

AWS MANAGED SERVICES SERVICE LEVEL AGREEMENT

Last updated November 6, 2019

This AWS Managed Services Service Level Agreement (“SLA”) is a policy governing the use of AWS Managed Services (“AMS”) under the terms of the Amazon Web Services Customer Agreement or Enterprise Agreement (the “AWS Agreement”) between Amazon Web Services, Inc. and its affiliates (“AWS”, “us” or “we”) and users of AWS’ services (“you”). This SLA applies separately to each account using AWS Managed Services. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Service Commitments

AWS will use commercially reasonable efforts to meet the following Service Commitments:

- **AWS Console/API Availability** - AWS will make the AMS console and AMS APIs available as set out in the Service Commitment & Credit Table (SCCT) below.
- **Incident Response Time** – Once an Incident is reported by you, AWS Managed Services will send an initial response to you concerning the Incident via the AMS console, e-mail, service center, or telephone within the timeframes set out in the SCCT below.
- **Incident Resolution Time** – AWS Managed Services will Resolve Incidents reported by AWS Managed Services or you within the timeframes set out in the SCCT below.
- **Patch Management** – AWS Managed Services will attempt to apply or install new updates to EC2 instances and provision AWS Managed Services AMIs with new updates, as applicable, within your Managed Environment within the timeframes set out in the SCCT below. This Service Commitment only applies to vendor updates for supported operating systems and software pre-installed with supported operating systems. A list of supported operating systems is available in the AWS Managed Services Service Description¹.
- **Environment Recovery Initiation Time**– AWS will initiate a customer-authorized Environment Recovery, as needed, within the timeframes set out in the SCCT below.

If AWS Managed Services does not meet a Service Commitment in Conformance with the Service Commitment & Credit Table, you will be eligible to receive a Service Credit as described below.

¹ See [Service Description](#) for more information.

Service Commitment & Credit Table (SCCT)

Service Commitment Category	Key Performance Indicator	Service Commitment ²		Conformance	Service Credits (% of Total AMS Monthly Fee)
		Plus ³	Premium ⁴		
AMS API and Console Availability	1. API Availability Percentage	>=99.95%*	>=99.95%*	100%	0.5%
	2. Console Availability Percentage	>=99.95%*	>=99.95%*	100%	0.5%
Incident Management - Response Time**	3. Incident Response Time for Priority 1 Incident	<=4 hours	<=15 min	95%	3%
	4. Incident Response Time for Priority 2 Incident	<=8 hours	<=4 hours	95%	2%
	5. Incident Response Time for Priority 3 Incident	<=24 hours	<=12 hours	90%	1%
Incident Management - Resolution Time**, ***	6. Incident Resolution Time for Priority 1 Incident	<=12 hours	<=4 hours	95%	6%
	7. Incident Resolution Time for Priority 2 Incident	<=24 hours	<=8 hours	95%	4%
	8. Incident Resolution Time for Priority 3 Incident	<=48 hours	<=24 hours	90%	2%
Patch Management***	9. Patching Time for a Critical Security Update	Within 10 business days of release by the vendor	Within 8 calendar days of release by the vendor	100%	4%
	10. Patching Time for an Important Update	Within two months of release by the vendor	Within two months of release by the vendor	95%	3%
Continuity Management - Environment Recovery ***	11. Environment Recovery Initiation Time	<=12 hours	<=4 hours	100%	4%

² References to minutes or hours within the table below refer to “Business Hours” as defined in the AWS Managed Services Service Description. The AWS Managed Services Maintenance Window is excluded from all Service Commitment time calculations.

³ See AWS Managed Services Pricing Plan for Plus and Premium options.

⁴ See AWS Managed Services Pricing Plan for Plus and Premium options.

* API Availability Percentage and Console Availability Percentage are each calculated by subtracting from 100%, the average Unavailability rate from each five minute period in the monthly billing cycle. The Unavailability rate is (i) the total number of Unavailable responses divided by (ii) the total number of requests for the applicable request type during the five minute period.

** If five (5) or more Priority 1 Incidents, caused due to application issues, are reported on any individual Stack during any rolling 30 day period, any subsequent Incidents for the same Stack will be excluded for the purposes of calculating Service Credits until AWS Managed Services determines otherwise. AWS Managed Services will escalate the issue with you in the monthly service review meetings to determine what, if any, changes are needed before the Stacks are included in Service Credit Calculations.

*** Does not apply to resources provisioned through Self-service Provisioning. See AWS Managed Services [Service Description](#) for more information.

Definitions

Capitalized terms are defined below or in the AWS Managed Services Service Description⁵.

- **“Unavailable”** and **“Unavailability”** mean:
 - For AWS Managed Services APIs, if a request submitted by you results in an error or other response that is not an expected output for your request.
 - For AWS Managed Services console, if an HTTP request submitted by you results in a 5xx HTTP response (where “x” represents any single digit number).
 - For AWS Stacks and resources, if any of the AWS Service Offerings that constitute the Stack(s) or resource(s) are in a state of “Service Disruption” as indicated in <http://status.aws.amazon.com/>.
 - Unavailability resulting directly or indirectly from an AWS Managed Services Exclusion will not be considered in determining eligibility for Service Credits.
 - Services are considered available unless they meet the criteria for being Unavailable.
- The **“AWS Managed Services Maintenance Window”** or **“Maintenance Window”** is a recurring time window that occurs on the second Thursday of every month from 3pm to 4pm Pacific Time, to perform maintenance activities for AWS Managed Services. AWS Managed Services may change the Maintenance Window by providing 48 hours notice.
- **“Infrastructure Restore”** means re-deploying existing Stack(s), based on templates of impacted Stack(s), and initiating a data restore based on the last known restore point, unless otherwise specified by the customer. Ephemeral data that is not part of the Stack template or data restore will be lost. AWS Managed Services will use reasonable efforts to perform Infrastructure Restore while AWS Service Offerings are Unavailable. Infrastructure Restore will be completed once AWS Service Offerings are available.
- **“Incident Resolution”** or **“Resolved”** Incident means that either (1) AWS Managed Services has restored all Unavailable services or resources pertaining to that Incident to an available state, or (2) where AWS Managed Services determines that Unavailable Stack(s) or resource(s) cannot be restored to an available state, AWS Managed Services has initiated a customer-authorized Infrastructure Restore. If you do not authorize an Infrastructure Restore as recommended by AWS when an Infrastructure Restore will bring all the resources pertaining to that Incident to an available state, you will not be eligible for a Service Credit for the associated Incident Resolution Time Service Commitment.
- **“Incident Response Time”** means the difference in time between when you create an Incident, and when AWS Managed Services provides an initial response via console, e-mail, service center, or telephone.
- **“Incident Resolution Time”** means the difference in time between when either AWS Managed Services or you create an Incident, and when the Incident is Resolved. Time spent waiting for inputs or approvals from you is excluded from Incident Resolution Time calculations. For Incidents that AWS Managed Services creates, the Incident creation time is the time of the initial customer notification.
- **“Incident Priority”** – Incidents will be categorized by AWS Managed Services or you as either Priority 1, 2, or 3.
 - **“Priority 1”** means that either (1) the AWS Managed Services Console, or one or more AWS Managed Services APIs within your Managed Environment are Unavailable; or (2) one or more AWS Managed Services Stacks or resources within your Managed Environment are Unavailable and the Unavailability prevents your application from performing its normal function.
 - **“Priority 2”** means that an AWS service within your Managed Environment is available but is not performing as intended (per the applicable service description).
 - **“Priority 3”** includes any Incident that is not categorized as Priority 1 or Priority 2.
 - AWS Managed Services may re-categorize Incidents in accordance with the above guidelines.

⁵ See [Service Description](#) for more information.

- **“Patching”** means attempting to apply or install new updates to EC2 instances, and provisioning AWS Managed Services AMIs with new updates, as applicable, within your Managed Environment based on your infrastructure configuration. AWS will only perform Patching for vendor updates to supported operating systems and software pre-installed with supported operating systems. Once an update is released by the vendor, AWS Managed Services will create an RFC to apply the update. You must approve the RFC before the update can be applied to your EC2 instances.
- **“Patching Time”** means the difference in time between the release of an update(s) for a supported operating system or software pre-installed with a supported operating system, and completion of Patching. Time spent waiting for inputs or approval from you is excluded from Patching Time. You must approve RFCs for Critical Security Updates within 14 days, and RFCs for Important Updates within 21 days, to be eligible for Service Credits for the respective updates.
- **“Environment Recovery”** – In case of Availability Zone (AZ) Unavailability in a Region used by your AWS Managed Services account, “Environment Recovery” is the process of restoring one or more AWS subnets in your Managed Environment by re-deploying existing stacks, based on templates of impacted stacks, and initiating a data restore based on a last known restore point, unless otherwise advised by the customer.
- **“Environment Recovery Initiation Time”** means the difference in time between when you request or authorize an Environment Recovery and the time AWS Managed Services initiates the Environment Recovery process. Time spent waiting for inputs or approvals from you is excluded from Environment Recovery Initiation Time calculations.
- **“Conformance”** is the percentage of times that AWS Managed Services must meet a Service Commitment in any monthly billing cycle. If AWS Managed Services does not meet the Conformance percentage for any Service Commitment, you will be eligible for a Service Credit.
 - For the purpose of determining Conformance for the Patch Management Service Commitment, each release of an update or multiple updates released simultaneously by an AWS Managed Services-supported operating system vendor will be considered as a single update.
- A **“Service Credit”** is a dollar credit, calculated as set forth below, that we may credit back to an eligible AWS Managed Services account.
- A **“Service Request”** is a request from a customer for information or advice.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for AWS Managed Services for the billing cycle in which the missed Service Commitment occurred in accordance with the Service Commitment & Credit Table and as further specified below:

- The Service Credit percent indicated in the SCCT may only be recovered once per monthly billing cycle for each Service Commitment.
- Separately reported Incidents that have the same Incident Resolution will be combined into one Incident for the purposes of calculating Service Credits. Service Credits will be due for the individual Incident that provides the highest Service Credits for the customer.

We will apply any Service Credits only against future AWS Managed Services payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card or other payment method you used to pay for the billing cycle in which the error occurred. Service Credits will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar (\$1 USD). Service Credits may be transferred or applied to any other account enrolled by you in AWS Managed Services. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any Unavailability, non-performance, or other failure by us to provide AWS Managed Services is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Notwithstanding the above, Service Credits may not individually or cumulatively exceed 30% of the total charges paid by you for AWS Managed Services for the billing cycle in which the missed Service Commitment (s) occurred.

Credit Request and Payment Process

To receive a Service Credit, you must submit a claim by opening a service request in the AWS Managed Services Interface. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the missed Service Commitment occurred and must include:

1. the words "SLA Credit Request" in the subject line;
2. the dates and times of each missed Service Commitment that you are claiming; and
3. logs and other documents that corroborate your claim (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

Once we review your Service Credit Request and confirm your eligibility, we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

SLA Exclusions

The Service Commitments do not apply to any Unavailability, suspension, or termination of AWS Managed Services, or any other AWS Managed Services performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of AWS Managed Services; (iii) that result from any actions or inactions of you or any third party, including your decision to postpone or not to authorize AWS Managed Services to perform or implement a change, update, patch, or other action recommended by AWS Managed Services; (iv) that result from your equipment, software, or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); (v) arising from our suspension and termination of your right to use AWS Managed Services in accordance with the AWS Agreement; (vi) that result from resources developed using non-AWS Managed Services approved AMIs; (vii) that result from the Unavailability or degraded performance of AWS Service Offerings; (viii) that result from unauthorized use of account credentials by you or any third party. SLAs are not applicable once off-boarding assistance commences as described in the AWS Managed Services Service Description. If availability is impacted by factors other than those included herein, then we may issue a Service Credit considering such factors at our discretion.