

Volume
Licensing

Service Level Agreement for Microsoft Online Services July 1, 2019



Table of Contents

TABLE OF CONTENTS	2	
INTRODUCTION	4	
GENERAL TERMS	5	
SERVICE SPECIFIC TERMS	7	
MICROSOFT DYNAMICS 365	7	
DYNAMICS 365 FOR CUSTOMER SERVICE ENTERPRISE; DYNAMICS 365 FOR CUSTOMER SERVICE PROFESSIONAL; DYNAMICS 365 CUSTOMER SERVICE INSIGHTS	7	
DYNAMICS 365 BUSINESS CENTRAL	7	
DYNAMICS 365 FOR FINANCE AND OPERATIONS (ENTERPRISE EDITION)	7	
DYNAMICS 365 FOR RETAIL	8	
DYNAMICS 365 FOR SALES ENTERPRISE; DYNAMICS 365 FOR SALES PROFESSIONAL	9	
DYNAMICS 365 FOR TALENT; DYNAMICS 365 FOR TALENT: ATTRACT; DYNAMICS 365 FOR TALENT: ONBOARD	9	
OFFICE 365 SERVICES.....	9	
DUET ENTERPRISE ONLINE	9	
EXCHANGE ONLINE	10	
EXCHANGE ONLINE ARCHIVING	10	
EXCHANGE ONLINE PROTECTION	11	
MICROSOFT MYANALYTICS	11	
MICROSOFT STREAM	12	
MICROSOFT TEAMS	12	
OFFICE 365 BUSINESS	12	
OFFICE 365 ADVANCED COMPLIANCE	13	
OFFICE 365 PROPLUS.....	13	
OFFICE ONLINE.....	13	
OFFICE 365 VIDEO.....	14	
ONEDRIVE FOR BUSINESS	14	
PROJECT ONLINE.....	14	
SHAREPOINT ONLINE	15	
SKYPE FOR BUSINESS ONLINE	15	
MICROSOFT TEAMS – CALLING PLANS AND AUDIO CONFERENCING	15	
MICROSOFT TEAMS – VOICE QUALITY	16	
WORKPLACE ANALYTICS	16	
YAMMER ENTERPRISE	17	
MICROSOFT AZURE SERVICES	17	
AD DOMAIN SERVICES	17	
ANALYSIS SERVICES.....	17	
API MANAGEMENT SERVICES	18	
APP SERVICE	18	
APPLICATION GATEWAY.....	19	
APPLICATION INSIGHTS (QUERY AVAILABILITY SLA)	19	
AUTOMATION SERVICE – DESIRED STATE CONFIGURATION (DSC).....	20	
AUTOMATION SERVICE – PROCESS AUTOMATION	20	
AZURE ADVANCED THREAT PROTECTION	21	
AZURE BOT SERVICE	21	
AZURE CONTAINER INSTANCES	21	
AZURE COSMOS DB	22	
AZURE DATABASE FOR MYSQL	25	
AZURE DATABASE FOR POSTGRESQL.....	25	
AZURE DDoS PROTECTION	25	
AZURE DNS	26	
AZURE FIREWALL	26	
AZURE FUNCTIONS.....	26	
AZURE LAB SERVICES	27	
AZURE LOAD BALANCER	27	
AZURE MAPS API.....	28	
AZURE MONITOR.....	28	
AZURE MONITOR ALERTS	29	
AZURE MONITOR NOTIFICATION DELIVERY	29	
AZURE SECURITY CENTER.....	29	
AZURE VIRTUAL WAN	30	
BATCH SERVICE	30	
BACKUP SERVICE	31	
BIZTALK SERVICES.....	31	
CACHE SERVICES.....	32	
CDN SERVICE	32	
CLOUD SERVICES	33	
CONTAINER REGISTRY	33	
DATA CATALOG	34	
DATA FACTORY – ACTIVITY RUNS	34	
DATA FACTORY – API CALLS	35	
DATA LAKE ANALYTICS	35	
DATA LAKE STORE.....	35	
EVENT GRID	36	
EXPRESSROUTE	36	
FUNCTION APP ON CONSUMPTION PLAN	37	
FUNCTION APP ON SERVICE PLAN	37	
HDINSIGHT	37	
HOCKEYAPP	38	
IOT CENTRAL	38	
IOT HUB	39	
KEY VAULT	39	
LOG ANALYTICS (QUERY AVAILABILITY SLA)	40	
LOGIC APPS	40	
AZURE MACHINE LEARNING STUDIO – BATCH EXECUTION SERVICE (BES) AND MANAGEMENT APIS SERVICE.....	40	
AZURE MACHINE LEARNING STUDIO – REQUEST RESPONSE SERVICE (RRS) ..	41	
MEDIA SERVICES – CONTENT PROTECTION SERVICE.....	41	
MEDIA SERVICES – ENCODING SERVICE.....	42	
MEDIA SERVICES – MEDIA INDEXER SERVICE.....	42	
MEDIA SERVICES – LIVE CHANNELS.....	42	
MEDIA SERVICES – STREAMING SERVICE	43	
MEDIA SERVICES – VIDEO INDEXER SERVICE	43	
MICROSOFT COGNITIVE SERVICES	44	
MICROSOFT GENOMICS	44	
MOBILE ENGAGEMENT	45	
MOBILE SERVICES	45	
NETWORK WATCHER.....	45	
REMOTEAPP.....	46	
SAP HANA ON AZURE	46	
SCHEDULER	47	
SEARCH	48	
SERVICE-BUS SERVICE – EVENT HUBS.....	48	

SERVICE-BUS SERVICE – NOTIFICATION HUBS	49
SERVICE-BUS SERVICE – QUEUES AND TOPICS	49
SERVICE-BUS SERVICE – RELAYS	50
SIGNALR SERVICE	50
SQL DATA WAREHOUSE DATABASE	51
SQL DATABASE SERVICE (BASIC, STANDARD AND PREMIUM TIERS)	51
SQL DATABASE SERVICE (WEB AND BUSINESS TIERS)	51
SQL SERVER STRETCH DATABASE	52
STORAGE SERVICE	52
STREAM ANALYTICS – API CALLS	54
STREAM ANALYTICS – JOBS	54
TRAFFIC MANAGER SERVICE	55
VIRTUAL MACHINES	55
VPN GATEWAY	56
VISUAL STUDIO APP CENTER BUILD SERVICE	57
VISUAL STUDIO APP CENTER TEST SERVICE	57
VISUAL STUDIO APP CENTER PUSH NOTIFICATION SERVICE	58
AZURE DEV OPS SERVICES – AZURE PIPELINES	58
AZURE DEVOPS TEST PLANS – LOAD TESTING SERVICE	59
AZURE DEVOPS SERVICES – USER PLANS SERVICE	59
MICROSOFT AZURE PLANS	60
AZURE ACTIVE DIRECTORY BASIC	60
AZURE ACTIVE DIRECTORY B2C	60
AZURE ACTIVE DIRECTORY PREMIUM	61
AZURE INFORMATION PROTECTION PREMIUM	61
AZURE SITE RECOVERY SERVICE – ON-PREMISES-TO-AZURE	61
AZURE SITE RECOVERY SERVICE – ON-PREMISES-TO-ON-PREMISES	62
AZURE SITE RECOVERY SERVICE – AZURE-TO-AZURE FAILOVER	62
MULTI-FACTOR AUTHENTICATION SERVICE	63
STORSIMPLE SERVICE	63
STORSIMPLE DATA MANAGER	64
OTHER ONLINE SERVICES	64
BING MAPS ENTERPRISE PLATFORM	64
BING MAPS MOBILE ASSET MANAGEMENT	65
MICROSOFT CLOUD APP SECURITY	65
MICROSOFT FLOW	65
MICROSOFT INTUNE	66
MICROSOFT KAIZALA PRO	66
MICROSOFT POWERAPPS	67
MINECRAFT: EDUCATION EDITION	67
POWER BI EMBEDDED	67
POWER BI PREMIUM	68
POWER BI PRO	68
TRANSLATOR API	68
WINDOWS DESKTOP OPERATING SYSTEM	69
APPENDIX A – SERVICE LEVEL COMMITMENT FOR VIRUS DETECTION AND BLOCKING, SPAM EFFECTIVENESS, OR FALSE POSITIVE	70
APPENDIX B - SERVICE LEVEL COMMITMENT FOR UPTIME AND EMAIL DELIVERY	71

Introduction

About this Document

This Service Level Agreement for Microsoft Online Services (this “SLA”) is a part of your Microsoft volume licensing agreement (the “Agreement”). Capitalized terms used but not defined in this SLA will have the meaning assigned to them in the Agreement. This SLA applies to the Microsoft Online Services listed herein (a “Service” or the “Services”), but does not apply to separately branded services made available with or connected to the Services or to any on-premise software that is part of any Service.

If we do not achieve and maintain the Service Levels for each Service as described in this SLA, then you may be eligible for a credit towards a portion of your monthly service fees. We will not modify the terms of your SLA during the initial term of your subscription; however, if you renew your subscription, the version of this SLA that is current at the time of renewal will apply throughout your renewal term. We will provide at least 90 days’ notice for adverse material changes to this SLA. You can review the most current version of this SLA at any time by visiting <http://www.microsoftvolumelicensing.com/SLA>.

Prior Versions of this Document

This SLA provides information on Services currently available. Earlier versions of this document are available at <http://www.microsoftvolumelicensing.com>. To find the needed version, a customer may contact its reseller or Microsoft Account Manager.

Clarifications and Summary of Changes to this Document

Below are recent additions, deletions and other changes to this SLA. Also listed below, are clarifications of Microsoft policy in response to common customer questions.

Additions	Deletions
None	None

Commented [A1]: Page 4

Service Specific Terms

[Storage Service](#): Minor updates to align with current service offerings.
[VPN Gateway](#): Minor updates to align with current service offerings.

[Table of Contents](#) / [Definitions](#)

General Terms

Definitions

“**Applicable Monthly Period**” means, for a calendar month in which a Service Credit is owed, the number of days that you are a subscriber for a Service.

“**Applicable Monthly Service Fees**” means the total fees actually paid by you for a Service that are applied to the month in which a Service Credit is owed.

“**Downtime**” is defined for each Service in the Services Specific Terms below. Except for Microsoft Azure Services, Downtime does not include Scheduled Downtime. Downtime does not include unavailability of a Service due to limitations described below and in the Services Specific Terms.

“**Error Code**” means an indication that an operation has failed, such as an HTTP status code in the 5xx range.

“**External Connectivity**” is bi-directional network traffic over supported protocols such as HTTP and HTTPS that can be sent and received from a public IP address.

“**Incident**” means (i) any single event, or (ii) any set of events, that result in Downtime.

“**Management Portal**” means the web interface, provided by Microsoft, through which customers may manage the Service.

“**Scheduled Downtime**” means periods of Downtime related to network, hardware, or Service maintenance or upgrades. We will publish notice or notify you at least five (5) days prior to the commencement of such Downtime.

“**Service Credit**” is the percentage of the Applicable Monthly Service Fees credited to you following Microsoft’s claim approval.

“**Service Level**” means the performance metric(s) set forth in this SLA that Microsoft agrees to meet in the delivery of the Services.

“**Service Resource**” means an individual resource available for use within a Service.

“**Success Code**” means an indication that an operation has succeeded, such as an HTTP status code in the 2xx range.

“**Support Window**” refers to the period of time during which a Service feature or compatibility with a separate product or service is supported.

“**User Minutes**” means the total number of minutes in a month, less all Scheduled Downtime, multiplied by the total number of users.

Terms

Claims

In order for Microsoft to consider a claim, you must submit the claim to customer support at Microsoft Corporation including all information necessary for Microsoft to validate the claim, including but not limited to: (i) a detailed description of the Incident; (ii) information regarding the time and duration of the Downtime; (iii) the number and location(s) of affected users (if applicable); and (iv) descriptions of your attempts to resolve the Incident at the time of occurrence.

For a claim related to Microsoft Azure, we must receive the claim within two months of the end of the billing month in which the Incident that is the subject of the claim occurred. For claims related to all other Services, we must receive the claim by the end of the calendar month following the month in which the Incident occurred. For example, if the Incident occurred on February 15th, we must receive the claim and all required information by March 31st.

We will evaluate all information reasonably available to us and make a good faith determination of whether a Service Credit is owed. We will use commercially reasonable efforts to process claims during the subsequent month and within forty-five (45) days of receipt. You must be in compliance with the Agreement in order to be eligible for a Service Credit. If we determine that a Service Credit is owed to you, we will apply the Service Credit to your Applicable Monthly Service Fees.

If you purchased more than one Service (not as a suite), then you may submit claims pursuant to the process described above as if each Service were covered by an individual SLA. For example, if you purchased both Exchange Online and SharePoint Online (not as part of a suite), and during the term of the subscription an Incident caused Downtime for both Services, then you could be eligible for two separate Service Credits (one for each Service), by submitting two claims under this SLA. In the event that more than one Service Level for a particular Service is not met because of the same Incident, you must choose only one Service Level under which to make a claim based on the Incident. Unless as otherwise provided in a specific SLA, only one Service Credit is permitted per Service for an Applicable Monthly Period.

Service Credits

Service Credits are your sole and exclusive remedy for any performance or availability issues for any Service under the Agreement and this SLA. You may not unilaterally offset your Applicable Monthly Service Fees for any performance or availability issues.

Service Credits apply only to fees paid for the particular Service, Service Resource, or Service tier for which a Service Level has not been met. In cases where Service Levels apply to individual Service Resources or to separate Service tiers, Service Credits apply only to fees paid for the affected

Service Resource or Service tier, as applicable. The Service Credits awarded in any billing month for a particular Service or Service Resource will not, under any circumstance, exceed your monthly service fees for that Service or Service Resource, as applicable, in the billing month.

If you purchased Services as part of a suite or other single offer, the Applicable Monthly Service Fees and Service Credit for each Service will be pro-rated.

If you purchased a Service from a reseller, you will receive a service credit directly from your reseller and the reseller will receive a Service Credit directly from us. The Service Credit will be based on the estimated retail price for the applicable Service, as determined by us in our reasonable discretion.

Limitations

- This SLA and any applicable Service Levels do not apply to any performance or availability issues:
- 1. Due to factors outside our reasonable control (for example, natural disaster, war, acts of terrorism, riots, government action, or a network or device failure external to our data centers, including at your site or between your site and our data center);
 - 2. That result from the use of services, hardware, or software not provided by us, including, but not limited to, issues resulting from inadequate bandwidth or related to third-party software or services;
 - 3. That results from failures in a single Microsoft Datacenter location, when your network connectivity is explicitly dependent on that location in a non-geo-resilient manner;
 - 4. Caused by your use of a Service after we advised you to modify your use of the Service, if you did not modify your use as advised;
 - 5. During or with respect to preview, pre-release, beta or trial versions of a Service, feature or software (as determined by us) or to purchases made using Microsoft subscription credits;
 - 6. That result from your unauthorized action or lack of action when required, or from your employees, agents, contractors, or vendors, or anyone gaining access to our network by means of your passwords or equipment, or otherwise resulting from your failure to follow appropriate security practices;
 - 7. That result from your failure to adhere to any required configurations, use supported platforms, follow any policies for acceptable use, or your use of the Service in a manner inconsistent with the features and functionality of the Service (for example, attempts to perform operations that are not supported) or inconsistent with our published guidance;
 - 8. That result from faulty input, instructions, or arguments (for example, requests to access files that do not exist);
 - 9. That result from your attempts to perform operations that exceed prescribed quotas or that resulted from our throttling of suspected abusive behavior;
 - 10. Due to your use of Service features that are outside of associated Support Windows; or
 - 11. For licenses reserved, but not paid for, at the time of the Incident.

Services purchased through Open, Open Value, and Open Value Subscription volume licensing agreements, and Services in an Office 365 Small Business Premium suite purchased in the form of a product key are not eligible for Service Credits based on service fees. For these Services, any Service Credit that you may be eligible for will be credited in the form of service time (i.e., days) as opposed to service fees, and any references to “Applicable Monthly Service Fees” is deleted and replaced by “Applicable Monthly Period.”

Service Specific Terms

Microsoft Dynamics 365

Dynamics 365 for Customer Service Enterprise; Dynamics 365 for Customer Service Professional; Dynamics 365 Customer Service Insights

Downtime: Any period of time when end users are unable to read or write any Service data for which they have appropriate permission but this does not include non-availability of Service add-on features.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Dynamics 365 Business Central

Downtime: Any period of time when end users are unable to login to their instance.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Dynamics 365 for Finance and Operations (Enterprise edition)

Additional Definitions:

“**Active Tenant**” means a tenant with an active high availability production topology in the Management Portal that (A) has been deployed to a Partner Application Service; and (B) has an active database that users can log into.

“**Partner Application Service**” means a partner application built on top of and combined with the Platform that (A) is used for processing your organization’s actual business transactions; and (B) has reserve compute and storage resources equal to or greater than one of the Scale Units your partner selected for the applicable partner application.

“**Maximum Available Minutes**” means the total accumulated minutes during a billing month in which an Active Tenant was deployed in a Partner Application Service using an active high availability production topology.

“Platform” means the Service’s client forms, SQL server reports, batched operations, and API endpoints, or the Service’s retail APIs that are used for commerce or retail purposes only.

“Scale Unit” means the increments by which compute and storage resources are added to or removed from a Partner Application Service.

“Service Infrastructure” means the authentication, computing, and storage resources that Microsoft provides in connection with the Service.

Downtime: Any period of time when end users are unable to login to their Active Tenant, due to a failure in the unexpired Platform or the Service Infrastructure as Microsoft determines from automated health monitoring and system logs. Downtime does not include Scheduled Downtime, the unavailability of Service add-on features, the inability to access the Service due to your modifications of the Service, or periods where the Scale Unit capacity is exceeded.

Monthly Uptime Percentage: The Monthly Uptime Percentage for a given Active Tenant in a calendar month is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Dynamics 365 for Retail

Additional Definitions:

“Active Tenant” means a tenant with an active high availability production topology in the Management Portal that (A) has been deployed to a Partner Application Service; and (B) has an active database that users can log into.

“Partner Application Service” means a partner application built on top of and combined with the Platform that (A) is used for processing your organization’s actual business transactions; and (B) has reserve compute and storage resources equal to or greater than one of the Scale Units your partner selected for the applicable partner application.

“Maximum Available Minutes” means the total accumulated minutes during a billing month in which an Active Tenant was deployed in a Partner Application Service using an active high availability production topology.

“Platform” means the Service’s client forms, SQL server reports, batched operations, and API endpoints, or the Service’s retail APIs that are used for commerce or retail purposes only.

“Scale Unit” means the increments by which compute and storage resources are added to or removed from a Partner Application Service.

“Service Infrastructure” means the authentication, computing, and storage resources that Microsoft provides in connection with the Service.

Downtime: Any period of time when end users are unable to access their Active Tenant, due to a failure in the unexpired Platform or the Service Infrastructure as Microsoft determines from automated health monitoring and system logs. Downtime does not include Scheduled Downtime, the unavailability of Service add-on features, the inability to access the Service due to your modifications of the Service, or periods where the Scale Unit capacity is exceeded.

Monthly Uptime Percentage: The Monthly Uptime Percentage for a given Active Tenant in a calendar month is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%

Monthly Uptime Percentage	Service Credit
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Dynamics 365 for Sales Enterprise; Dynamics 365 for Sales Professional

Downtime: Any period of time when end users are unable to read or write any Service data for which they have appropriate permission but this does not include non-availability of Service add-on features.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Dynamics 365 for Talent; Dynamics 365 for Talent: Attract; Dynamics 365 for Talent: Onboard

Additional Definitions:

“Active Tenant” means a tenant with an active high availability production topology in the Management Portal that has an active database that users can log into.

Downtime: Any period of time when end users are unable to read or write any Service data for which they have appropriate permission. Downtime does not include Scheduled Downtime.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.5%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Office 365 Services

Duet Enterprise Online

Downtime: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply when the inability to read or write any portion of a SharePoint Online site is caused by any failure of third party software, equipment, or services that are not controlled by Microsoft, or Microsoft software that is not being run by Microsoft itself as part of the Service.

Additional Terms: You will be eligible for a Service Credit for Duet Enterprise Online only when you are eligible for a Service Credit for the SharePoint Online Plan 2 User SLs that you have purchased as a prerequisite for your Duet Enterprise Online User SLs.

[Table of Contents](#) / [Definitions](#)

Exchange Online

Downtime: Any period of time when users are unable to send or receive email with Outlook Web Access. There is no Scheduled Downtime for this service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Additional Terms: See Appendix 1 – Service Level Commitment for Virus Detection and Blocking, Spam Effectiveness, or False Positive.

[Table of Contents](#) / [Definitions](#)

Exchange Online Archiving

Downtime: Any period of time when users are unable to access the email messages stored in their archive. There is no Scheduled Downtime for this service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to the Enterprise CAL suite purchased through Open Value and Open Value Subscription volume licensing agreements.

Table of Contents / Definitions

Exchange Online Protection

Downtime: Any period of time when the network is not able to receive and process email messages. There is no Scheduled Downtime for this service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to the Enterprise CAL suite purchased through Open Value and Open Value Subscription volume licensing agreements.

Additional Terms: See (i) Appendix 1 – Service Level Commitment for Virus Detection and Blocking, Spam Effectiveness, or False Positive and (ii) Appendix 2 – Service Level Commitment for Uptime and Email Delivery.

Table of Contents / Definitions

Microsoft MyAnalytics

Downtime: Any period of time when users are unable to access the MyAnalytics dashboard.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Table of Contents / Definitions

Microsoft Stream

Downtime: Any period of time when users are unable to upload, playback, delete video or edit video metadata when they have appropriate permissions and content is valid excluding unsupported scenarios¹.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Level Commitment:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: No SLA is provided for any free of charge tier of Microsoft Stream.

¹Unsupported Scenarios could include playback on unsupported devices / OS, client side network issues, and user errors.

Microsoft Teams

Downtime: Any period of time when end users are unable to see presence status, conduct instant messaging conversations, or initiate online meetings.¹

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

¹Online meeting functionality applicable only to users licensed for the Skype for Business Online Plan 2 Service.

Office 365 Business

Downtime: Any period of time when Office applications are put into reduced functionality mode due to an issue with Office 365 activation.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Office 365 Advanced Compliance

Downtime: Any period of time when Customer Lockbox component of Office 365 Advanced Compliance is put into reduced functionality mode due to an issue with Office 365.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Office 365 ProPlus

Downtime: Any period of time when Office applications are put into reduced functionality mode due to an issue with Office 365 activation.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Office Online

Downtime: Any period of time when users are unable to use the Web Applications to view and edit any Office document stored on a SharePoint Online site for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Office 365 Video

Downtime: Any period of time when users are unable to upload, view or edit videos in the video portal when they have appropriate permissions and valid content.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Level Commitment:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

OneDrive for Business

Downtime: Any period of time when users are unable to view or edit files stored on their personal OneDrive for Business storage.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Project Online

Downtime: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection with Project Web App for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Table of Contents / Definitions

SharePoint Online

Downtime: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Table of Contents / Definitions

Skype for Business Online

Downtime: Any period of time when end users are unable to see presence status, conduct instant messaging conversations, or initiate online meetings.¹

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

¹Online meeting functionality applicable only to Skype for Business Online Plan 2 Service.

Microsoft Teams – Calling Plans and Audio Conferencing

Downtime: Any period of time when end users are unable to initiate a PSTN call or unable to dial into conference audio via the PSTN.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

Where Downtime is measured in user-minutes; that is, for each month Downtime is the sum of the length (in minutes) of each incident that occurs during that month multiplied by the number of users impacted by that incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Microsoft Teams – Voice Quality

This SLA applies to any eligible call placed by any voice service user within the subscription (enabled for making any type of call VOIP or PSTN).

Additional Definitions:

“**Eligible Call**” is a Microsoft Teams placed call (within a subscription) that meets both conditions below:

- The call was placed from a Microsoft Teams Certified IP Desk phones on wired Ethernet
- Packet Loss, Jitter and Latency issues on the call were due to networks managed by Microsoft.

“**Total Calls**” is the total number of Eligible Calls

“**Poor Quality Calls**” is the total number of Eligible Calls that are classified as poor based on numerous factors that could impact call quality in the networks managed by Microsoft. While the current Poor Call classifier is built primarily on network parameters like RTT (Roundtrip Time), Packet Loss Rate, Jitter and Packet Loss-Delay Concealment Factors, it is dynamic and continually updated based on new learnings from analysis using millions of Skype, Skype for Business, and Microsoft Teams calls and evolution of Devices, Algorithms and end user ratings.

Monthly Good Call Rate: The Monthly Good Call Rate is calculated using the following formula:

$$\frac{Total\ Calls - Poor\ Quality\ Calls}{Total\ Calls} \times 100$$

Service Credit:

Monthly Good Call Rate	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Workplace Analytics

Downtime: Any period of time when users are unable to access the Workplace Analytics website.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%

Monthly Uptime Percentage	Service Credit
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Yammer Enterprise

Downtime: Any period of time greater than ten minutes when more than five percent of end users are unable to post or read messages on any portion of the Yammer network for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Microsoft Azure Services

AD Domain Services

Additional Definitions:

“**Managed Domain**” refers to an Active Directory domain that is provisioned and managed by Azure Active Directory Domain Services.
“**Maximum Available Minutes**” is the total number of minutes that a given Managed Domain has been deployed by Customer in Microsoft Azure during a billing month in a given Microsoft Azure subscription.
“**Downtime**” is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Managed Domain is unavailable. A minute is considered unavailable if all requests for domain authentication of user accounts belonging to the Managed Domain, LDAP bind to the root DSE, or DNS lookup of records, made from within the virtual network where the Managed Domain is enabled, either return an Error Code or fail to return a Success Code within 30 seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Active Directory Domain Services:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Analysis Services

Additional Definitions:

“**Server**” means any Azure Analysis Services server.
“**Maximum Available Minutes**” is the total number of minutes that a given Server has been deployed in Microsoft Azure during a billing month in a given Microsoft Azure subscription.
“**Client Operations**” is the set of all documented operations supported by Azure Analysis Services.

Downtime: is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Server is unavailable. A minute is considered unavailable for a given Server if more than 1% of all Client Operations completed during the minute return an Error Code.

Monthly Uptime Percentage: The Monthly Uptime Percentage for a given Server is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

API Management Services

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given API Management instance has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all API Management instances deployed by you in a given Microsoft Azure subscription during a billing month.

“Proxy” is the component of the API Management Service responsible for receiving API requests and forwarding them to the configured dependent API.

Downtime: The total accumulated Deployment Minutes, across all API Management instances deployed by you in a given Microsoft Azure subscription, during which the API Management Service is unavailable. A minute is considered unavailable for a given API Management instance if all continuous attempts to perform operations through the Proxy throughout the minute result in either an Error Code or do not return a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit for Basic Tier, Standard Tier and Premium Tier deployments scaled within a single region:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Credit for Premium Tier deployments scaled across two or more regions:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

App Service

Additional Definitions:

“App” is an API App, Logic App, Web App or Mobile App deployed by Customer within the App Service, excluding web apps in the Free and Shared tiers.

“Deployment Minutes” is the total number of minutes that a given App has been set to running in Microsoft Azure during a billing month. Deployment Minutes is measured from when the App was created or the Customer initiated an action that would result in running the App to the time the Customer initiated an action that would result in stopping or deleting the App.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Apps deployed by Customer in a given Microsoft Azure subscription during a billing month

Downtime: is the total accumulated Deployment Minutes, across all Apps deployed by Customer in a given Microsoft Azure subscription, during which the App is unavailable. A minute is considered unavailable for a given App when there is no connectivity between the App and Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Additional Terms: Service Credits are applicable only to fees attributable to your use of Web Apps or Mobile Apps and not to fees attributable to other types of apps available through the App Service, which are not covered by this SLA.

[Table of Contents / Definitions](#)

Application Gateway

Additional Definitions:

“Application Gateway Cloud Service” refers to a collection of one or more Application Gateway instances configured to perform HTTP load balancing services.

“Maximum Available Minutes” is the total accumulated minutes during a billing month during which an Application Gateway Cloud Service comprising two or more medium or larger Application Gateway instances has been deployed in a Microsoft Azure subscription.

Downtime: is the total accumulated Maximum Available Minutes during a billing month for a given Application Gateway Cloud Service during which the Application Gateway Cloud Service is unavailable. A given minute is considered unavailable if all attempts to connect to the Application Gateway Cloud Service throughout the minute are unsuccessful.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Application Insights (Query Availability SLA)

Additional Definitions:

“Application Insights Resource” is the container in Application Insights that collects, processes and stores the data for a single instrumentation key.

“Maximum Available Minutes” is the total number of minutes that a given Application Insights Resource has been deployed by Customer within a Microsoft Azure subscription during a billing month.

"Monthly Query Availability Percentage" for a given Application Insights Resource is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100.

“Downtime” is the total number of minutes within Maximum Available Minutes that data within an Application Insights Resource are unavailable. A minute is considered unavailable for a given Application Insights Resource during which no HTTP operations resulted in a Success Code.

Monthly Query Availability Percentage: The Monthly Query Availability Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits:

Monthly Query Availability Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Automation Service – Desired State Configuration (DSC)

Additional Definitions:
“**Deployment Minutes**” is the total number of minutes that a given Automation account has been deployed in Microsoft Azure during a billing month.
“**DSC Agent Service**” is the component of the Automation Service responsible for receiving and responding to pull, registration, and reporting requests from DSC nodes.
“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Automation accounts deployed in a given Microsoft Azure subscription during a billing month

Downtime: The total accumulated Deployment Minutes, across all Automation accounts deployed in a given Microsoft Azure subscription, during which the DSC Agent Service is unavailable. A minute is considered unavailable for a given Automation account if all continuous pull, registration, and reporting requests from DSC nodes associated with the Automation account to the DSC Agent Service throughout the minute either result in an Error Code or do not return a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Automation Service – Process Automation

Additional Definitions:
“**Delayed Jobs**” is the total number of Jobs, for a given Microsoft Azure subscription, that fail to start within thirty (30) minutes of their Planned Start Times.
“**Job**” means the execution of a Runbook.
“**Planned Start Time**” is a time at which a Job is scheduled to begin executing.
“**Runbook**” means a set of actions specified by you to execute within Microsoft Azure.
“**Total Jobs**” is the total number of Jobs scheduled for execution during a given billing month, for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Jobs} - \text{Delayed Jobs}}{\text{Total Jobs}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Azure Advanced Threat Protection

Additional Definitions:

“**Downtime**” is Any period of time when the admin is unable to access the Azure ATP portal.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Azure Bot Service

Additional Definitions:

“**Azure Bot Service Premium Channel**” is a Bot Framework channel in the premium category.

“**Bot**” is the developer’s Internet facing conversational application which is registered with and is configured to send and receive messages from the Azure Bot Service.

“**Bot Framework**” is a platform for building, connecting, testing, and deploying powerful and intelligent bots.

“**Client**” is the end user facing portion of a Bot.

“**Premium Channels API Endpoint**” is a Bot Framework REST API endpoint for Azure Bot Service Premium Channels

“**Total API Requests**” is the total number of requests made by the Bot or the Client to the Premium Channel’s API Endpoint in a Microsoft Azure subscription during a billing month.

“**Failed API Requests**” are the total number of requests within Total API Requests that return an Error Code or do not respond within 2 minutes.

“**Monthly Uptime Percentage**” is calculated as Total API Requests less Failed API Requests divided by Total API Requests multiplied by 100.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ API\ Requests - Failed\ API\ Requests}{Total\ API\ Requests} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the Azure Bot Service Premium Channels.

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Container Instances

Additional Definitions:

“**Connectivity**” is bi-directional network traffic between the Container Group and other IP addresses using TCP or UDP network protocols in which the Container Group is configured for allowed traffic.

“**Container Group**” is a collection of co-located containers that shares the same lifecycle and networking resources.

“**Maximum Available Minutes**” is the total number of minutes that a given Container Group has been deployed by Customer in a Microsoft Azure subscription during a billing month. Maximum Available Minutes is measured from Customer action that results in starting a given Container Group to the time Customer action that results in stopping or deleting a given Container Group.

“Downtime” is the total number of minutes within Maximum Available Minutes that have no Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Container Instances:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Cosmos DB

Additional Definitions:

“**Container**” is a container of data items, and a unit of scale for transactions and queries.

“**Consumed RUs**” is the sum of the Request Units consumed by all the requests which are processed by the Azure Cosmos DB Container in a given second.

“**Database Account**” is the top-level resource of the Azure Cosmos DB resource model. A Azure Cosmos DB Database Account contains one or more databases.

“**Failed Requests**” are requests within Total Requests that either return an Error Code or fail to return a Success Code within the maximum upper bounds documented in the table below.

“**Failed Read Requests**” are requests within Total Read Requests that either return an Error Code or fail to return a Success Code within the maximum upper bounds documented in the table below.

Operation	Maximum Upper Bound on Processing Latency
All Database Account configuration operations	2 Minutes
Add a new Region	60 Minutes
Manual Failover	5 Minutes
Resource Operations	5 Seconds
Media Operations	60 Seconds

“**Provisioned RUs**” is the total provisioned Request Units for a given Azure Cosmos DB Container for a given second.

“**Rate Limited Requests**” are requests that return a 429 status code from the Azure Cosmos DB Container, indicating that Consumed RUs have exceeded the Provisioned RUs for a partition in the Container for a given second.

“**Request Unit (RU)**” is a measure of throughput in Azure Cosmos DB.

“**Resource**” is a set of URI addressable entities associated with a Database Account.

“**Successful Requests**” are Total Requests minus Failed Requests.

“**Total Read Requests**” is the set of all the read requests, including Rate Limited Requests and all the Failed Read Requests, issued against Resources within a one-hour interval within a given Azure subscription during a billing month.

“**Total Requests**” is the set of all requests, including Rate Limited Requests and all Failed Requests, issued against Resources within a one-hour interval within a given Azure subscription during a billing month.

Availability SLA

“**Read Error Rate**” is the total number of Failed Read Requests divided by Total Read Requests, across all Resources in a given Azure subscription, during a given one-hour interval. If the Total Read Requests in a given one-hour interval is zero, the Read Error Rate for that interval is 0%.

“**Error Rate**” is the total number of Failed Requests divided by Total Requests, across all Resources in a given Azure subscription, during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

“**Average Error Rate**” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“**Average Read Error Rate**” for a billing month is the sum of Read Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

Monthly Availability Percentage: For the Azure Cosmos DB Service deployed via Database Accounts scoped to a single Azure region configured with any of the five Consistency Levels or Database Accounts spanning multiple regions, configured with any of the four relaxed Consistency Levels is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. The Monthly Availability Percentage is represented by the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Availability Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Monthly Read Availability Percentage: For the Azure Cosmos DB Service deployed via Database Account configured to span two or more regions is calculated by subtracting from 100% the Average Read Error Rate for a given Microsoft Azure subscription in a billing month. Monthly Read Availability Percentage is represented by the following formula:

$$100\% - \text{Average Read Error Rate}$$

Service Credit:

Monthly Read Availability Percentage	Service Credit
< 99.999%	10%
< 99%	25%

Monthly Multiple Write Locations Availability Percentage: For the Azure Cosmos DB Service deployed via Database Accounts configured to span multiple Azure regions with multiple writable locations, is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. Monthly Availability Percentage is represented by the following formula:

$$\text{Monthly Uptime \%} = 100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Multiple Write Locations Availability Percentage	Service Credit
< 99.999%	10%
< 99%	25%

Throughput SLA

“Throughput Failed Requests” are Rate-Limited Requests resulting in an Error Code, before Consumed RUs have exceeded the Provisioned RUs for a partition in the Container for a given second.

“Error Rate” is the total number of Throughput Failed Requests divided by Total Requests, across all Resources in a given Azure subscription, during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

“Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“Monthly Throughput Percentage” for the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. Monthly Throughput Percentage is represented by the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Throughput Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Consistency SLA

“K” is the number of versions of a given data item for which the reads lag behind the writes.

“T” is a given time interval.

“Consistency Level” is the setting for a particular read request that supports consistency guarantees. The following table captures the guarantees associated with the Consistency Levels. Note that Session, Bounded Staleness, Consistent Prefix and Eventual Consistency Levels are all referred to as “relaxed”.

[Table of Contents](#)

→

[Introduction](#)

→

[General Terms](#)

→

[Service Specific Terms](#)

→

[Appendices](#)

Consistency Level	Consistency Guarantees
Strong	Linearizability
Sessions	Read Your Own Write (within write region) Monotonic Read Consistent Prefix
Bounded Staleness	Read Your Own Write (within write region) Monotonic Read (within a region) Consistent Prefix Staleness Bound < K,T
Consistent Prefix	Consistent Prefix
Eventual	Eventual

“**Consistency Violation Rate**” is Successful Requests that could not be delivered when performing the consistency guarantees specified for the chosen Consistency Level divided by Total Requests, across all Resources in a given Azure subscription, during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Consistency Violation Rate for that interval is 0%.

“**Average Consistency Violation Rate**” for a billing month is the sum of Consistency Violation Rates for each hour in the billing month divided by the total number of hours in the billing month.

“**Monthly Consistency Attainment Percentage**” for the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Consistency Violation Rate for a given Microsoft Azure subscription in a billing month.

Monthly Consistency Percentage: For the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Consistency Violation Rate for a given Microsoft Azure subscription in a billing month. The Monthly Consistency Percentage is represented by the following formula:

$$100\% - \text{Average Consistency Violation Rate}$$

Service Credit:

Consistency Attainment Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Latency SLA

“**Application**” is a Azure Cosmos DB application deployed within a local Azure region with accelerated networking enabled and using the Azure Cosmos DB client SDK configured with TCP direct connectivity for a given Microsoft Azure subscription in a billing month.

“**N**” is the number of Successful Requests for a given Application performing either a data item read or data item write operations with a payload size less than or equal to 1 KB in a given hour.

“**S**” is the latency-sorted set of Successful Request response times in ascending order for a given Application performing data item read or data item write operations with a payload size less than or equal to 1 KB in a given hour.

“**Ordinal Rank**” is the 99th percentile using the nearest rank method represented by the following formula:

$$\text{Ordinal Rank} = \frac{99}{100} \times N$$

“**P99 Latency**” is the value at the Ordinal Rank of S.

“**Excessive Latency Hours**” is the total number of one-hour intervals during which Successful Requests submitted by an Application resulted in a P99 Latency greater than or equal to 10ms for data item read or 10ms for data item write operations. If the number of Successful Requests in a given one-hour interval is zero, the Excessive Latency Hours for that interval is 0.

“**Average Excessive Latency Rate**” for a billing month is the sum of Excessive Latency Hours divided by the total number of hours in the billing month.

“**Monthly P99 Latency Attainment Percentage**” for a given Azure Cosmos DB Application deployed via Database Accounts scoped to a single Azure region configured with any of the five Consistency Levels or Database Accounts spanning multiple regions, configured with any of the four relaxed Consistency Levels is calculated by subtracting from 100% the Average Excessive Latency Rate for a given Microsoft Azure subscription in a billing month. Monthly P99 Latency Attainment Percentage is represented by the following formula:

$$100\% - \text{Average Excessive Latency Rate}$$

[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Service Credit:

Monthly P99 Latency Attainment Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Database for MySQL

Additional Definitions:
“**Server**” is any given Azure Database for MySQL server.
“**Maximum Available Minutes**” is the total number of minutes for a given Server deployed by Customer in a Microsoft Azure subscription during a billing month.
“**Downtime**” is the total number of minutes within Maximum Available Minutes during which a Server is unavailable. A minute is considered unavailable if all continuous attempts by Customer to establish a connection to the Server returned an Error Code.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Database for MySQL:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Database for PostgreSQL

Additional Definitions:
“**Server**” is any given Azure Database for PostgreSQL server.
“**Maximum Available Minutes**” is the total number of minutes for a given Server deployed by Customer in a Microsoft Azure subscription during a billing month.
“**Downtime**” is the total number of minutes within Maximum Available Minutes during which a Server is unavailable. A minute is considered unavailable if all continuous attempts by Customer to establish a connection to the Server returned an Error Code.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Database for PostgreSQL:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure DDoS Protection

Additional Definitions:
“**Maximum Available Minutes**” is the total number of minutes DDoS Protection Service is enabled for a given Microsoft Azure subscription during a billing month.
“**Downtime**” is the total number of minutes within Maximum Available Minutes where protected Azure resources were not available. A minute is considered unavailable when DDoS Protection did not mitigate an attack which directly resulted in underlying Azure resources not meeting respective SLA.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure DDoS Protection:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99.95%	25%

[Table of Contents](#) / [Definitions](#)

Azure DNS

Additional Definitions:
“**DNS Zone**” refers to a deployment of the Azure DNS Service containing a DNS zone and record sets.
“**Deployment Minutes**” is the total number of minutes that a given DNS Zone has been deployed in Microsoft Azure during a billing month.
“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all DNS Zones deployed in a given Microsoft Azure subscription during a billing month.
“**Valid DNS Request**” means a DNS request to an Azure DNS Service name server associated with a DNS Zone for a matching record set within the DNS Zone.
“**Downtime**” is the total accumulated Maximum Available Minutes during which the DNS Zone is unavailable. A minute is considered unavailable for a given DNS Zone if a DNS response is not received within two seconds to a valid DNS Request, provided that the valid DNS Request is made to all name servers associated with the DNS Zone and retries are continually attempted for at least 60 consecutive seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
<100	10%
< 99.99%	25%
< 99.5%	100%

[Table of Contents](#) / [Definitions](#)

Azure Firewall

Additional Definitions:
“**Azure Firewall Service**” refers to a logical firewall instance deployed in a customer Virtual Network.
“**Maximum Available Minutes**” is the total accumulated minutes in a billing month during which the Azure Firewall Service has been deployed in a Microsoft Azure subscription.
“**Downtime**” is the total accumulated Maximum Available Minutes in a billing month for a given Azure Firewall Service during which the Azure Firewall Service is unavailable. A given minute is considered unavailable if all attempts to connect to the Azure Firewall Service throughout the minute are unsuccessful.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Functions

For Function Apps running on App Service Plans we guarantee that the associated Functions compute will be available 99.95% of the time. No SLA is provided for Functions Apps running under Consumption Plans.

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Function App is available to be triggered during a billing month. Deployment Minutes are measured based on the total time that the service is available to trigger a function execution and not based on the potential number of Function executions that might be triggered during a given month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Function Apps deployed by Customer in a given Microsoft Azure subscription during a billing month.

“Function App” is an individual Function deployed on an App Service Plan with an associated trigger.

“Downtime” The total accumulated Deployment Minutes, across the Function App deployed by a customer in a given Microsoft Azure subscription, during which the Function App is unavailable to be triggered. A minute is considered unavailable for a given Function App when there is no connectivity between the App Service Plan on which the Function App is hosted and Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Lab Services

Additional Definitions:

“Lab Virtual Machine(s)” is defined as any virtual machine provisioned inside a lab in Azure Lab Services.

“Lab Virtual Machine Connectivity” is bi-directional network traffic between the Lab Virtual Machine and other IP addresses using TCP or UDP network protocols in which the Lab Virtual Machine is configured for allowed traffic. The IP addresses can be IP addresses within the same virtual network as the Lab Virtual Machine or public, routable IP addresses.

“Minutes in the Month” is the total number of minutes in a given month.

“Downtime” is the total accumulated minutes that are part of Minutes in the Month that have no Lab Virtual Machine Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Minutes\ in\ the\ Month - Downtime}{Minutes\ in\ the\ Month} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Lab Services:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Load Balancer

Additional Definitions:

“Load Balanced Endpoint” is an IP address and associated IP transport port definition.

“Healthy Virtual Machine” is a Virtual Machine which returns a Success Code for the health probe sent by the Azure Standard Load Balancer. The Virtual Machine must have Network Security Group rules permitting communication with the load balanced port.

“Connectivity” is bi-directional network traffic over supported IP transport protocols that can be sent and received from any IP address configured to allow traffic.

“Maximum Available Minutes” is the total number of minutes that a given Azure Standard Load Balancer (serving two or more Healthy Virtual Machines) has been deployed by Customer in a Microsoft Azure subscription during a billing month.

“**Downtime**” is the total number of minutes within Maximum Available Minutes during which the given Azure Standard Load Balancer is unavailable. A minute is considered unavailable if all Healthy Virtual Machines have no Connectivity through the Load Balanced Endpoint. Downtime does not include minutes resulting from SNAT port exhaustions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Load Balancer:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99.9%	25%

Service Level Exceptions: The Basic Load Balancer is not covered by this SLA.

[Table of Contents / Definitions](#)

Azure Maps API

Additional Definitions:

“**Total Transaction Attempts**” is the total number of authenticated API requests made by Customer for a given Azure Map API during a billing month in a given Microsoft Azure subscription. Total Transaction Attempts do not include API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received.

“**Failed Transactions**” is the set of all requests within Total Transaction Attempts that result in an Error Code or otherwise do not return a Success Code within 60 seconds after receipt by the Service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Maps API:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Azure Monitor

Additional Definitions:

“**Action Group**” is a collection of actions deployed by Customer in a given Microsoft Azure subscription which defines preferred notification delivery methods.

“**Deployment Minutes**” is the total number of minutes that a given Action Group has been deployed by Customer in Microsoft Azure subscription during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Action Groups deployed by Customer in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment Minutes, across all Action Groups, during which the Action Group is unavailable. A minute is considered unavailable for a given Action Group if all continuous attempts to send alerts or perform registration management operations with respect to the Action Group throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Also refer to Log Analytics and Application Insights.

[Table of Contents](#) / [Definitions](#)

Azure Monitor Alerts

Additional Definitions:

“Alert Rule” is a collection of signal criteria used to generate alerts using monitoring event data already available to Alert Service for analysis.
“Maximum Available Minutes” is the total number of minutes which Alert Rule(s) are deployed by Customer in a given Microsoft Azure subscription during a billing month.
“Downtime” is the total number of minutes within Maximum Available Minutes during which the Alert Rule is unavailable. A minute is considered unavailable for a given Alert Rule if all continuous attempts to analyze telemetry signals for resources defined within the Alert Rule throughout the minute either return an Error Code or do not result in a Success Code within five minutes from scheduled Alert Rule start time.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Monitor Alerts:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Monitor Notification Delivery

Additional Definitions:

“Action Group” is a collection of actions that defines preferred notification delivery methods.
“Maximum Available Minutes” is the total number of minutes which Action Group(s) are deployed by Customer in a given Microsoft Azure subscription during a billing month.
“Downtime” is the total number of minutes within Maximum Available Minutes during which the Action Group is unavailable. A minute is considered unavailable for a given Action Group if all continuous attempts to send alerts or perform registration management operations with respect to the Action Group throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Azure Monitor Notification Delivery:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Security Center

Additional Definitions:

“Protected Node” is a Microsoft Azure resource, counted as a node for billing purposes that is configured for the Azure Security Center Standard Tier
“Security Monitoring” is the assessment of a Protected Node resulting in potential findings such as security health status, recommendations, and security alerts, exposed in Azure Security Center.

“Maximum Available Minutes” is the total number of minutes during a billing month that a given Protected Node has been deployed and configured for Security Monitoring.

“Downtime” is the total accumulated minutes during a billing month for which Security Monitoring information of a given Protected Node is unavailable. A minute is considered unavailable for a given Protected Node if all continuous attempts to retrieve Security Monitoring information throughout the minute result in either an Error Code or do not return a Success Code within two minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Virtual WAN

Additional Definitions:

“Maximum Available Minutes” is the total accumulated minutes in a billing month during which a given Azure Virtual WAN has been deployed in a Microsoft Azure subscription.

“Downtime” is the total accumulated Maximum Available Minutes during which an Azure Virtual WAN is unavailable. A given minute is considered unavailable if all attempts to connect to the Azure Virtual WAN throughout the minute are unsuccessful.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Batch Service

Additional Definitions:

“Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“Error Rate” is the total number of Failed Requests divided by Total Requests during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

“Excluded Requests” are requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

“Failed Requests” is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or fail to return a Success Code within 5 seconds.

“Total Requests” is the total number of authenticated REST API requests, other than Excluded Requests, to perform operations against Batch accounts attempted within a one-hour interval within a given Azure subscription during a billing month.

Monthly Uptime Percentage: for the Batch Service is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. The “Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month. Monthly Uptime Percentage is represented by the following formula:

$$\text{Monthly Uptime \%} = 100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Backup Service

Additional Definitions:

“**Backup**” or “**Back Up**” is the process of copying computer data from a registered server to a Backup Vault.

“**Backup Agent**” refers to the software installed on a registered server that enables the registered server to Back Up or Restore one or more Protected Items.

“**Backup Vault**” refers to a container in which you may register one or more Protected Items for Backup.

“**Deployment Minutes**” is the total number of minutes during which a Protected Item has been scheduled for Backup to a Backup Vault.

“**Failure**” means that either the Backup Agent or the Service fails to fully complete a properly configured Backup or Recovery operation due to unavailability of the Backup Service.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Protected Items for a given Microsoft Azure subscription during a billing month.

“**Protected Item**” refers to a collection of data, such as a volume, database, or virtual machine that has been scheduled for Backup to the Backup Service such that it is enumerated as a Protected Item in the Protected Items tab in the Recovery Services section of the Management Portal.

“**Recovery**” or “**Restore**” is the process of restoring computer data from a Backup Vault to a registered server.

Downtime: The total accumulated Deployment Minutes across all Protected Items scheduled for Backup by you in a given Microsoft Azure subscription during which the Backup Service is unavailable for the Protected Item. The Backup Service is considered unavailable for a given Protected Item from the first Failure to Back Up or Restore the Protected Item until the initiation of a successful Backup or Recovery of a Protected Item, provided that retries are continually attempted no less frequently than once every thirty minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

BizTalk Services

Additional Definitions:

“**BizTalk Service Environment**” refers to a deployment of the BizTalk Services created by you, as represented in the Management Portal, to which you may send runtime message requests.

“**Deployment Minutes**” is the total number of minutes that a given BizTalk Service Environment has been deployed in Microsoft Azure during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all BizTalk Service Environments deployed by you in a given Microsoft Azure subscription during a billing month.

“**Monitoring Storage Account**” refers to the Azure Storage account used by the BizTalk Services to store monitoring information related to the execution of the BizTalk Services.

Downtime: The total accumulated Deployment Minutes, across all BizTalk Service Environments deployed by you in a given Microsoft Azure subscription, during which the BizTalk Service Environment is unavailable. A minute is considered unavailable for a given BizTalk Service Environment when there is no connectivity between your BizTalk Service Environment and Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Basic, Standard, and Premium tiers of the BizTalk Services. The Developer tier of the Microsoft Azure BizTalk Services is not covered by this SLA.

Additional Terms: When submitting a claim, you must ensure that complete monitoring data is maintained within the Monitoring Storage Account and is made available to Microsoft.

[Table of Contents / Definitions](#)

Cache Services

Additional Definitions:

“**Cache**” refers to a deployment of the Cache Service created by you, such that its Cache Endpoints are enumerated in the Cache tab in the Management Portal.

“**Cache Endpoints**” refers to endpoints through which a Cache may be accessed.

“**Deployment Minutes**” is the total number of minutes that a given Cache has been deployed in Microsoft Azure during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Caches deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes, across all Caches deployed by you in a given Microsoft Azure subscription, during which the Cache is unavailable. A minute is considered unavailable for a given Cache when there is no connectivity throughout the minute between one or more Cache Endpoints associated with the Cache and Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Cache Service, which includes the Azure Managed Cache Service or the Standard tier of the Azure Redis Cache Service. The Basic tier of the Azure Redis Cache Service is not covered by this SLA.

[Table of Contents / Definitions](#)

CDN Service

Downtime: To assess Downtime, Microsoft will review data from any commercially reasonable independent measurement system used by you.

You must select a set of agents from the measurement system’s list of standard agents that are generally available and represent at least five geographically diverse locations in major worldwide metropolitan areas (excluding PR of China).

Measurement System tests (frequency of at least one test per hour per agent) will be configured to perform one HTTP GET operation according to the model below:

1. A test file will be placed on your origin (e.g., Azure Storage account).
2. The GET operation will retrieve the file through the CDN Service, by requesting the object from the appropriate Microsoft Azure domain name hostname.
3. The test file will meet the following criteria:

- i. The test object will allow caching by including explicit “Cache-control: public” headers, or lack of “Cache-Control: private” header.
- ii. The test object will be a file at least 50KB in size and no larger than 1MB.
- iii. Raw data will be trimmed to eliminate any measurements that came from an agent experiencing technical problems during the measurement period.

Monthly Uptime Percentage: The percentage of HTTP transactions in which the CDN responds to client requests and delivers the requested content without error. Monthly Uptime Percentage of the CDN Service is calculated as the number of times the object was delivered successfully divided by the total number of requests (after removing erroneous data).

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99.5%	25%

[Table of Contents](#) / [Definitions](#)

Cloud Services

Additional Definitions:

“**Cloud Services**” refers to a set of compute resources utilized for Web and Worker Roles.

“**Role Instance Connectivity**” is bi-directional network traffic between the role instance and other IP addresses using TCP or UDP network protocols in which the role instance is configured for allowed traffic. The IP addresses can be IP addresses in the same Cloud Service as the virtual machine, IP addresses within the same virtual network as the virtual machine or public, routable IP addresses.

“**Maximum Available Minutes**” is the total accumulated minutes during a billing month for all Internet facing roles that have two or more instances deployed in different Update Domains. Maximum Available Minutes is measured from when the Tenant has been deployed and its associated roles have been started resultant from action initiated by Customer to the time Customer has initiated an action that would result in stopping or deleting the Tenant.

“**Tenant**” represents one or more roles each consisting of one or more role instances that are deployed in a single package.

“**Update Domain**” refers to a set of Microsoft Azure instances to which platform updates are concurrently applied.

“**Web Role**” is a Cloud Services component run in the Azure execution environment that is customized for web application programming as supported by IIS and ASP.NET.

“**Worker Role**” is a Cloud Services component run in the Azure execution environment that is useful for generalized development, and may perform background processing for a Web Role.

Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no Role Instance Connectivity.

Monthly Uptime Percentage: Monthly Uptime Percentage is represented by the following formula:

$$\text{Monthly Uptime \%} = \frac{(\text{Maximum Available Minutes} - \text{Downtime})}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Container Registry

Additional Definitions:

“**Managed Registry**” is any instance of Basic, Standard or Premium Container Registry.

“**Registry Endpoint**” is the host name from which a given Managed Registry is accessed by clients to perform Container Registry related operations.

“**Registry Transactions**” is the set of transaction requests sent from the client to the Registry Endpoint.

“**Maximum Available Minutes**” is the total number of minutes that a given Managed Container Registry has been deployed by Customer in a Microsoft subscription during a billing month.

“**Downtime**” is the total number of minutes within Maximum Available Minutes during which Managed Registry is unavailable. A minute is considered unavailable if all continuous attempts to send Registry Transactions receive an Error Code or do not respond within the Maximum Processing Time outlined in the table below.

Transaction Types	Maximum Processing Time
List (Repository, Manifests, Tags)	8 Minutes
Others	1 Minute

“Monthly Uptime Percentage” for Managed Container Registry is calculated using the following formula:

$$\text{Monthly Uptime \%} = \frac{(\text{Maximum Available Minutes} - \text{Downtime})}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Data Catalog

Additional Definitions:

“Deployment Minutes” is the total number of minutes for which a Data Catalog has been purchased during a billing month.

“Entries” means any catalog object registration in the Data Catalog (such as a table, view, measure, cluster or report).

“Maximum Available Minutes” is the sum of all Deployment Minutes for the Data Catalog associated with a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment minutes, during which the Data Catalog is unavailable. A minute is considered unavailable for a given Data Catalog if all attempts by administrators to add or remove users to the Data Catalog or all attempts by users to execute API calls to the Data Catalog for registering, searching, or deleting Entries either result in an Error Code or do not return a response within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Data Factory – Activity Runs

Additional Definitions:

“Activity Run” means the execution or attempted execution of an activity

“Delayed Activity Runs” is the total number of attempted Activity Runs in which an activity fails to begin executing within four (4) minutes after the time at which it is scheduled for execution and all dependencies that are prerequisite to execution have been satisfied.

“Total Activity Runs” is the total number of Activity Runs attempted during in a billing month for a given Microsoft Azure Subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Activity Runs} - \text{Delayed Activity Runs}}{\text{Total Activity Runs}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Data Factory – API Calls

Additional Definitions:

“**Excluded Requests**” is the set of requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

“**Failed Requests**” is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or otherwise fail to return a Success Code within two minutes.

“**Resources**” means integration runtimes (including Azure, SSIS and self-hosted Integration Runtimes), triggers, pipelines, data sets, and linked services created within a Data Factory.

“**Total Requests**” is the set of all requests, other than Excluded Requests, to perform operations against Resources during a billing month for a given Microsoft Azure subscription.

Monthly Uptime Percentage: of the API calls made to the Data Factory Services is calculated as Total Requests less Failed Requests divided by Total Requests in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

Monthly Uptime % =

(Total Requests – Failed Requests)

Total Requests

Service Credit:

The following Service Credits are applicable to Customer’s use of API calls within the Data Factory Service

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Data Lake Analytics

Additional Definitions:

“**Total Operations**” is the total number of authenticated operations attempted within a one-hour interval across all Data Lake Analytics accounts in a given Azure subscription during a billing month.

“**Failed Operations**” is the set of all operations within Total Operations that either return an Error Code or fail to return a Success Code within 5 minutes for account creation and deletion and 25 seconds for all other operations with an additional 2 seconds per MB for operations with payload.

“**Error Rate**” is the total number of Failed Operations divided by Total Operations during a given one-hour interval. If the Total Operations in a one-hour interval is zero, the Error Rate for that interval is 0%.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

100% - Average Error Rate

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Data Lake Store

Additional Definitions:

“**Total Operations**” is the total number of authenticated operations attempted within a one-hour interval across all Data Lake Store accounts in a given Azure subscription during a billing month.

“**Failed Operations**” is the set of all operations within Total Operations that either return an Error Code or fail to return a Success Code within 5 minutes for account creation and deletion, 2 seconds per file for operations on multiple files, 2 seconds per MB for data transfer operations, and 2 seconds for all other operations.

“**Error Rate**” is the total number of Failed Operations divided by Total Operations during a given one-hour interval. If the Total Operations in a one-hour interval is zero, the Error Rate for that interval is 0%.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

100% - Average Error Rate

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Event Grid

Additional Definitions:

“Maximum Available Minutes” is the total number of minutes that an Event Grid has been deployed by Customer in a Microsoft Azure subscription during a billing month.

“Downtime” is the total number of minutes within Maximum Available Minutes across all Event Grids deployed by Customer in a given Microsoft Azure subscription during which Event Grid is unavailable. A minute is considered unavailable for a given Event Grid if all requests to publish a message either return an Error Code or do not result in a Success Code within one minute.

“Monthly Uptime Percentage”: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
<99.99%	10%
<99%	25%

[Table of Contents](#) / [Definitions](#)

ExpressRoute

Additional Definitions:

“Dedicated Circuit” means a logical representation of connectivity offered through the ExpressRoute Service between your premises and Microsoft Azure through an ExpressRoute connectivity provider, where such connectivity does not traverse the public Internet.

“Maximum Available Minutes” is the total number of minutes that a given Dedicated Circuit is linked to one or more Virtual Networks in Microsoft Azure during a billing month in a given Microsoft Azure subscription.

“Virtual Network” refers to a virtual private network that includes a collection of user-defined IP addresses and subnets that form a network boundary within Microsoft Azure.

“VPN Gateway” refers to a gateway that facilitates cross-premises connectivity between a Virtual Network and a customer on-premises network.

“Downtime” is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which the Dedicated Circuit is unavailable. A minute is considered unavailable for a given Dedicated Circuit if all attempts by you within the minute to establish IP-level connectivity to the VPN Gateway associated with the Virtual Network fail for longer than thirty seconds.

“Monthly Uptime Percentage” is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit The following Service Levels and Service Credits are applicable to Customer’s use of each Dedicated Circuit within the ExpressRoute Service.

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Function App on Consumption Plan

Additional Definitions:

"Function App" is a collection of one or more functions deployed with an associated trigger.

"Total Triggered Executions" is the total number of all Function App executions triggered by Customer in a given Microsoft Azure subscription during a billing month.

"Unavailable Executions" is the total number of executions within Total Triggered Executions which failed to run. An execution failed to run when the given Function App history log did not capture any output five (5) minutes after the trigger is successfully fired.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Triggered Executions} - \text{Unavailable Executions}}{\text{Total Triggered Executions}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Function App on Consumption Plan:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Function App on Service Plan

Additional Definitions:

"Deployment Minutes" is the total number of minutes that a given Function App is available to be triggered during a billing month. Deployment Minutes are measured based on the total time that the service is available to trigger a function execution and not based on the potential number of function executions that might be triggered during a given month.

"Maximum Available Minutes" is the sum of all Deployment Minutes for a given Function App deployed by Customer in a given Microsoft Azure subscription during a billing month.

"Downtime" is the total number of minutes within Maximum Available Minutes, during which the Function App is unavailable to be triggered. A minute is considered unavailable for a given Function App when there is no connectivity between the App Service Plan on which the Function App is hosted and Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Levels and Service Credits are applicable to Customer’s use of Function App on Service Plan:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

HDInsight

Additional Definitions:

"Cluster Internet Gateway" means a set of virtual machines within an HDInsight Cluster that proxy all connectivity requests to the Cluster.

"Deployment Minutes" is the total number of minutes that a given HDInsight Cluster has been deployed in Microsoft Azure.

"HDInsight Cluster" or **"Cluster"** means a collection of virtual machines running a single instance of the HDInsight Service.

"Maximum Available Minutes" is the sum of all Deployment Minutes across all Clusters deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes when the HDInsight Service is unavailable. A minute is considered unavailable for a given Cluster if all continual attempts within the minute to establish a connection to the Cluster Internet Gateway fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

HockeyApp

Additional Definitions:

“**HockeyApp Dashboard**” means the web interface provided to developers to view and manage applications using the HockeyApp Service.

“**Maximum Available Minutes**” is the total number of minutes in a billing month.

Downtime: is the total accumulated minutes in a billing month during which the HockeyApp Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the HockeyApp Dashboard or to the HockeyApp API throughout the minute either result in an Error Code or do not return a response within one minute. For purposes of the HockeyApp API, HTTP response codes 408, 429, 500, 503, and 511 are not considered Error Codes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

IoT Central

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes that a given IoT Central application has been deployed in given Microsoft Azure Subscription during a billing month.

“**Device Identity Operations**” refers to create, read, update, and delete operations performed on the devices of an IoT Central application.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all IoT Central applications deployed in a given Microsoft Azure subscription during a billing month.

“**Message**” refers to any content sent by a deployed IoT Central application to a device registered to the IoT Central application or received by the IoT Central application from a registered device.

Downtime: The total accumulated Maximum Available Minutes during which IoT Central is unavailable. A minute is considered unavailable for a given IoT Central application if all continuous attempts to send or receive Messages or perform Device Identity Operations on the IoT Central application throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

IoT hub

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes that a given IoT hub has been deployed in Microsoft Azure during a billing month.

“**Device Identity Operations**” refers to create, read, update, and delete operations performed on the device identity registry of an IoT hub.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all IoT hubs deployed in a given Microsoft Azure subscription during a billing month.

“**Message**” refers to any content sent by a deployed IoT hub to a device registered to the IoT hub or received by the IoT hub from a registered device, using any protocol supported by the Service.

Downtime: The total accumulated Deployment Minutes, across all IoT hubs deployed in a given Microsoft Azure subscription, during which the IoT hub is unavailable. A minute is considered unavailable for a given IoT hub if all continuous attempts to send or receive Messages or perform Device Identity Operations on the IoT hub throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Key Vault

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes that a given key vault has been deployed in Microsoft Azure during a billing month.

“**Excluded Transactions**” are transactions for creating, updating, or deleting key vaults, keys, or secrets.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Key Vaults deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment Minutes, across all key vaults deployed by Customer in a given Microsoft Azure subscription, during which the key vault is unavailable. A minute is considered unavailable for a given key vault if all continuous attempts to perform transactions, other than Excluded Transactions, on the key vault throughout the minute either return an Error Code or do not result in a Success Code within 5 seconds from Microsoft’s receipt of the request.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Log Analytics (Query Availability SLA)

Additional Definitions:

“**Maximum Available Minutes**” is the total number of minutes that a given Log Analytics Workspace has been deployed by Customer in a Microsoft Azure subscription during a billing month.

“**Downtime**” is the total number of minutes within Maximum Available Minutes that data in a Log Analytics Workspace are unavailable. A minute is considered unavailable for a given Log Analytics Workspace during which no HTTP operations resulted in a Success Code.

"Monthly Query Availability Percentage" for a given Log Analytics Workspace calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100.

Monthly Query Availability Percentage: The Monthly Query Availability Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Query Availability Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Logic Apps

Additional Definitions:

"Deployment Minutes" is the total number of minutes that a given Logic App has been set to running in Microsoft Azure during a billing month. Deployment Minutes is measured from when the Logic App was created or Customer initiated an action that would result in running the Logic App to the time Customer initiated an action that would result in stopping or deleting the Logic App.

"Maximum Available Minutes" is the sum of all Deployment Minutes across all Logic Apps deployed by Customer in a given Microsoft Azure subscription during a billing month.

"Downtime" The total accumulated Deployment Minutes, across all Logic Apps deployed by Customer in a given Microsoft Azure subscription, during which the Logic App is unavailable. A minute is considered unavailable for a given Logic App when there is no connectivity between the Logic App and Microsoft's Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Azure Machine Learning Studio – Batch Execution Service (BES) and Management APIs Service

Additional Definitions:

"Failed Transactions" is the set of all requests within Total Transaction Attempts that return an Error Code.

"Total Transaction Attempts" is the total number of authenticated REST BES and Management API requests by you during a billing month for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: Service Levels and Service Credits are applicable to your use of the Azure Machine Learning Studio BES and Management API Service. The Free Azure Machine Learning Studio tier is not covered by this SLA.

Azure Machine Learning Studio – Request Response Service (RRS)

Additional Definitions:
“**Failed Transactions**” is the set of all requests within Total Transaction Attempts that return an Error Code.
“**Total Transaction Attempts**” is the total number of authenticated REST RRS and Management API requests by you during a billing month for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Service Level Exceptions: Service Levels and Service Credits are applicable to your use of the Azure Machine Learning Studio RRS and Management API Service. The Free Azure Machine Learning Studio tier is not covered by this SLA.

Media Services – Content Protection Service

Additional Definitions:
“**Failed Transactions**” are all Valid Key Requests included in Total Transaction Attempts that result in an Error Code or otherwise do not return a Success Code within 30 seconds after receipt by the Content Protection Service.
“**Total Transaction Attempts**” are all Valid Key Requests made by you during a billing month for a given Azure subscription.
“**Valid Key Requests**” are all requests made to the Content Protection Service for existing content keys in a Customer’s Media Service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Media Services – Encoding Service

Additional Definitions:
“**Encoding**” means the processing of media files per subscription as configured in the Media Services Tasks.
“**Failed Transactions**” is the set of all requests within Total Transaction Attempts that do not return a Success Code within 30 seconds from Microsoft’s receipt of the request.
“**Media Service**” means an Azure Media Services account, created in the Management Portal, associated with your Microsoft Azure subscription. Each Microsoft Azure subscription may have more than one associated Media Service.
“**Media Services Task**” means an individual operation of media processing work as configured by you. Media processing operations involve encoding and converting media files.
“**Total Transaction Attempts**” is the total number of authenticated REST API requests with respect to a Media Service made by you during a billing month for a subscription. Total Transaction Attempts does not include REST API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Media Services – Media Indexer Service

Additional Definitions:

“**Failed Transactions**” is the set of Indexer Tasks within Total Transaction Attempts that either, a) do not complete within a time period that is 3 times the duration of the input file, or b) do not start processing within 5 minutes of the time that a Media Reserved Unit becomes available for use by the Indexer Task.

“**Indexer Task**” means a Media Services Task that is configured to extract the speech content from an MP3 input file with a minimum five-minute duration.

“**Media Reserved Unit**” means reserved units purchased by the customer in an Azure Media Services account.

“**Total Transaction Attempts**” is the total number of Indexer Tasks attempted to be executed using an available Media Reserved Unit by Customer during a billing month for a subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Media Services – Live Channels

Additional Definitions:

“**Channel**” means an end point within a Media Service that is configured to receive media data.

“**Deployment Minutes**” is the total number of minutes that a given Channel has been purchased and allocated to a Media Service and is in a running state during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Channels purchased and allocated to a Media Service during a billing month.

“**Media Service**” means an Azure Media Services account, created in the Management Portal, associated with your Microsoft Azure subscription. Each Microsoft Azure subscription may have more than one associated Media Service.

Downtime: The total accumulated Deployment Minutes when the Live Channels Service is unavailable. A minute is considered unavailable for a given Channel if the Channel has no External Connectivity during the minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%

Monthly Uptime Percentage	Service Credit
< 99%	25%

[Table of Contents / Definitions](#)

Media Services – Streaming Service

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Streaming Unit has been purchased and allocated to a Media Service during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Streaming Units purchased and allocated to a Media Service during a billing month.

“Media Service” means an Azure Media Services account, created in the Management Portal, associated with your Microsoft Azure subscription. Each Microsoft Azure subscription may have more than one associated Media Service.

“Media Service Request” means a request issued to your Media Service.

“Streaming Unit” means a unit of reserved egress capacity purchased by you for a Media Service.

“Valid Media Services Requests” are all qualifying Media Service Requests for existing media content in a customer’s Azure Storage account associated with its Media Service when at least one Streaming Unit has been purchased and allocated to that Media Service. Valid Media Services Requests do not include Media Service Requests for which total throughput exceeds 80% of the Allocated Bandwidth.

Downtime: The total accumulated Deployment Minutes when the Streaming Service is unavailable. A minute is considered unavailable for a given Streaming Unit if all continuous Valid Media Service Requests made to the Streaming Unit throughout the minute result in an Error Code.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Media Services – Video Indexer Service

Additional Definitions:

“Failed Transactions” is the set of all requests within Total Transaction Attempts that return an Error Code, or do not send a response within 360 seconds from the completion of client sending the request.

“Total Transaction Attempts” is the total number of authenticated Video Indexer API requests made by Customer during a billing month for a subscription. Total Transaction Attempts do not include Video Indexer API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received, or Upload POST requests that send the file as byte array content.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Microsoft Cognitive Services

Additional Definitions:

[Table of Contents](#) →
 [Introduction](#) →
 [General Terms](#) →
 [Service Specific Terms](#) →
 [Appendices](#)

“**Total Transaction Attempts**” is the total number of authenticated API requests by Customer during a billing month for a given Cognitive Service API. Total Transaction Attempts do not include API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received.

“**Failed Transactions**” is the set of all requests to the Cognitive Service API within Total Transaction Attempts that return an Error Code . Failed Transaction Attempts do not include API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received.

“**Monthly Uptime Percentage**” for each API Service is calculated as Total Transaction Attempts less Failed Transactions divided by Total Transaction Attempts in a billing month for a given API subscription. Monthly Uptime Percentage is represented by the following formula:

Monthly Uptime % = (Total Transaction Attempts - Failed Transactions) / Total Transaction Attempts * 100

$$\text{Monthly Uptime \%} = \frac{(\text{Total Transaction Attempts} - \text{Failed Transactions})}{\text{Total Transaction Attempts}} \times 100$$

Service Credit

The following Service Levels and Service Credits are applicable to Cognitive Services APIs:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: No SLA is provided to free tier or offerings in preview.

[Table of Contents](#) / [Definitions](#)

Microsoft Genomics

Additional Definitions:

“**Maximum Available Minutes**” is the total accumulated minutes for all Microsoft Genomics accounts created by Customer and active during a billing month for a given Microsoft Azure Subscription.

“**Downtime**” is the total number of minutes within Maximum Available Minutes during which Microsoft Genomics is unavailable. A minute is considered unavailable if all continuous attempts to send authenticated Genomics service REST API requests throughout the minute either return an Error Code or do not respond with an acknowledgement within the minute.

“**Monthly Uptime Percentage**” for Microsoft Genomics is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Mobile Engagement

Additional Definitions:

“**Average Error Rate**” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“**Error Rate**” is the total number of Failed Requests divided by Total Requests during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

“**Excluded Requests**” is the set of REST API requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

“**Failed Requests**” is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or fail to return a Success Code within 30 seconds.

“**Mobile Engagement Application**” is an Azure Mobile Engagement service instance.

“**Total Requests**” is the total number of authenticated REST API requests, other than Excluded Requests, made to Mobile Engagement Applications within a given Azure subscription during a billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

100% – Average Error Rate

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

The Free Mobile Engagement tier is not covered by this SLA.

Table of Contents / Definitions

Mobile Services

Additional Definitions:

“Failed Transactions” include any API calls included in Total Transaction Attempts that result in either an Error Code or do not return a Success Code.

“Total Transaction Attempts” are the total accumulated API calls made to the Azure Mobile Services during a billing month for a given Microsoft Azure subscription for which the Azure Mobile Services are running.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Standard and Premium Mobile Services tiers. The Free Mobile Services tier is not covered by this SLA.

Table of Contents / Definitions

Network Watcher

Additional Definitions:

“Network Diagnostic Tools” is a collection of network diagnostic and topology tools.

“Maximum Diagnostic Checks” is the total number of diagnostic actions performed by the Network Diagnostic Tool as configured by Customer in a billing month for a given Microsoft Azure subscription.

“Failed Diagnostic Checks” is the total number of diagnostic actions within Maximum Diagnostic Checks that returns an Error Code or does not return a response within the Maximum Processing Time documented in the table below.

Diagnostic Tool	Maximum Processing Time
IPFlow Verify	2 minutes
NextHop	
Packet Capture	
Security Group View	
Topology	
VPN Troubleshoot	10 minutes

“Monthly Uptime Percentage” is calculated by using the following formula:

$$\frac{\text{Maximum Diagnostic Checks} - \text{Failed Diagnostic Checks}}{\text{Maximum Diagnostic Checks}} \times 100$$

Service Levels:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents → Introduction → General Terms → Service Specific Terms → Appendices

RemoteApp

Additional Definitions:

“Application” means a software application that is configured for streaming to a device using the RemoteApp Service.

“Maximum Available Minutes” is the sum of all User Application Minutes across all Users granted access to one or more Applications in a given Azure subscription during a billing month.

“User” means a specific user account that is able to stream an Application using the RemoteApp Service, as enumerated in the Management Portal.

“User Application Minutes” is the total number of minutes in a billing month during which you have granted a User access to an Application.

Downtime: The total accumulated User Minutes during which the RemoteApp Service is unavailable. A minute is considered unavailable for a given User when the User is unable to establish connectivity to an Application.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the RemoteApp Service. The RemoteApp free trial is not covered by this SLA.

SAP HANA on Azure

Additional Definitions:

“Announced Single Instance Maintenance” means periods of Downtime related to network, hardware, or Service maintenance or upgrades impacting Single Instances. We will publish notice or notify you at least five (5) days prior to the commencement of such Downtime.

“High Availability Pair” refers to two or more identical SAP HANA on Azure large instances deployed in the same region and configured by the customer for system replication at the application layer. Customer must declare the members of a High Availability Pair to Microsoft during the architecture design process.

“SAP HANA on Azure Connectivity” is bi-directional network traffic between the SAP HANA on Azure large instance and other IP addresses using TCP or UDP network protocols in which the instance is configured for allowed traffic. The IP addresses must be IP addresses on the Virtual Network of the associated Azure subscription.

“Single Instance” is defined as any single Microsoft SAP HANA on Azure Large Instance machine that is not deployed in an High Availability Pair.

Monthly Uptime Calculation and Service Levels for SAP HANA on Azure High Availability Pair

“Maximum Available Minutes” is the total accumulated minutes during a billing month for all SAP HANA on Azure instances deployed in the same High Availability Pair. Maximum Available Minutes is measured from when two or more instances in the same High Availability Pair have both been started resultant from an action initiated by Customer to the time Customer has initiated an action that would result in stopping the instances.

“Downtime” is the total accumulated minutes that are part of Maximum Available Minutes that have no SAP HANA on Azure Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage for SAP HANA on Azure High Availability Pair is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit for SAP HANA on Azure High Availability Pair:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99.9%	25%

Monthly Uptime Calculation and Service Levels for SAP HANA on Azure Single Instance

“**Maximum Available Minutes**” is the total accumulated minutes for all SAP HANA on Azure Single Instances deployed by Customer during a billing month for a given Microsoft Azure subscription.

“**Downtime**” is the total accumulated minutes that are part of Maximum Available Minutes that have no SAP HANA on Azure Connectivity. Downtime excludes Announced Single Instance Maintenance.

Monthly Uptime Percentage: The Monthly Uptime Percentage for SAP HANA on Azure Single Instance is calculated using the following formula

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of SAP HANA on Azure Single Instances:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%
<95%	100%

[Table of Contents / Definitions](#)

Scheduler

Additional Definitions:

“**Maximum Available Minutes**” is the total number of minutes in a billing month.

“**Planned Execution Time**” is a time at which a Scheduled Job is scheduled to begin executing.

“**Scheduled Job**” means an action specified by you to execute within Microsoft Azure according to a specified schedule.

Downtime: The total accumulated minutes in a billing month during which one or more of your Scheduled Jobs is in a state of delayed execution. A given Scheduled Job is in a state of delayed execution if it has not begun executing after a Planned Execution Time, provided that such delayed execution time shall not be considered Downtime if the Scheduled Job begins executing within thirty (30) minutes after a Planned Execution Time.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Search

Additional Definitions:

“**Average Error Rate**” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“**Error Rate**” is the total number of Failed Requests divided by Total Requests, across all Search Service Instances in a given Azure subscription, during a given one-hour interval. If the Total Requests in a one-hour interval is zero, the Error Rate for that interval is 0%.

“**Excluded Requests**” are all requests that are throttled due to exhaustion of resources allocated for a Search Service Instance, as indicated by an HTTP 503 status code and a response header indicating the request was throttled.

“**Failed Requests**” is the set of all requests within Total Requests that fail to return either a Success Code or HTTP 4xx response.

“**Replica**” is a copy of a search index within a Search Service Instance.

“**Search Service Instance**” is an Azure Search service instance containing one or more search indexes.

“**Total Requests**” is the set of (i) all requests to update a Search Service Instance having three or more Replicas, plus (ii) all requests to query a Search Service Instance having two or more Replicas, other than Excluded Requests, within a one-hour interval within a given Azure subscription during a billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \textit{Average Error Rate}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Free Search tier is not covered by this SLA.

[Table of Contents / Definitions](#)

Service-Bus Service – Event Hubs

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes that a given Event Hub has been deployed in Microsoft Azure during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Event Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Event Hubs tiers during a billing month.

“**Message**” refers to any user-defined content sent or received through Service Bus Relays, Queues, Topics, or Notification Hubs, using any protocol supported by Service Bus.

Downtime: The total accumulated Deployment Minutes, across all Event Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Event Hubs tiers, during which the Event Hub is unavailable. A minute is considered unavailable for a given Event Hub if all continuous attempts to send or receive Messages or perform other operations on the Event Hub throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\textit{Maximum Available Minutes-Downtime}}{\textit{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Basic and Standard Event Hubs tiers. The Free Event Hubs tier is not covered by this SLA.

[Table of Contents / Definitions](#)

Service-Bus Service – Notification Hubs

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes that a given Notification Hub has been deployed in Microsoft Azure during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Notification Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Notification Hubs tiers during a billing month.

Downtime: The total accumulated Deployment Minutes, across all Notification Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Notification Hubs tiers, during which the Notification Hub is unavailable. A minute is considered unavailable for a given Notification Hub if all continuous attempts to send notifications or perform registration management operations with respect to the Notification Hub throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Basic and Standard Notification Hubs tiers. The Free Notification Hubs tier is not covered by this SLA.

[Table of Contents](#) / [Definitions](#)

Service-Bus Service – Queues and Topics

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Queue or Topic has been deployed in Microsoft Azure during a billing month.
“Maximum Available Minutes” is the sum of all Deployment Minutes across all Queues and Topics deployed by you in a given Microsoft Azure subscription during a billing month.
“Message” refers to any user-defined content sent or received through Service Bus Relays, Queues, Topics, or Notification Hubs, using any protocol supported by Service Bus.

Downtime: The total accumulated Deployment Minutes, across all Queues and Topics deployed by you in a given Microsoft Azure subscription, during which the Queue or Topic is unavailable. A minute is considered unavailable for a given Queue or Topic if all continuous attempts to send or receive Messages or perform other operations on the Queue or Topic throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Service-Bus Service – Relays

Additional Definitions:

“Message” refers to any user-defined content sent or received through Service Bus Relays, Queues, or Topics, using any protocol supported by Service Bus.
“Deployment Minutes” is the total number of minutes that a given Relay has been deployed in Microsoft Azure during a billing month.
“Maximum Available Minutes” is the sum of all Deployment Minutes across all Relays deployed by Customer in a given Microsoft Azure subscription during a billing month.

Downtime: Is the total accumulated Deployment Minutes, across all Relays deployed by Customer in a given Microsoft Azure subscription, during which the Relay is unavailable. A minute is considered unavailable for a given Relay if all continuous attempts to establish a connection to the Relay throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime percentage for Relays is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\frac{\textit{Maximum Available Minutes-Downtime}}{\textit{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

SignalR Service

Additional Definitions:

“Downtime” is the total accumulated Maximum Available Minutes during a billing month for the SignalR Service during which the SignalR Service is unavailable. A given minute is considered unavailable if all attempts to send SignalR Transactions throughout the minute either return an Error Code or do not result in a Success Code within one minute.

“Maximum Available Minutes” is the total number of minutes that the SignalR Service has been deployed by the Customer in a given Microsoft Azure subscription during a billing month.

“SignalR Service Endpoint” is the host name from which the SignalR Service is accessed by servers or clients to perform SignalR Transactions.

“SignalR Transactions” is the set of transaction requests sent from client to server or from server to client through a SignalR Service Endpoint.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\textit{Maximum Available Minutes-Downtime}}{\textit{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the SignalR Service Standard tiers. The SignalR Service Free tier is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

SQL Data Warehouse Database

Additional Definitions:

“Database” means any SQL Data Warehouse Database.

“Maximum Available Minutes” is the total number of minutes that a given Database has been deployed in Microsoft Azure during a billing month in a given Microsoft Azure subscription.

“Client Operations” is the set of all documented operations supported by SQL Data Warehouse.

Downtime: is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Database is unavailable. A minute is considered unavailable for a given Database if more than 1% of all Client Operations completed during the minute return an Error Code.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\textit{Maximum Available Minutes-Downtime}}{\textit{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

SQL Database Service (Basic, Standard and Premium Tiers)

Additional Definitions:

“**Database**” means any single or elastic Basic, Standard, or Premium Microsoft Azure SQL Database.

“**Maximum Available Minutes**” is the total number of minutes that a given Database has been deployed in in Microsoft Azure during a billing month in a given Microsoft Azure subscription.

Downtime: is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Database is unavailable. A minute is considered unavailable for a given Database if all continuous attempts to establish a connection to the Database within the minute fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

SQL Database Service (Web and Business Tiers)

Additional Definitions:

“**Database**” means any Web or Business Microsoft Azure SQL Database.

“**Deployment Minutes**” is the total number of minutes that a given Web or Business Database has been deployed in Microsoft Azure during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Web and Business Databases for a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes across all Web and Business Databases deployed by you in a given Microsoft Azure subscription during which the Database is unavailable. A minute is considered unavailable for a given Database if all continuous attempts by you to establish a connection to the Database within the minute fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

SQL Server Stretch Database

Additional Definitions:

“**Database**” means one instance of SQL Server Stretch Database.

“**Maximum Available Minutes**” is the total number of minutes that a given Database has been deployed in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated minutes across all Databases deployed by Customer in a given Microsoft Azure subscription during which the Database is unavailable. A minute is considered unavailable for a given Database if all continuous attempts by Customer to establish a connection to the Database within the minute fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Storage Service

Additional Definitions:

“Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“Blob Storage Account” is a storage account specialized for storing data as blobs and provides the ability to specify an access tier indicating how frequently the data in that account is accessed.

“Block Blob Storage Account” is a storage account specialized for storing data as block or append blobs on solid-state drives.

“Cool Access Tier” is an attribute of a blob or account indicating it is infrequently accessed and has a lower availability service level than blobs in Hot Access Tier.

“Hot Access Tier” is an attribute of a blob or account indicating it is frequently accessed.

“Excluded Transactions” are storage transactions that do not count toward either Total Storage Transactions or Failed Storage Transactions. Excluded Transactions include pre-authentication failures; authentication failures; attempted transactions for storage accounts over their prescribed quotas; creation or deletion of containers, file shares, tables, or queues; clearing of queues; and copying blobs or files between storage accounts.

“Error Rate” is the total number of Failed Storage Transactions divided by the Total Storage Transactions during a set time interval (currently set at one hour). If the Total Storage Transactions in a given one-hour interval is zero, the error rate for that interval is 0%.

“Failed Storage Transactions” is the set of all storage transactions within Total Storage Transactions that are not completed within the Maximum Processing Time associated with their respective transaction type, as specified in the table below. Maximum Processing Time includes only the time spent processing a transaction request within the Storage Service and does not include any time spent transferring the request to or from the Storage Service.

Commented [A2]: Page 52

Transaction Type	Maximum Processing Time
PutBlob and GetBlob (includes blocks and pages) Get Valid Page Blob Ranges	Two (2) seconds multiplied by the number of MBs transferred in the course of processing the request
PutFile and GetFile	Two (2) seconds multiplied by the number of MBs transferred in the course of processing the request
Copy Blob	Ninety (90) seconds (where the source and destination blobs are within the same storage account)
Copy File	Ninety (90) seconds (where the source and destination files are within the same storage account)
PutBlockList GetBlockList	Sixty (60) seconds
Table Query List Operations	Ten (10) seconds (to complete processing or return a continuation)
Batch Table Operations	Thirty (30) seconds
All Single Entity Table Operations All other Blob, File, and Message Operations	Two (2) seconds

Commented [A3]: Page 52

These figures represent maximum processing times. Actual and average times are expected to be much lower.

Failed Storage Transactions do not include:

1.

Transaction requests that are throttled by the Storage Service due to a failure to obey appropriate back-off principles.
2.

Transaction requests having timeouts set lower than the respective Maximum Processing Times specified above.
3.

Read transactions requests to RA-GRS Accounts for which you did not attempt to execute the request against Secondary Region associated with the storage account if the request to the Primary Region was not successful.
4.

Read transaction requests to RA-GRS Accounts that fail due to Geo-Replication Lag.

“Geo Replication Lag” for GRS and RA-GRS Accounts is the time it takes for data stored in the Primary Region of the storage account to replicate to the Secondary Region of the storage account. Because GRS and RA-GRS Accounts are replicated asynchronously to the Secondary Region, data

written to the Primary Region of the storage account will not be immediately available in the Secondary Region. You can query the Geo Replication Lag for a storage account, but Microsoft does not provide any guarantees as to the length of any Geo Replication Lag under this SLA.

“Geographically Redundant Storage (GRS) Account” is a storage account for which data is replicated synchronously within a Primary Region and then replicated asynchronously to a Secondary Region. You cannot directly read data from or write data to the Secondary Region associated with GRS Accounts.

“Locally Redundant Storage (LRS) Account” is a storage account for which data is replicated synchronously only within a Primary Region.

“Primary Region” is a geographical region in which data within a storage account is located, as selected by you when creating the storage account. You may execute write requests only against data stored within the Primary Region associated with storage accounts.

“Read Access Geographically Redundant Storage (RA-GRS) Account” is a storage account for which data is replicated synchronously within a Primary Region and then replicated asynchronously to a Secondary Region. You can directly read data from, but cannot write data to, the Secondary Region associated with RA-GRS Accounts.

“Secondary Region” is a geographical region in which data within a GRS or RA-GRS Account is replicated and stored, as assigned by Microsoft Azure based on the Primary Region associated with the storage account. You cannot specify the Secondary Region associated with storage accounts.

“Total Storage Transactions” is the set of all storage transactions, other than Excluded Transactions, attempted within a one-hour interval across all storage accounts in the Storage Service in a given subscription.

“Zone Redundant Storage (ZRS) Account” is a storage account for which data is replicated across multiple facilities. These facilities may be within the same geographical region or across two geographical regions.

Monthly Uptime Percentage: Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit – hot blobs in LRS, ZRS, GRS and RA-GRS (write requests) Accounts and blobs in LRS Block Blob Storage Accounts:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Commented [A4]: Page 53 header on first Monthly update graph

Service Credit – RA-GRS (read requests) Accounts:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Service Credit – LRS, GRS and RA-GRS (write requests) Blob Storage Accounts (Cool Access Tier):

Monthly Uptime Percentage	Service Credit
< 99%	10%
< 98%	25%

Service Credit – RA-GRS (read requests) Blob Storage Accounts (Cool Access Tier):

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 98%	25%

[Table of Contents / Definitions](#)

Stream Analytics – API Calls

Additional Definitions:

“Total Transaction Attempts” is the total number of authenticated REST API requests to manage a streaming job within the Stream Analytics Service by Customer during a billing month for a given Microsoft Azure subscription.

“Failed Transactions” is the set of all requests within Total Transaction Attempts that return an Error Code or otherwise do not return a Success Code within five minutes from Microsoft’s receipt of the request.

“Monthly Uptime Percentage” for API calls within the Stream Analytics Service is represented by the following formula:

$$\text{Monthly Uptime \%} = \frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Stream Analytics – Jobs

Additional Definitions:
“**Deployment Minutes**” is the total number of minutes that a given job has been deployed within the Stream Analytics Service during a billing month.
“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all jobs deployed by Customer in a given Microsoft Azure subscription during a billing month.

Downtime is the total accumulated Deployment Minutes, across all jobs deployed by Customer in a given Microsoft Azure subscription, during which the job is unavailable. A minute is considered unavailable for a deployed job if the job is neither processing data nor available to process data throughout the minute.

Monthly Uptime Percentage for jobs within the Stream Analytics Service is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Traffic Manager Service

Additional Definitions:
“**Deployment Minutes**” is the total number of minutes that a given Traffic Manager Profile has been deployed in Microsoft Azure during a billing month.
“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Traffic Manager Profiles deployed by you in a given Microsoft Azure subscription during a billing month.
“**Traffic Manager Profile**” or “**Profile**” refers to a deployment of the Traffic Manager Service created by you containing a domain name, endpoints, and other configuration settings, as represented in the Management Portal.
“**Valid DNS Response**” means a DNS response, received from at least one of the Traffic Manager Service name server clusters, to a DNS request for the domain name specified for a given Traffic Manager Profile.

Downtime: The total accumulated Deployment Minutes, across all Profiles deployed by you in a given Microsoft Azure subscription, during which the Profile is unavailable. A minute is considered unavailable for a given Profile if all continual DNS queries for the DNS name specified in the Profile that are made throughout the minute do not result in a Valid DNS Response within two seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Virtual Machines

Additional Definitions:

“**Availability Set**” refers to two or more Virtual Machines deployed across different Fault Domains to avoid a single point of failure.

“**Availability Zone**” is a fault-isolated area within an Azure region, providing redundant power, cooling, and networking.

“**Data Disk**” is a persistent virtual hard disk, attached to a Virtual Machine, used to store application data.

“**Fault Domain**” is a collection of servers that share common resources such as power and network connectivity.

“**Operating System Disk**” is a persistent virtual hard disk, attached to a Virtual Machine, used to store the Virtual Machine’s operating system.

“**Single Instance**” is defined as any single Microsoft Azure Virtual Machine that either is not deployed in an Availability Set or has only one instance deployed in an Availability Set.

“**Virtual Machine**” refers to persistent instance types that can be deployed individually or as part of an Availability Set.

“**Virtual Machine Connectivity**” is bi-directional network traffic between the Virtual Machine and other IP addresses using TCP or UDP network protocols in which the Virtual Machine is configured for allowed traffic. The IP addresses can be IP addresses in the same Cloud Service as the Virtual Machine, IP addresses within the same virtual network as the Virtual Machine or public, routable IP addresses.

Monthly Uptime Calculation and Service Levels for Virtual Machines in Availability Zones

“**Maximum Available Minutes**” is the total accumulated minutes during a billing month that have two or more instances deployed across two or more Availability Zones in the same region. Maximum Available Minutes is measured from when at least two Virtual Machines across two Availability Zones in the same region have both been started resultant from action initiated by Customer to the time Customer has initiated an action that would result in stopping or deleting the Virtual Machines.

“**Downtime**” is the total accumulated minutes that are part of Maximum Available Minutes that have no Virtual Machine Connectivity in the region.

“**Monthly Uptime Percentage**” for Virtual Machines in Availability Zones is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\text{Monthly Uptime \%} = \frac{(\text{Maximum Available Minutes}-\text{Downtime})}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

The following Service Levels and Service Credits are applicable to Customer’s use of Virtual Machines deployed across two or more Availability Zones in the same region:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%
< 95%	100%

Monthly Uptime Calculation and Service Levels for Virtual Machines in an Availability Set

“**Maximum Available Minutes**”: The total accumulated minutes during a billing month for all Internet facing Virtual Machines that have two or more instances deployed in the same Availability Set. Maximum Available Minutes is measured from when at least two Virtual Machines in the same Availability Set have both been started resultant from action initiated by you to the time you have initiated an action that would result in stopping or deleting the Virtual Machines.

“**Downtime**”: The total accumulated minutes that are part of Maximum Available Minutes that have no Virtual Machine Connectivity.

“**Monthly Uptime Percentage**”: for Virtual Machines is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\text{Monthly Uptime \%} = \frac{(\text{Maximum Available Minutes}-\text{Downtime})}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

The following Service Levels and Service Credits are applicable to Customer’s use of Virtual Machines in an Availability Set:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%
< 95%	100%

Monthly Uptime Calculation and Service Levels for Single-Instance Virtual Machines

“Minutes in the Month” is the total number of minutes in a given month.

Downtime: is the total accumulated minutes that are part of Minutes in the Month that have no Virtual Machine Connectivity.

Monthly Uptime Percentage: is calculated by subtracting from 100% the percentage of Minutes in the Month in which any Single Instance Virtual Machine using premium storage for all Operating System Disks and Data Disks had Downtime.

$$\text{Monthly Uptime \%} = \frac{(\text{Minutes in the Month} - \text{Downtime})}{\text{Minutes in the Month}} \times 100$$

Service Credit:
The following Service Levels and Service Credits are applicable to Customer’s use of Single-Instance Virtual Machines:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

VPN Gateway

Additional Definitions:
“Maximum Available Minutes” is the total accumulated minutes during a billing month during which a given VPN Gateway has been deployed in a Microsoft Azure subscription.

Downtime: Is the total accumulated Maximum Available Minutes during which a VPN Gateway is unavailable. A minute is considered unavailable if all attempts to connect to the VPN Gateway within a thirty-second window within the minute are unsuccessful.

Monthly Uptime Percentage: The Monthly Uptime Percentage for a given VPN Gateway is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for the VPN Gateway. The Uptime Percentage is represented by the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of each VPN Gateway:
Basic Gateway for VPN or ExpressRoute Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Gateway for VPN and Gateway for ExpressRoute SKUs excluding Basic Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Visual Studio App Center Build Service

Additional Definitions:
“Build Service” is a feature that allows customers to build their mobile applications in Visual Studio App Center.
“Maximum Available Minutes” is the total number of minutes for which Build Service has been deployed by Customer for a given Microsoft Azure subscription during a billing month.

Commented [A5]: Page 56

“Downtime” is the total number of minutes within Maximum Available Minutes during which the Build Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Build Service to perform operations initiated by Customer throughout the minute either result in an Error Code or do not return a response within one minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Visual Studio App Center Build Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the Visual Studio App Center Build Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Visual Studio App Center Test Service

Additional Definitions:

“Test Service” is a feature that allows customers to upload and run tests for their mobile applications on physical devices running in Visual Studio App Center.

“Maximum Available Minutes” is the total number of minutes for which Test Service has been deployed by Customer for a given Microsoft Azure subscription during a billing month.

Downtime: The total number of minutes within Maximum Available Minutes during which the Test Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Test Service to perform operations initiated by Customer throughout the minute either result in an Error Code or do not return a response within one minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Visual Studio App Center Test Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the Visual Studio App Center Test Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Visual Studio App Center Push Notification Service

Additional Definitions:

“Push Notification Service” is a feature that enables customers to push messages to specific devices configured to receive such messages using Visual Studio App Center.

“Maximum Available Minutes” is the total number of minutes for which Push Notification Service has been deployed by Customer for a given Microsoft Azure subscription during a billing month.

Downtime: The total number of minutes within Maximum Available Minutes during which Push Notification Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to Push Notification Service to perform operations initiated by Customer throughout the minute either result in an Error Code or do not return a response within one minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Visual Studio App Center Push Notification Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the Visual Studio App Center Push Notification Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure Dev Ops Services – Azure Pipelines

Additional Definitions:

“**Azure Pipelines**” is a feature that allows customers to build and deploy their applications in Azure DevOps Services.

“**Maximum Available Minutes**” is the total number of minutes for which the paid Azure Pipelines has been enabled for a given Microsoft Azure subscription during a billing month.

“**Downtime**” is the total accumulated minutes for a given Microsoft Azure subscription during which the Azure Pipelines Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Azure Pipelines Service to perform operations initiated by you throughout the minute either result in an Error Code or do not return a response.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Azure Pipelines Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the Azure Pipelines Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure DevOps Test Plans – Load Testing Service

Additional Definitions:

“**Azure DevOps Test Plans Load Testing Service**” is a feature that allows customers to generate automated tasks to test the performance and scalability of applications.

“**Maximum Available Minutes**” is the total number of minutes for which the paid Azure DevOps Test Plans Load Testing Service has been enabled for a given Microsoft Azure subscription during a billing month.

“**Downtime**” is the total accumulated minutes for a given Microsoft Azure subscription during which the Azure DevOps Test Plans Load Testing Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Azure DevOps Test Plans Load Testing Service to perform operations initiated by you throughout the minute either result in an Error Code or do not return a response.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Azure DevOps Test Plans Load Testing Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

The following Service Levels and Service Credits are applicable to Customer’s use of the Azure Test Plans Load Testing Service.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Azure DevOps Services – User Plans Service

Additional Definitions:

“**Azure DevOps Test Plans Load Testing Service**” is a feature that allows customers to generate automated tasks to test the performance and scalability of applications.

“**Azure DevOps Services Users**” refers to the set of features and capabilities available to a user within an Azure DevOps Services account in a Customer subscription. The features and capabilities available are described on the [Azure DevOps](#) website.

“**Azure Pipelines**” is a feature that allows customers to build and deploy their applications in Azure DevOps Services.

“**Deployment Minutes**” is the total number of minutes for which a User Plan has been purchased during a billing month.

“**Downtime**” is the total accumulated Deployment Minutes, across all User Plans for a given Microsoft Azure subscription, during which the Service Plan is unavailable. A minute is considered unavailable for a given User Plan if all continuous HTTP requests to perform operations, other than operations pertaining to the Azure Pipelines Service or the Azure DevOps Test Plans Load Testing Service, throughout the minute either result in an Error Code or do not return a response.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all User Plans for a given Microsoft Azure subscription during a billing month.

“**User-Based Extensions**” means the ser of Azure DevOps Services extensions published by Microsoft which are sold on a per-user basis via the Azure DevOps Marketplace.

“**User Plans**” refer to Azure DevOps Services Users and User-Based Extensions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

In the event Azure DevOps Services are unavailable, Service Credits are applicable to Azure DevOps Services Users and User-Based Extensions. The following Service Levels and Service Credits are applicable to Customer’s use of Azure DevOps Services User Plans:

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Microsoft Azure Plans

Azure Active Directory Basic

Downtime: Any period of time when users are not able to log in to the service, log in to the Access Panel, access applications on the Access Panel and reset passwords; or any period of time IT administrators are not able to create, read, write and delete entries in the directory and/or provision/de-provision users to applications in the directory.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Azure Active Directory B2C

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes for which an Azure AD B2C directory has been deployed during a billing month.
“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Azure AD B2C directories in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated minutes across all Azure AD B2C directories deployed by Customer in a given Microsoft Azure subscription during which the Azure AD B2C service is unavailable. A minute is considered unavailable if either all attempts to process user sign-up, sign-in, profile editing, password reset and multi-factor authentication requests, or all attempts by developers to create, read, write and delete entries in a directory, fails to return tokens or valid Error Codes, or do not return responses within two minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: No SLA is provided for the Free tier of Azure Active Directory B2C.

[Table of Contents](#) / [Definitions](#)

Azure Active Directory Premium

Downtime: Any period of time when users are not able to log in to the service, log in to the Access Panel, access applications on the Access Panel and reset passwords; or any period of time IT administrators are not able to create, read, write and delete entries in the directory and/or provision/de-provision users to applications in the directory.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Azure Information Protection Premium

Downtime: Any period of time when end users cannot create or consume IRM documents and email.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Azure Site Recovery Service – On-Premises-to-Azure

Additional Definitions:

“Failover” is the process of transferring control, either simulated or actual, of a Protected Instance from a primary site to a secondary site.

“On-Premises-to-Azure Failover” is the Failover of a Protected Instance from a non-Azure primary site to an Azure secondary site.

“Protected Instance” refers to a virtual or physical machine configured for replication by the Azure Site Recovery Service from a primary site to a secondary site. Protected Instances are enumerated in the Protected Items tab in the Recovery Services section of the Management Portal.

“Recovery Time Objective (RTO)” means the period of time beginning when you initiate a Failover of a Protected Instance experiencing either a planned or unplanned outage for On-Premises-to-Azure replication to the time when the Protected Instance is running as a virtual machine in Microsoft Azure, excluding any time associated with manual action or the execution of your scripts.

“Monthly Recovery Time Objective”: For a specific Protected Instance configured for On-Premises-to-Azure replication in a given billing month is two hours.

Service Credit:

Monthly Recovery Time Objective	Service Credit
> 2 hours	100%

Additional Terms: Monthly Recovery Time Objective and Service Credits are calculated for each Protected Instance used by you.

[Table of Contents](#) / [Definitions](#)

Azure Site Recovery Service – On-Premises-to-On-Premises

Additional Definitions:

“Failover” is the process of transferring control, either simulated or actual, of a Protected Instance from a non-Azure primary site to a non-Azure secondary site.

“Failover Minutes” is the total number of minutes in a billing month during which a Failover of a Protected Instance configured for On-Premises-to-On-Premises replication has been attempted but not completed.

“Maximum Available Minutes” is the total number of minutes that a given Protected Instance has been configured for On-Premises-to-On-Premises replication by the Azure Site Recovery Service during a billing month.

“On-Premises-to-On-Premises Failover” is the Failover of a Protected Instance from a non-Azure primary site to a non-Azure secondary site.

“Protected Instance” refers to a virtual or physical machine configured for replication by the Azure Site Recovery Service from a primary site to a secondary site. Protected Instances are enumerated in the Protected Items tab in the Recovery Services section of the Management Portal.

Downtime: Is the total accumulated Failover Minutes in which the Failover of a Protected Instance is unsuccessful due to unavailability of the Azure Site Recovery Service, provided that retries are continually attempted no less frequently than once every thirty minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Additional Terms: Monthly Recovery Time Objective and Service Credits are calculated for each Protected Instance used by you.

Table of Contents / Definitions

Azure Site Recovery Service – Azure-to-Azure Failover

Additional Definitions:

“Failover” is the process of transferring control, either simulated or actual, of a Protected Instance from a primary site to a secondary site.

“Azure-to-Azure Failover” is the Failover of a Protected Instance from an Azure primary site to an Azure secondary site.

“Protected Instance” refers to a virtual or physical machine configured for replication by the Azure Site Recovery Service from a primary site to a secondary site. Protected Instances are enumerated in the Protected Items tab in the Recovery Services section of the Management Portal.

“Recovery Time Objective (RTO)” means the period of time beginning when Customer initiates a Failover of a Protected Instance for Azure-to-Azure replication to the time when the Protected Instance is running as a virtual machine in secondary Azure region, excluding any time associated with manual action or the execution of Customer scripts.

“Monthly Recovery Time Objective” for a specific Protected Instance configured for Azure-to-Azure replication in a given billing month is 2 hours.

Service Credit:

Monthly Recovery Time Objective	Service Credit
>2 hours	100%

Additional Terms: Monthly Recovery Time Objective and Service Credits are calculated for each Protected Instance used by you.

Table of Contents / Definitions

Multi-Factor Authentication Service

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Multi-Factor Authentication provider has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Multi-Factor Authentication providers deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes, across all Multi-Factor Authentication providers deployed by you in a given Microsoft Azure subscription, during which the Multi-Factor Authentication Service is unable to receive or process authentication requests for the Multi-Factor Authentication provider.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

StorSimple Service

Additional Definitions:

“**Backup**” is the process of backing up data stored on a registered StorSimple device to one or more associated cloud storage accounts within Microsoft Azure.

“**Cloud Tiering**” is the process of transferring data from a registered StorSimple device to one or more associated cloud storage accounts within Microsoft Azure.

“**Deployment Minutes**” is the total number of minutes during which a Managed Item has been configured by Customer for Backup or Cloud Tiering to a StorSimple storage account in Microsoft Azure.

“**Failure**” means the inability to fully complete a properly configured Backup, Tiering, or Restoring operation due to unavailability of the StorSimple Service.

“**Managed Item**” refers to a volume that has been configured to Backup to the cloud storage accounts using the StorSimple Service.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Managed Items for a given Microsoft Azure subscription during a billing month.

“**Restoring**” is the process of copying data to a registered StorSimple device from its associated cloud storage account(s).

Downtime: The total number of minutes within Maximum Available Minutes during which the StorSimple Service is unavailable for the Managed Item. The StorSimple Service is considered unavailable for a given Managed Item from the first Failure of a Backup, Cloud Tiering, or Restoring operation with respect to the Managed Item until the initiation of a successful Backup, Cloud Tiering, or Restoring operation of the Managed Item, provided that retries are continually attempted no less frequently than once every thirty minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

StorSimple Data Manager

Additional Definitions:

“**Total Requests**” is the set of all requests, other than Excluded Requests, to perform operations against StorSimple Data Manager service during a billing month for a given Microsoft Azure subscription.

“**Excluded Requests**” is the set of requests that result in an HTTP 4xx status code.

“**Failed Requests**” is the set of all requests within Total Requests that either return an Error Code or fail to return a Success Code within 60 seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ Requests - Failed\ Requests}{Total\ Requests} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Other Online Services

Bing Maps Enterprise Platform

Downtime: Any period of time when the Service is not available as measured in Microsoft’s data centers, provided that you access the Service using the methods of access, authentication and tracking methods documented in the Bing Maps Platform SDKs.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured as the total number of minutes during the month when the aspects of the Service set forth above are unavailable.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to Bing Maps Enterprise Platform purchased through Open Value and Open Value Subscription volume licensing agreements.

Service Credits will not apply if: (i) you fail to implement any Services updates within the time specified in the Bing Maps Platform API’s Terms of Use; and (ii) you do not provide Microsoft with at least ninety (90) days’ advance notice of any known significant usage volume increase, with significant usage volume increase defined as 50% or more of the previous month’s usage.

Bing Maps Mobile Asset Management

Downtime: Any period of time when the Service is not available as measured in Microsoft’s data centers, provided that you access the Service using the methods of access, authentication and tracking methods documented in the Bing Maps Platform SDKs.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured as the total number of minutes during the month when the aspects of the Service set forth above are unavailable.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to Bing Maps Enterprise Platform purchased through Open Value and Open Value Subscription volume licensing agreements.

Service Credits will not apply if: (i) you fail to implement any Services updates within the time specified in the Bing Maps Platform API’s Terms of Use; and (ii) you do not provide Microsoft with at least ninety (90) days’ advance notice of any known significant usage volume increase, with significant usage volume increase defined as 50% or more of the previous month’s usage.

Microsoft Cloud App Security

Downtime: Any period of time when the Customer’s IT administrator or users authorized by Customer are unable to log on with proper credentials. Scheduled Downtime will not exceed 10 hours per calendar year.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: This Service Level does not apply to any: (i) On-premises software licensed as part of the Service subscription, or (ii) Internet-based services (excluding Microsoft Cloud App Security) that provide updates via API (application programming interface) to any services licensed as part of the Service subscription.

Table of Contents / Definitions

Microsoft Flow

Downtime: Any period of time when users’ flows have no connectivity to Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month - Downtime}{Total\ number\ of\ minutes\ in\ a\ month} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: No SLA is provided for any free of charge tier of Microsoft Flow.

Table of Contents / Definitions

Microsoft Intune

Downtime: Any period of time when the Customer’s IT administrator or users authorized by Customer are unable to log on with proper credentials. Scheduled Downtime will not exceed 10 hours per calendar year.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This Service Level does not apply to any: (i) On-premises software licensed as part of the Service subscription, or (ii) Internet-based services (excluding Microsoft Intune Service) that provide updates to any on-premise software licensed as part of the Service subscription.

[Table of Contents](#) / [Definitions](#)

Microsoft Kaizala Pro

Downtime: Any period of time when end users are unable to read or post message in organizations groups for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Microsoft PowerApps

Downtime: Any period of time when users are unable to read or write any portion of data in Microsoft PowerApps to which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month - Downtime}{Total\ number\ of\ minutes\ in\ a\ month} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: No SLA is provided for any free of charge tier of Microsoft PowerApps.

[Table of Contents](#) / [Definitions](#)

Minecraft: Education Edition

Downtime: Any period of time when users are unable to access Minecraft: Education Edition.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Power BI Embedded

Deployment Minutes: is the total number of minutes for which a given workspace collection has been provisioned during a billing month.

Maximum Available Minutes: is the sum of all Deployment Minutes across all workspace collections provisioned by a customer in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment Minutes, during which the workspace collection is unavailable. A minute is considered unavailable for a given workspace collection if all continuous attempts within the minute to read or write any portion of Power BI Embedded data result in an Error Code or do not return a response within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Power BI Premium

“**Capacity**” means a named capacity provisioned by an admin through the Power BI Premium capacity admin portal. A Capacity is a grouping of one or more nodes.

“**Maximum Available Minutes**” is the total number of minutes that a given Capacity has been instantiated during a billing month in a given tenant.

Downtime: The total accumulated minutes during a billing month for a given Capacity during which a given Capacity is unavailable. A minute is considered unavailable for a given Capacity if all attempts to view Power BI reports or dashboards within the minute fail due to system errors.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Power BI Pro

Downtime: Any period of time when users are unable to read or write any portion of Power BI data to which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Translator API

Downtime: Any period of time when users are not able to perform translations.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured as the total number of minutes during the month when the aspects of the Service set forth above are unavailable.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Windows Desktop Operating System

Additional Definitions:

“**Maximum Available Minutes**” is the total accumulated minutes during a billing month for Windows Defender Advanced Threat Protection portal. Maximum Available Minutes is measured from when the Tenant has been created resultant from successful completion of the on-boarding process.

“**Tenant**” represents Windows Defender Advanced Threat Protection customer specific cloud environment.

Downtime: The total accumulated minutes that are part of Maximum Available Minutes in which the Customer unable to access any portion of a Windows Defender Advanced Threat Protection portal site collections for which they have appropriate permissions and customer has a valid, active, license.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: This SLA does not apply to any trial/preview version Tenants.

Appendix A – Service Level Commitment for Virus Detection and Blocking, Spam Effectiveness, or False Positive

With respect to Exchange Online and EOP licensed as a standalone Service or via ECAL suite, or Exchange Enterprise CAL with Services, you may be eligible for Service Credits if we do not meet the Service Level described below for: (1) Virus Detection and Blocking, (2) Spam Effectiveness, or (3) False Positive. If any one of these individual Service Levels is not met, you may submit a claim for a Service Credit. If one Incident causes us to fail more than one SLA metric for Exchange Online or EOP, you may only make one Service Credit claim for that incident per Service.

1. **Virus Detection and Blocking Service Level**

- a. “Virus Detection and Blocking” is defined as the detection and blocking of Viruses by the filters to prevent infection. “Viruses” is broadly defined as known malware, which includes viruses, worms, and Trojan horses.
- b. A Virus is considered known when widely used commercial virus scanning engines can detect the virus and the detection capability is available throughout the EOP network.
- c. Must result from a non-purposeful infection.
- d. The Virus must have been scanned by the EOP virus filter.
- e. If EOP delivers an email that is infected with a known virus to you, EOP will notify you and work with you to identify and remove it. If this results in the prevention of an infection, you won’t be eligible for a Service Credit under the Virus Detection and Blocking Service Level.
- f. The Virus Detection and Blocking Service Level shall not apply to:
 - i. Forms of email abuse not classified as malware, such as spam, phishing and other scams, adware, and forms of spyware, which due to its targeted nature or limited use is not known to the anti-virus community and thus not tracked by anti-virus products as a virus.
 - ii. Corrupt, defective, truncated, or inactive viruses contained in NDRs, notifications, or bounced emails.
- g. The Service Credit available for the Virus Detection and Blocking Service is: 25% Service Credit of Applicable Monthly Service Fee if an infection occurs in a calendar month, with a maximum of one claim allowed per calendar month.

2. **Spam Effectiveness Service Level**

- a. “Spam Effectiveness” is defined as the percentage of inbound spam detected by the filtering system, measured on a daily basis.
- b. Spam effectiveness estimates exclude false negatives to invalid mailboxes.
- c. The spam message must be processed by our service and not be corrupt, malformed, or truncated.
- d. The Spam Effectiveness Service Level does not apply to email containing a majority of non-English content.
- e. You acknowledge that classification of spam is subjective and accept that we will make a good faith estimation of the spam capture rate based on evidence timely supplied by you.
- f. The Service Credit available for the Spam Effectiveness Service is:

% of Calendar Month that Spam Effectiveness is below 99%	Service Credit
>25%	25%
> 50%	50%
100%	100%

3. **False Positive Service Level**

- a. “False Positive” is defined as the ratio of legitimate business email incorrectly identified as spam by the filtering system to all email processed by the service in a calendar month.
- b. Complete, original messages, including all headers, must be reported to the abuse team.
- c. Applies to email sent to valid mailboxes only.
- d. You acknowledge that classification of false positives is subjective and understand that we will make a good faith estimation of the false positive ratio based on evidence timely supplied by you.
- e. This False Positive Service Level shall not apply to:
 - i. bulk, personal, or pornographic email
 - ii. email containing a majority of non-English content
 - iii. email blocked by a policy rule, reputation filtering, or SMTP connection filtering
 - iv. email delivered to the junk folder
- f. The Service Credit available for the False Positive Service is:

False Positive Ratio in a Calendar Month	Service Credit
> 1:250,000	25%
> 1:10,000	50%
> 1:100	100%

Appendix B - Service Level Commitment for Uptime and Email Delivery

With respect to EOP licensed as a standalone Service, ECAL suite, or Exchange Enterprise CAL with Services, you may be eligible for Service Credits if we do not meet the Service Level described below for (1) Uptime and (2) Email Delivery.

1. **Monthly Uptime Percentage:**
If the Monthly Uptime Percentage for EOP falls below 99.999% for any given month, you may be eligible for the following Service Credit:
- | Monthly Uptime Percentage | Service Credit |
|---------------------------|----------------|
| <99.999% | 25% |
| <99.0% | 50% |
| <98.0% | 100% |
2. **Email Delivery Service Level:**

a. "Email Delivery Time" is defined as the average of email delivery times, measured in minutes over a calendar month, where email delivery is defined as the elapsed time from when a business email enters the EOP network to when the first delivery attempt is made.
b. Email Delivery Time is measured and recorded every 5 minutes, then sorted by elapsed time. The fastest 95% of measurements are used to create the average for the calendar month.
c. We use simulated or test emails to measure delivery time.
d. The Email Delivery Service Level applies only to legitimate business email (non-bulk email) delivered to valid email accounts.
e. This Email Delivery Service Level does not apply to:

1. Delivery of email to quarantine or archive
2. Email in deferral queues
3. Denial of service attacks (DoS)
4. Email loops

f. The Service Credit available for the Email Delivery Service is:

Average Email Delivery Time (as defined above)	Service Credit
> 1	25%
> 4	50%
> 10	100%