downstream services have been identified as impacted, with targeted communications being distributed via Azure Service Health.

If your workloads are protected by Azure Site Recovery or Azure Backup, and you need critical services back online before this datacenter is fully recovered, we recommend to either initiate a failover to the recovery region or recover using Cross Region Restore. Note that any new allocations for resources in the Australia East region will automatically avoid the impacted scale units.

Current Status: We are not able to provide a mitigation ETA at this time, but power to all hardware has been restored, and temperatures in the impacted datacenter have been stabilized. We are now making progress on service recovery - cautiously, starting with core services. The vast majority of network infrastructure is back online, and storage device recovery has started. Due to the nature of this issue our storage scale units are expected to require additional recovery efforts to ensure that all resources return in a consistent state. As service recovery continues, some customers may start experiencing signs of recovery. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 01:58:46 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one of the Availability Zones. While working to restore the cooling units, temperatures in the datacenter increased so we have proactively powered down a small subset of selected compute and storage scale units to avoid damage to hardware and reduce cooling system load. All impacted storage and compute scale units are in the same datacenter, within one of the region's three Availability Zones (AZs). Multiple downstream services have been identified as impacted.

Current Status: We do not have an exact ETA at this time, but temperatures in the impacted datacenter have been stabilized. The Azure service recovery process has commenced and is expected to progress over a number of hours. Due to the nature of this issue our storage scale units are expected to require additional recovery efforts to ensure all resources return in a consistent state. Note that any new allocations for resources will automatically avoid the impacted scale units. If your workloads are protected by Azure Site Recovery or Azure Backup, we recommend to either initiate a failover to the recovery region or recover using Cross Region Restore.

All impacted hardware previously powered off as a preventative measure has now been powered back on. We continue our mitigation efforts to bring all affected services back online, and as this continues, some customers may start experiencing signs of recovery. Further updates will be provided in 60 minutes, or as events warrant.

Thu Aug 31 2023 01:01:02 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one of the Availability Zones. While working to restore the cooling

units, temperatures in the datacenter increased so we have proactively powered down a small subset of selected compute and storage scale units to avoid damage to hardware and reduce cooling system load. All impacted storage and compute scale units are in the same datacenter, within one of the region's three Availability Zones (AZs). Multiple downstream services have been identified as impacted.

Current Status: We do not have an exact ETA at this time, but temperatures in the impacted datacenter have been stabilized. The Azure service recovery process has commenced and is expected to progressively return over a number of hours. Due to the nature of this issue our storage scale units are expected to require additional recovery efforts to ensure all resources return in a consistent state. Note that any new allocations for resources will automatically avoid the impacted scale units. If your workloads are protected by Azure Site Recovery or Azure Backup, we recommend to either initiate a failover to the recovery region or recover using Cross Region Restore.

Our mitigation efforts to restore our services are ongoing, and further updates will be provided in 1 hour or as events warrant.

Wed Aug 30 2023 23:34:17 GMT+1000 (Australian Eastern Standard Time)

Impact Statement: Starting at approximately 08:30 UTC on 30 August 2023, a utility power surge in the Australia East region tripped a subset of the cooling units offline in one of the Availability Zones. While working to restore the cooling units, temperatures in the datacenter increased so we have proactively powered down a small subset of selected compute and storage scale units to avoid damage to hardware and reduce cooling system load. All impacted storage and compute scale units are in the same datacenter, within one of the region's three Availability Zones (AZs). Multiple downstream services have been identified as impacted.

Current Status: We do not have an exact ETA at this time, but temperature in the impacted datacenter have been stabilized. The Azure service recovery process has commenced and is expected to progressively return over a number of hours. Due to the nature of this issue our storage scale units are expected to require additional recovery efforts to ensure all resources return in a consistent state. Note that any new allocations for resources will automatically avoid the impacted scale units. If your workloads are protected by Azure Site Recovery or Azure Backup, we recommend to either initiate a failover to the recovery region or recover using Cross Region Restore.

Further updates will be provided in an hour or as events warrant.

Wed Aug 30 2023 21:48:25 GMT+1000 (Australian Eastern Standard Time)

Summary of Impact: Starting at 08:30 UTC on 30 August 2023, a subset of customers with workloads hosted in the Australia East may be experiencing difficulties accessing and managing resources deployed in this region.

Current Status: We are experiencing impact related to a cooling issue in a section of a datacentre in Australia East. This is resulting in connectivity and availability issues for services with dependencies on storage or compute in this section. More updates will be provided in an hour or as events warrant.

Wed Aug 30 2023 21:36:34 GMT+1000 (Australian Eastern Standard Time)

Impact statement: Starting at 08:30 UTC on 30 August 2023, a subset of customers with workloads hosted in the Australia East region may be experiencing difficulties accessing and managing resources deployed in these regions.

Current Status: We are actively working on a cooling event which has impacted a specific datacenter in the Australia East region. The result of this event is that resources might experience unexpected restarts. Further updates will be provided in an hour or as events warrant.