Preliminary RCA - Authentication errors across multiple Microsoft services (Tracking ID LN01-P8Z)

Summary of impact: Between 19:00 UTC (approx) on March 15, 2021, and 09:25 UTC on March 16, 2021 customers may have encountered errors performing authentication operations for any Microsoft and third-party applications that depend on Azure Active Directory (Azure AD) for authentication.

Azure Admin Portal, Teams, Exchange, Azure KeyVault, SharePoint, Storage and other major applications have recovered. Any customers experiencing residual impact will continue to receive updates regarding these via their Azure Service Health notifications.

Preliminary Root Cause: The preliminary analysis of this incident shows that an error occurred in the rotation of keys used to support Azure AD's use of OpenID, and other, Identity standard protocols for cryptographic signing operations. As part of standard security hygiene, an automated system, on a time-based schedule, removes keys that are no longer in use. Over the last few weeks, a particular key was marked as "retain" for longer than normal to support a complex cross-cloud migration. This exposed a bug where the automation incorrectly ignored that "retain" state, leading it to remove that particular key.
Metadata about the signing keys is published by Azure AD to a global location in line with Internet Identity standard protocols. Once the public metadata was changed at 19:00 UTC, applications using these protocols with Azure AD began to pick up the new metadata and stopped trusting tokens/assertions signed with the key that was removed. At that point, end users were no longer able to access those applications.

Mitigation: Service telemetry identified the problem, and the engineering team was automatically engaged. The key removal operation was identified as the cause, and the key metadata was rolled back to its prior state at 21:05 UTC.

Applications need to pick up the rolled back metadata and refresh their caches with the correct metadata. Time to mitigation for individual applications varies due to a variety of server implementations that handle caching differently. Azure Admin Portal, Teams, Exchange, Azure Key Vault, SharePoint and other major applications have recovered. A subset of Storage resources experienced residual impact due to cached metadata, and we pushed an update to invalidate these entries and force a refresh. This process completed and mitigation for the residually impacted customers was declared at 09:25 UTC

Azure AD is in a multi-phase effort to apply additional protections to the backend Safe Deployment Process (SDP) system to prevent a class of risks including this problem. The first phase does provide protections for adding a new key, but the remove key component is in the second phase which is scheduled to be finished by mid-year. A previous Azure AD incident occurred on September 28th, 2020 and both incidents are in the class of risks that will be prevented once the multi-phase SDP effort is completed.

Next Steps: We understand how incredibly impactful and unacceptable this is and apologize deeply. We are continuously taking steps to improve the Microsoft Azure Platform and our processes to help ensure such incidents do not occur in the future. In the September incident we indicated our plans to "apply additional protections to the Azure AD service backend SDP system to prevent the class of issues identified here."

- The first phase of those SDP changes is finished, and the second phase is in a very carefully staged deployment that will finish mid-year. The initial analysis does indicate that once that is fully deployed, it will prevent the type of outage that happened today, as well as the related incident in September 2020. In the meantime, additional safeguards have been added to our key removal process which will remain until the second phase of the SDP deployment is completed.
- In that September incident we also referred to our rollout of Azure AD backup authentication. That effort is progressing well. Unfortunately, it did not help in this case as it provided coverage for token issuance but did not provide coverage for token validation as that was dependent on the impacted metadata endpoint.

The Root Cause Analysis investigation relating to this incident is ongoing, and a full RCA will be published when this is completed, or if any other substantive details emerge in the interim.