
Amazon Mechanical Turk

Developer Guide

API Version 2017-01-17



Amazon Mechanical Turk: Developer Guide

Copyright © Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is Amazon Mechanical Turk?	1
Mechanical Turk marketplace	1
Marketplace rules	2
The sandbox marketplace	2
Tasks that work well on Mechanical Turk	3
Tasks can be completed within a web browser	3
Work can be broken into distinct, bite-sized tasks	3
Task supports clear instructions and outcomes	3
Examples of common uses of Mechanical Turk	3
Core concepts	4
Requesters and workers	4
Marketplace	4
Task or HIT	4
Assignment	5
Reward and bonus	5
Qualifications	5
Best practices	5
Allow workers to be as efficient as possible	5
Build tasks with family and friends in mind	5
Include an optional feedback field	5
Test your HITs	6
Start small	6
Keep HIT type attributes consistent	6
Specify that links open new browser, windows or tabs	7
Limit your use of worker blocks	7
Include clear reasons for rejections and blocks	7
Frequently asked questions	7
Why aren't my tasks being completed?	7
How do I pull down HITs I created by mistake?	8
I expired my HITs. Why am I still getting submissions from workers?	8
Why are some of my task fields missing from my results?	8
Can I make some fields in my task interface required?	8
How can I test my task interface?	8
What is the difference between a HIT and an assignment?	8
Can I view the HITs I create with the API in the requester website?	9
I published HITs in the sandbox environment. Why they aren't being completed?	9
I incorrectly rejected some assignments. Can I reverse the rejection?	9
How do I filter the workers eligible to work on my task?	9
How do I create a custom qualification?	9
Can I restrict how many HITs a worker can complete for my project?	9
Can I post HITs in languages other than English?	9
Additional Mechanical Turk Resources	10
Set up Mechanical Turk	11
Step 1: Create a Mechanical Turk account	11
Step 2: Link your AWS account to your Mechanical Turk requester account	11
Step 3: Select a payment option	12
Step 4: Get an AWS access key	12
Step 5: Configure Your Credentials	13
Step 6: Set up the developer sandbox	13
Access Mechanical Turk	14
Use the Mechanical Turk Requester UI	14
Use the Mechanical Turk API	14
Use the AWS CLI	14
Download and configure the AWS CLI	15

Use the AWS CLI with Mechanical Turk	15
Get Started with Mechanical Turk	17
Prerequisites	17
Step 1: Create a task	17
Question definition	17
Task attributes	18
Post the task	18
Step 2: Check task status	20
Step 3: Retrieve results	23
Step 4: Approve Assignments	27
Creating and managing tasks (HITs)	29
Define questions	29
Use HTML to define questions	29
Question definitions	33
Requester website layouts	35
HIT attributes	35
Creating HITs	37
Create a HIT Directly	37
Create a HIT using a HIT Type	37
Modifying HITs	38
Modify expiration time	38
Add additional assignments	39
Modify the HIT Type	39
Delete	39
Using the sandbox	39
API calls using the CLI	39
API calls using the Python SDK (boto3)	40
Testing HITs	40
HIT references using RequesterAnnotation	40
Retrieving results	41
Assignment attributes	41
Assignment answer	41
Retrieving HIT status	44
Managing workers	45
Approving and rejecting work	45
Awarding a bonus	45
Blocking workers	46
Selecting eligible workers	46
Qualifications and qualification types	46
Qualification requirements	47
System qualification types	48
Working with custom qualification types	51
Create a qualification type	52
Assign or remove a worker qualification	52
Qualification requests	52
Tutorial: Require workers to be in a group	53
Tutorial: Require workers to meet accuracy level	53
Tutorial: Exclude workers from selecting tasks	54
Communicate with workers	55
Use Mechanical Turk notifications	56
Notification event types	56
Notification destination	56
Handling notifications using AWS Lambda	57
Sending test events	59
HIT references using requester annotation	60
Use request tokens	61
CORS configuration requirement	62

What is Amazon Mechanical Turk?

Amazon Mechanical Turk (Mechanical Turk) is a crowdsourcing marketplace that connects you with an on-demand, scalable, human workforce to complete tasks. Using Mechanical Turk, you can programmatically direct tasks to the Mechanical Turk marketplace, where they can be completed by workers around the world. Mechanical Turk allows you to access the intelligence, skills, and insights of a global workforce for tasks as varied as data categorization, moderation, data collection and analysis, behavioral studies, and image annotation.

Mechanical Turk is built around the concept of *microtasks*, which are small, atomic tasks that workers can complete in their web browser. When you submit work to Mechanical Turk, you typically start by breaking it into smaller tasks on which workers can work independently. In this way, a project involving categorizing 10,000 images becomes 10,000 individual microtasks that workers can complete. By breaking tasks down atomically, hundreds of workers can work on portions of your project at the same time, which increases how quickly the work can be completed. In addition, you can specify that each task be completed by multiple workers to allow you to check for quality or identify biases in subjective questions.

Important

If you do not add a CORS configuration to the Amazon S3 buckets that contain your image input data, HITs that you create using those input images will fail. To learn more, see [CORS configuration requirement \(p. 62\)](#).

Use this guide to learn how you can interact with Mechanical Turk programmatically. We recommend you begin by reading the following topics. To get started quickly with Mechanical Turk, see [Get Started with Amazon Mechanical Turk \(p. 17\)](#).

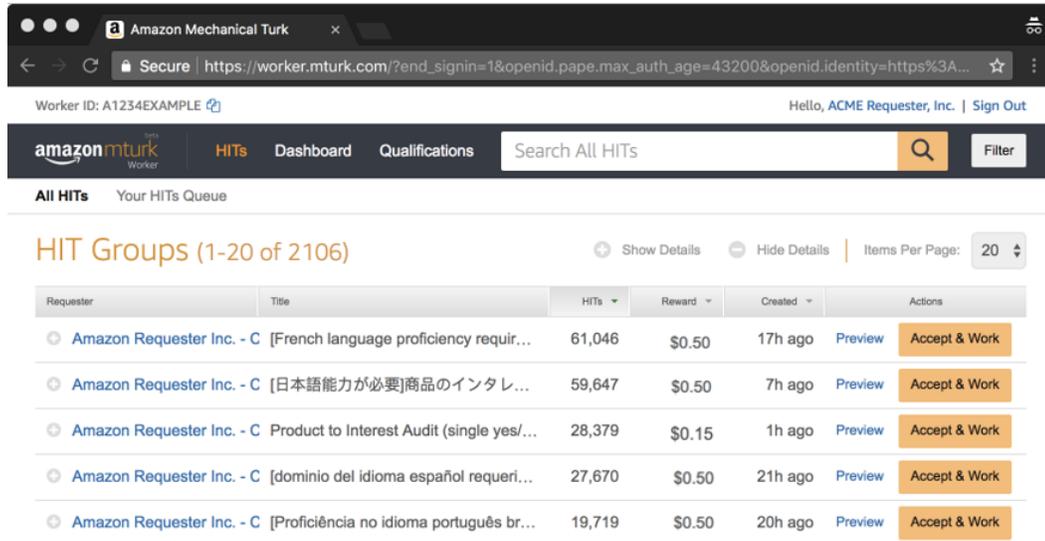
Topics

- [The Amazon Mechanical Turk marketplace \(p. 1\)](#)
- [Creating tasks that work well on Amazon Mechanical Turk \(p. 3\)](#)
- [Amazon Mechanical Turk core concepts \(p. 4\)](#)
- [Amazon Mechanical Turk best practices \(p. 5\)](#)
- [Frequently asked questions \(p. 7\)](#)

The Amazon Mechanical Turk marketplace

Mechanical Turk uses the *requester* and *worker* terms to describe the two participants in the marketplace. When you post new tasks to Amazon Mechanical Turk (Mechanical Turk), you are a *requester* asking *workers* to complete your tasks in exchange for the reward amount you offer. Workers can go to the [Mechanical Turk marketplace](#) to find and accept tasks.

As shown in the following image of the marketplace website, workers can see a list of available tasks, along with details about each task. Workers can review the title and description, reward amount, and time allotted to complete each task before accepting and working on it. In many cases, workers preview a task prior to accepting it, which allows them to decide if they want to work on it.



Submitting tasks to the Mechanical Turk marketplace does not guarantee that workers will complete them. If workers don't believe that the reward amount is reasonable for the effort required, or the work isn't something on which they want to work, they skip it and move on to other tasks. For this reason, we recommended that you put thought into how you describe your task so that workers can make an informed decision.

After workers complete your task, they submit their response and move on to additional tasks that you've posted or tasks from other requesters. You can review a worker's submission shortly after they submit the task. You have the option to *approve* or *reject* their submission. If you approve the work, the reward amount is distributed to the worker. Note that if you neither approve nor reject a task submission, it is automatically approved after a set time.

Marketplace rules

Prior to submitting tasks to Mechanical Turk, you should review the [Acceptable Use Policy](#) to ensure that your task adheres to the rules of the marketplace. Prohibited uses cover a range of activities such as violating the privacy or security of workers or others, abusive behavior, or any illegal activities. Violating these policies results in removal of your tasks from the Mechanical Turk marketplace and may result in the suspension of your account.

The sandbox marketplace

To experiment with Mechanical Turk without spending money on the Mechanical Turk marketplace, you can use the [sandbox environment for requestors](#) and [the one for workers](#). This is a mirror image of the *production* environment, but no money changes hands when work is completed. Many requesters create tasks here first and complete them themselves so that they can validate their task interface and ensure they get the results they expect back. You can find more information on using the sandbox in [Using the sandbox \(p. 39\)](#).

Note that there is no financial incentive to complete work in the sandbox marketplace, so you shouldn't expect tasks you post in the sandbox to be completed unless you do so yourself.

Creating tasks that work well on Amazon Mechanical Turk

Amazon Mechanical Turk (Mechanical Turk) can be used for an exceptionally wide range of tasks. Tasks that work well on Mechanical Turk generally meet the following criteria:

- Can be completed from within a web browser
- Can be broken into distinct, bite-sized tasks
- Can support clear instructions and outcomes

Most tasks that meet these criteria can be completed on Mechanical Turk, assuming you provide workers with a task interface that allows them to successfully perform the task. You should also keep in mind that Mechanical Turk workers excel at tasks that rely on general human knowledge and skills. While some workers have specialized experience such as legal or medical backgrounds, most do not. As a result, while Mechanical Turk can enable tasks such as labeling the location of people or animals in images, you are likely to have less success asking workers to apply expertise that would be associated with a radiologist.

Note that tasks must also conform to the rules in the [Mechanical Turk Acceptable Use Policy](#). Prohibited uses cover a range of activities such as violating the privacy or security of workers or others, abusive behavior, or any illegal activities.

Tasks can be completed within a web browser

Mechanical Turk tasks are built using HTML and presented to workers via the Mechanical Turk website. Most workers complete tasks on their computer without the need to use other devices or specialized software. Tasks that require workers to visit physical locations or leverage other devices aren't recommended.

Work can be broken into distinct, bite-sized tasks

Most Mechanical Turk tasks take less than five minutes to complete and almost all can be completed within an hour. This lets workers try new tasks without needing to commit a lot of time. Most workers appreciate the flexibility that Mechanical Turk provides in moving from task to-task without being locked in for an extended period of time.

Task supports clear instructions and outcomes

The most successful tasks on Mechanical Turk are those that provide the necessary information for a worker to imagine what a successful response would look like. Avoid tasks that are open-ended and could have multiple possible outcomes. For example, a task that asks workers to *identify all of the competitors of company X* would be frustrating for workers. By specifying that you want *all* competitors, workers are left wondering at what point they should draw a line and stop their research. It would also leave them wondering if you will reject their work if they aren't as comprehensive as you want them to be. In this example, you should instead be specific about the data that you need by describing your task as *identify the top 5 competitors of company X*.

Examples of common uses of Mechanical Turk

The following are examples of common Mechanical Turk use-cases:

- *Audio transcription*: Transcribe an audio clip.
- *Categorization*: Categorize products.
- *Data collection*: Identify the website for a business.
- *Writing*: Write a description of a product based on an image and details.
- *Market research*: Complete a market research survey.
- *Rating*: Evaluate and rate the quality of an image.
- *Usability testing*: Visit a website and complete a set of steps, providing feedback on each step.
- *Research study*: Participate in a study by responding to questions surrounding a scenario.
- *Computer vision*: Draw bounding boxes around animals in images.
- *Natural language processing*: Identify the named entities within a statement.
- *Matching*: Review two data records and confirm they relate to the same business.
- *Moderation*: Evaluate a set of images and identify any that don't meet the provided criteria.
- *Ranking*: Rank a list of products based on their relevance to a search query.
- *Data extraction*: Extract the names and prices of products in a receipt.
- *Text transcription*: Transcribe handwritten text.
- *Video transcription*: Transcribe a video clip.

Amazon Mechanical Turk core concepts

The following are the core concepts of Amazon Mechanical Turk (Mechanical Turk) that you need to understand to use it effectively.

Requesters and workers

A *requester* is a company, organization, or person that posts tasks (HITs) to Mechanical Turk for workers to perform. A *worker* is a person who performs the tasks specified by a requester in a HIT.

Marketplace

The [Mechanical Turk marketplace](#) is where workers can go to find and accept tasks. In addition to the *production* marketplace, there is a second [sandbox marketplace](#) where requesters can post development tasks without money changing hands.

More information can be found in Amazon Mechanical Turk marketplace.

Task or HIT

The base unit of work in Mechanical Turk is called a *Human Intelligence Task*, which is typically designated as a *HIT* or *task*. A HIT represents a single, self-contained task, such as *Identify the color of the car in the photo*, that a requester submits to Mechanical Turk for workers to complete.

Mechanical Turk is built around the concept of *microtasks*, which are small, atomic tasks that workers can complete in their web browser. When you submit work to Mechanical Turk, you typically start by breaking it into smaller tasks on which workers can work independently. In this way, a project involving categorizing 10,000 images becomes 10,000 individual microtasks that workers can complete. Hundreds of workers can work on portions of your project at the same time, which increases how quickly the work can be completed. In addition, you can specify that each task be completed by multiple workers to allow you to check for quality or identify biases in subjective questions.

Assignment

When creating a HIT, you can specify how many workers can accept and complete each task. Doing so allows you to collect multiple responses for each item and then compare them. This additional information can be valuable in managing quality, as well as in collecting multiple data points when responses are subjective.

When a worker accepts a HIT, Mechanical Turk creates an *assignment*, which belongs exclusively to the worker. The worker can submit results up until the expiration of the HIT. When retrieving results for a HIT, requesters retrieve all of the submitted assignments.

Reward and bonus

A *reward* is the money you, as a requester, pay workers for satisfactory work they do on your HITs. A *bonus* is the amount you award workers for high-quality performance. Rewards are transmitted to workers when assignment submissions are approved, either by approving the assignment or when the auto-approval threshold is reached. Bonuses can be sent to workers who have recently completed an assignment for you.

Qualifications

You can use *qualifications* to specify attributes of the workers eligible to work on your HITs. Qualifications can be either system-generated, such as qualifications based on location, or managed by you, based on past performance on your tasks.

To learn more, see [Selecting eligible workers \(p. 46\)](#).

Amazon Mechanical Turk best practices

Keep the following best practices in mind when you design and create your HITs.

Allow workers to be as efficient as possible

When you post tasks to Mechanical Turk, the reward amount you set is primarily for the worker's time and attention to your task. If your task interface is inefficient and requires multiple manual steps that require a lot of time, workers typically expect a higher reward amount to compensate for the time they need to spend performing those steps. Investing time to make your interface as efficient as possible pays dividends in higher accuracy and lower costs.

Build tasks with family and friends in mind

When building tasks, it's a common mistake to assume that workers have the same knowledge you do about your area of expertise. Very few workers have the expertise you do and will likely be confused if you use highly technical language or make assumptions about their skills. A great practice is to design your task interface with a member of your family or a friend in mind. Could they complete your task successfully? If you're not sure, share the interface with them and see if they can complete it without any additional instructions from you.

Include an optional feedback field

Whenever possible, include an optional feedback field at the end of your task interface, particularly when working with a new interface. Workers appreciate the opportunity to provide feedback and often share insights on how to improve it.

Test your HITs

Before posting your tasks to Mechanical Turk, it is always a good idea to take a few minutes to test your HITs to make sure they work as you expect. It allows you to validate that your interface does what you expect. Doing the task yourself also lets you get an idea of how long it takes to complete so that you can set an appropriate reward amount.

The easiest way to test your task interface is to save it to an HTML file and open it in a browser. From the browser, you can go through all of the steps that a worker would follow in completing the task. If your task interface is built around a standard form element, you won't be able to test submitting it, but can test to ensure it works as you expect. If you use the crowd-form element from [Crowd HTML Elements \(p. 31\)](#), you can test it by selecting **Submit**. When you submit from outside of Mechanical Turk, the results are displayed at the top of the window.

To fully test a task interface and the creation and retrieval of HITs, you can use the sandbox environment.

Start small

When you create or update a task interface, it's always best to start by posting a small number of HITs first to confirm that workers complete the task as you expect. It's a great way to understand how workers respond and gives you a chance to correct any issues before you post the remaining work. Nothing is worse than posting thousands of dollars of HITs, only to discover that the results are invalid because you made a mistake in your task interface.

Keep HIT type attributes consistent

When you create a HIT, you provide a number of attributes about the task that tell Mechanical Turk how to display it in the marketplace. These are separate from the content and question of the task itself, and include the title, description, reward amount, and attributes describing how long the task remains active. These attributes comprise the HIT type for your task. Mechanical Turk automatically creates a HIT type when you first call `CreateHIT` with those values. When you create multiple HITs, Mechanical Turk attempts to find an existing HIT type in your account that has the same attributes and reuse it. If you change any of these attributes—even if they are small changes to the title or description—it will force Mechanical Turk to create a new HIT type with each change.

Maintaining consistent attributes for your HIT type is important because it directly impacts how your HIT is displayed on the worker website. On the worker website, HITs are grouped together into HIT groups based on their HIT type values. As shown in the following image, each HIT group has thousands of HITs on which a worker can work because they all have the same attributes for title, description, reward, and other attributes. If workers accept a HIT from one of these HIT groups, they can automatically move to the next piece of work in the HIT group without needing to return to the list.

HIT Groups (1-20 of 907) Show Details Hide Details Items Per Page: 20

Requester	Title	HITs	Reward	Created	Actions
...	Find and submit data for high school clubs based on a list of high schools	16,504	...	4d ago	Preview Accept & Work
...	Market Research Survey	15,314	...	2h ago	Preview Accept & Work
...	Find a video advertisement on Youtube from metadata	9,852	...	8h ago	Preview Accept & Work

Description
We will provide you videos of advertisements as well as metadata of the advertisements, such as, campaign name, company name, release date, etc... We ask you to find the videos of the advertisements in the dataset on Youtube and provide us their links.

Time Allotted
5 Min
Expires in 7d

Qualifications Required
✔ HIT approval rate (%) is greater than 75

Your Values
100

If, however, each HIT has a unique HIT type, then workers see your HITs as a long list of options in the list and have to return to the list after completing each task.

Requester	Title	HITS	Reward	Created	Actions
Ad Tagger	Identify the brands featured in banner ads - 2020-10-26 21:19:18.	1	\$0.60	5d ago	Preview Quality
Ad Tagger	Identify the brands featured in banner ads - 2020-10-26 21:19:40.	1	\$0.60	5d ago	Preview Quality
Ad Tagger	Identify the brands featured in banner ads - 2020-10-26 21:19:54.	1	\$0.60	5d ago	Preview Quality
Ad Tagger	Identify the brands featured in banner ads - 2020-10-26 21:21:23.	1	\$0.60	5d ago	Preview Quality
Ad Tagger	Identify the brands featured in banner ads - 2020-10-26 21:21:49.	1	\$0.60	5d ago	Preview Quality

Specify that links open new browser, windows or tabs

When you add links to your task HTML, you should include a *target* attribute to let the browser know that it should open a new window or tab when workers click on it. This keeps the worker interface active in the existing window and prevents issues that sometimes occur when workers use the **Back** button to return to the worker interface. Add the `_blank` target to direct the browser to open a new window, as shown in the following example.

```
<a href="https://www.amazon.com" target="_blank">My link</a>
```

Limit your use of worker blocks

We recommend that you be judicious in your use of worker blocks and only block those workers who are clearly not making an attempt to correctly respond to your task (spamming). If a worker is simply misreading instructions or lacks the requisite skills to complete your task successfully, we advise you to use a custom qualification to exclude them from future tasks, rather than a block. Because the blocks a worker receives are a component of Mechanical Turk worker review policies, and frequent blocks may result in account suspension, workers are sensitive to being blocked by requesters. If the worker community believes that you are blocking workers unfairly, they may choose to avoid accepting your tasks in the future.

Include clear reasons for rejections and blocks

Workers take a lot of pride in the quality of their work and pay close attention to rejections and blocks they receive. When you decide to reject an assignment or block a worker, be as clear as possible about the reasons for the action. Simply providing a value such *incorrect* as the reason gives the worker no information they can use to improve in the future. Instead, be clear about what the worker did incorrectly. This allows workers to correct their mistakes in future tasks.

Frequently asked questions

Use the following sections to get answers to frequently asked questions. If you need additional support, use the following link to contact Amazon Mechanical Turk: www.mturk.com/contact-us.

Why aren't my tasks being completed?

There are a number of reasons why the tasks you post to Mechanical Turk aren't being completed. The most common reason is that the reward amount you specified isn't adequate to compensate workers for the time and effort they need to commit to your task to complete it. If you suspect this is the case, remove the HITS from Mechanical Turk by expiring them and experiment with reposting some of them at a higher reward amount.

Other common reasons include the following.

- The qualification requirements for the task are so narrow that few, if any, workers meet the criteria to be eligible for the task.

- The task interface has a technical issue that prevents workers from submitting it.
- The assignment duration is set too short for workers to successfully complete the task in the time allowed.

How do I pull down HITs I created by mistake?

Use the [UpdateExpirationForHIT](#) operation and set the `ExpiresAt` time to 0 to tell Mechanical Turk to immediately expire a HIT. Note that this won't prevent workers that have already accepted your HIT from completing and submitting it.

I expired my HITs. Why am I still getting submissions from workers?

If a worker accepts a HIT before it expires, they are still allowed to complete and submit the task until the assignment duration elapses. This protects the worker experience by letting them submit work on which they may have already spent a lot of time, even if you opt to take the HIT down.

Why are some of my task fields missing from my results?

A common mistake in building task interfaces is using the same *name* attribute for multiple form inputs. In those cases, only one of the input field values is returned. You should check your HTML to ensure that each input has a unique name.

Can I make some fields in my task interface required?

You can use HTML, JavaScript, or both to specify required fields and minimum or maximum values or perform other validations that prevent workers from submitting the task if it doesn't meet the requirements. To learn more about the types of form validation you can apply, see [Client-side form validation](#) on the Mozilla developer site.

How can I test my task interface?

The easiest way to test your task interface is to save it to an HTML file and open it in a browser. From the browser, you can go through all of the steps that a worker would perform in completing the task. If your task interface is built around a standard form element, you can't test submitting it, but you can test to ensure it works as you expect. If you use the `crowd-form` element from Crowd HTML Elements, you can test it by selecting **Submit**. When you submit from outside of Mechanical Turk, the results are displayed at the top of the window.

To fully test a task interface and the creation and retrieval of HITs, you can use the sandbox environment.

What is the difference between a HIT and an assignment?

A *HIT* is a single task that you create in Mechanical Turk. When workers accept a HIT, they get an *assignment* that gives them the right to submit their response. When you create a HIT, you can specify the maximum number of assignments that can be created for each HIT, which allows you to get multiple different worker responses for each task. For more information on HITs and assignments, see [Amazon Mechanical Turk core concepts](#) (p. 4).

Can I view the HITs I create with the API in the requester website?

No, the requester website only displays HITs that are created from the requester website.

I published HITs in the sandbox environment. Why they aren't being completed?

The sandbox environment is a great way to test HITs without spending any money. However, because no money changes hands, there isn't any incentive for workers to complete your tasks. To complete your testing, create an account in the [worker sandbox environment](#) to complete the tasks yourself. Then, publish them in the *production* environment.

I incorrectly rejected some assignments. Can I reverse the rejection?

In the event that you reject an assignment but then discover that the issue was not the worker's fault, you can call [ApproveAssignment](#) to reverse the rejection, but only for assignments submitted in the last 30 days that haven't been deleted.

How do I filter the workers eligible to work on my task?

Mechanical Turk provides a qualifications system that allows you to use system-managed or custom criteria to limit the workers that can work on a task. For more information, see [Selecting eligible workers](#) (p. 46).

How do I create a custom qualification?

You can create custom qualification types that allow you to filter workers eligible to work on your tasks using criteria based on their past performance on your tasks. For more information, see [Working with custom qualification types](#) (p. 51).

Can I restrict how many HITs a worker can complete for my project?

Mechanical Turk doesn't provide a native capability to limit the number of HITs that a worker can contribute to a project or batch. To learn how to accomplish this using custom qualification types, see [Working with custom qualification types](#) (p. 51). Before starting your project or batch, create a custom qualification type with a label such as **Completed Enough of Project A** and specify that this type doesn't exist (**DoesNotExist**) in your qualification requirements for each HIT. When a worker reaches your threshold of HITs they can submit, you can assign this qualification type to them, after which they can't accept any HITs for the project.

Can I post HITs in languages other than English?

You can post HITs using any language, provided you note the required language in the title of your task. However, the number of available workers who are fluent in a given language varies greatly. It may take

longer for your task to be completed or you may need to increase your reward amount if not enough workers are available.

Additional Mechanical Turk Resources

- The [Mechanical Turk API Reference](#) describes all the API operations for Mechanical Turk in detail.
- The [Mechanical Turk Requester User Interface Documentation](#) describes how to create Mechanical Turk tasks using a graphical user interface.
- Posts on the [Mechanical Turk Happenings Blog](#) address updates to the Mechanical Turk marketplace.
- [Blog Tutorials](#) provide instruction on using Mechanical Turk for a variety of tasks.
- The [Amazon Mechanical Turk Developer Forums](#) provide questions and answers about Mechanical Turk.
- [Mechanical Turk on Github](#) offers sample code and tutorials.

Set up Amazon Mechanical Turk

Use the following topics to learn how to use Amazon Mechanical Turk (Mechanical Turk) with APIs or AWS command line tools.

If you plan to interact with Mechanical Turk only through the Mechanical Turk requester user interface, you can skip these steps and instead follow the *Getting Started* steps described in the [Requester UI Guide](#).

Topics

- [Step 1: Create a Mechanical Turk account](#) (p. 11)
- [Step 2: Link your AWS account to your Mechanical Turk requester account](#) (p. 11)
- [Step 3: Select a payment option](#) (p. 12)
- [Step 4: Get an AWS access key](#) (p. 12)
- [Step 5: Configure Your Credentials](#) (p. 13)
- [Step 6: Set up the developer sandbox](#) (p. 13)

Step 1: Create a Mechanical Turk account

To create an Amazon Mechanical Turk account, go to the [Amazon Mechanical Turk Requester](#) website, choose **Create an account**, and follow the on-screen instructions.

Note that Mechanical Turk accounts use the same login credentials and profiles as Amazon retail websites such as [Amazon.com](#). Changes in the name or address on your account, on either Amazon.com or Mechanical Turk, are reflected in both locations.

To use Mechanical Turk programmatically, you must have an AWS account. If you don't already have an account, you are prompted to create one when you sign up. You're not charged for any AWS services that you sign up for unless you use them.

To create an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

Note your AWS account ID. You need it for the next step.

Step 2: Link your AWS account to your Mechanical Turk requester account

You need to link your AWS account to your Mechanical Turk requester account. This operation grants permission to your AWS account to access your requester account using the Mechanical Turk APIs.

- Go to <https://requester.mturk.com/developer/>.
- Choose **Link your AWS Account** and sign in with your AWS root user email address and password.

Step 3: Select a payment option

Before you can post HITs to the Mechanical Turk marketplace, you need to enable AWS Billing for your account to pay worker rewards and Mechanical Turk fees. These appear on the AWS Anniversary Bill for your linked AWS account.

Alternatively, you can prepay for the HITs you plan to create using a credit card payment.

To enable AWS Billing or prepay for HITs, go to the account section of the [Requester website](#).

Step 4: Get an AWS access key

Before you can access Mechanical Turk programmatically, you must have an AWS access key. Access keys consist of an access key ID and secret access key, which are used to sign programmatic requests that you make to AWS. If you don't have access keys, you can create them from the AWS Management Console. As a best practice, do not use the AWS account root user access keys for any task where they are not required. Instead, create a new administrator IAM user with access keys for yourself. To learn how, see [Creating your first IAM admin user and group](#) in the IAM User Guide. If you do not wish to grant administrator access to this account, you can choose either the `AmazonMechanicalTurkFullAccess` or `AmazonMechanicalTurkReadOnly` policy rather than `AdministratorAccess` when you attach a policy to the user.

The only time that you can view or download the secret access key is when you create the keys. You cannot recover them later. However, you can create new access keys at any time. You must also have permissions to perform the required IAM actions. For more information, see [Permissions Required to Access IAM Resources](#) in the *IAM User Guide*.

To create access keys for an IAM user:

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Users**.
3. Choose the name of the user whose access keys you want to create, and then choose the **Security credentials** tab.
4. In the **Access keys** section, choose **Create access key**.
5. To view the new access key pair, choose **Show**. You will not have access to the secret access key again after this dialog box closes. Your credentials should resemble the following example:

Access key ID: AKIAIOSFODNN7EXAMPLE

Secret access key: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY

6. To download the key pair, choose **Download .csv file**. Store the keys in a secure location. You will not have access to the secret access key again after this dialog box closes.

Keep the keys confidential in order to protect your AWS account. Never email them. Do not share them outside your organization, even if an inquiry appears to come from AWS or Amazon.com. No one who legitimately represents Amazon will ever ask you for your secret key.

- After you download the `.csv` file, choose **Close**. When you create an access key, the key pair is active by default, and you can use the pair right away.

Related topics

- [What Is IAM?](#) in the *IAM User Guide*
- [AWS Security Credentials](#) in *AWS General Reference*

Step 5: Configure Your Credentials

To access Mechanical Turk programmatically, you must configure your credentials to enable authorization for your applications.

There are several ways to do this. For example, you can manually create the credentials file to store your access key ID and secret access key. You also can use the `aws configure` command of the AWS CLI to automatically create the file. Alternatively, you can use environment variables. For more information about configuring your credentials, see the programming language-specific [AWS SDK developer guide](#).

The Mechanical Turk API endpoint is only available in the `us-east-1` Region so it is recommended that you configure your default Region as `us-east-1`. If you primarily work with a different default AWS Region, you can specify the `us-east-1` Region and endpoint as part of your CLI or SDK requests to Mechanical Turk.

To install and configure the AWS CLI, see [Installing, updating, and uninstalling the AWS CLI](#) and [Configuring the AWS CLI](#) in the *IAM User Guide*, respectively.

Step 6: Set up the developer sandbox

You may wish to test your HITs in the Amazon Mechanical Turk sandbox testing environment to make sure they work as expected before publishing them in the Mechanical Turk marketplace. The sandbox is an environment where you can publish and work on HITs at no cost before publishing them in the *production* Mechanical Turk marketplace. The sandbox consists of a [requester sandbox website](#) and a [worker sandbox website](#).

Create a requester account on the requester sandbox website, which is located at <https://requestersandbox.mturk.com>. This follows the same procedure as creating a Mechanical Turk account described in [Step 1: Create a Mechanical Turk account \(p. 11\)](#). You can use the same email address and account if you wish.

You also need to create a worker account on the worker sandbox website located at <https://workersandbox.mturk.com> to view your sandbox HITs as a worker. There is no charge for using the Mechanical Turk sandbox sites.

To create HITs in the sandbox using the Mechanical Turk APIs, you also need to link your AWS account to your sandbox requester account, as described in [Step 2: Link your AWS account to your Mechanical Turk requester account \(p. 11\)](#), on the [requester sandbox website](#).

To configure the AWS CLI or SDKs to access the sandbox instead of the production environment, you must set the API endpoint to be <https://mturk-requester-sandbox.us-east-1.amazonaws.com>. Refer to the [AWS CLI Command Reference](#) or [SDK documentation](#) for how best to do this.

Access Amazon Mechanical Turk

You can access Amazon Mechanical Turk (Mechanical Turk) using the Mechanical Turk requester user interface, the AWS Command Line Interface (AWS CLI), or the Mechanical Turk API.

Topics

- [Use the Mechanical Turk Requester UI \(p. 14\)](#)
- [Use the Mechanical Turk API \(p. 14\)](#)
- [Use the AWS CLI \(p. 14\)](#)

Use the Mechanical Turk Requester UI

The Mechanical Turk Requester User Interface (RUI) provides access to Mechanical Turk functionality using a graphical user interface. You can use the Mechanical Turk RUI to:

- Create projects that define tasks you want workers to complete
- Submit batches of tasks to the Mechanical Turk marketplace
- Monitor and review the results of batches
- Approve and reject task submissions
- Manage workers
- Create and modify qualification types

For more information on using the RUI, visit the [Amazon Mechanical Turk Requester UI Guide](#).

HITs created from the API or AWS CLI cannot be viewed from the Requester UI.

Use the Mechanical Turk API

The AWS SDKs provide broad support for Mechanical Turk in [Java](#), [JavaScript in the browser](#), [.NET](#), [Node.js](#), [PHP](#), [Python](#), [Ruby](#), [C++](#), and [Go](#). To get started quickly with many of these languages, see the [Amazon Mechanical Turk Code Samples](#).

Before you can use the AWS SDKs with Mechanical Turk, you must get an AWS access key ID and secret access key. For more information, see [Set up Amazon Mechanical Turk \(p. 11\)](#).

Use the AWS CLI

You can use the AWS Command Line Interface (AWS CLI) to control multiple AWS services from the command line and automate them through scripts. This includes posting Mechanical Turk HITs and retrieving results, either on an ad-hoc basis or within utility scripts.

Before you can use the AWS CLI with Mechanical Turk, you must get an access key ID and secret access key. For more information, see [Get an AWS access key](#).

For a complete listing of all the commands available for Mechanical Turk in the AWS CLI, see the [AWS CLI Command Reference](#).

Download and configure the AWS CLI

The AWS CLI is available at <https://aws.amazon.com/cli>. It runs on Windows, MacOS, or Linux. After you download the AWS CLI, follow these steps to install and configure it:

- Go to the [AWS Command Line Interface User Guide](#).
- Follow the instructions for [Installing the AWS CLI](#) and [Configuring the AWS CLI](#).

Use the AWS CLI with Mechanical Turk

The command line format consists of a Mechanical Turk operation name followed by the parameters for that operation. The AWS CLI supports a shorthand syntax for the parameter values, as well as JSON.

For example, the following command returns a list of the HITs that have been created for in your account.

```
$ aws mturk list-hits
```

The next command creates a new HIT in the Mechanical Turk marketplace (for easier readability, long commands in this section are broken into separate lines).

Important

Creating the following HIT results in a charge of \$0.12 to your account.

```
$ aws mturk create-hit \  
  --title "Describe the weather" \  
  --description "Describe the current weather where you live" \  
  --reward "0.10" \  
  --lifetime-in-seconds 14400 \  
  --assignment-duration-in-seconds 300 \  
  --question '<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/  
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">  
    <HTMLContent><![CDATA[  
      <!DOCTYPE html>  
      <script src="https://assets.crowd.aws/crowd-html-elements.js"></script>  
      <crowd-form>  
        <p>Describe the current weather where you live</p>  
        <p><textarea name="weather" cols="80" rows="3"></textarea></p>  
      </crowd-form>  
    ]]>  
    </HTMLContent>  
    <FrameHeight>0</FrameHeight>  
</HTMLQuestion>'
```

This returns the attributes of the HIT you created. After the task is completed, you can retrieve the results by using the [ListAssignmentsForHIT](#) operation.

```
$ aws mturk list-assignments-for-hit \  
  --hit-id <The HITId from the previous step>
```

On the command line, it can be difficult to compose valid HTML for your task when calling [CreateHIT](#). However, the AWS CLI can read XML files. For example, consider the following, which creates a HIT using a question stored in a file *question.xml*.

```
$ aws mturk create-hit \  
  --title "Describe the weather" \  
  --description "Describe the current weather where you live" \  
  --reward "0.10" \  
  --lifetime-in-seconds 14400 \  
  --assignment-duration-in-seconds 300 \  
  --question file://question.xml
```

Get Started with Amazon Mechanical Turk

Use the hands-on tutorials in this section to help you get started and learn more about Amazon Mechanical Turk.

Topics

- [Prerequisites \(p. 17\)](#)
- [Step 1: Create a task \(p. 17\)](#)
- [Step 2: Check task status \(p. 20\)](#)
- [Step 3: Retrieve results \(p. 23\)](#)
- [Step 4: Approve Assignments \(p. 27\)](#)

Prerequisites

Before you begin, you should familiarize yourself with the basic concepts in Amazon Mechanical Turk. For more information, see [The Amazon Mechanical Turk marketplace \(p. 1\)](#) and [Amazon Mechanical Turk core concepts \(p. 4\)](#).

Additional, complete the steps in [Set up Amazon Mechanical Turk \(p. 11\)](#) before completing this tutorial.

Step 1: Create a task

In this step we create a task in Mechanical Turk that asks workers to describe the current weather where they live. The task interface for this will be created using the [HTMLQuestion](#) data structure and we'll make use of [Crowd HTML Elements](#) to simplify the task HTML.

Important

Completing the steps in this tutorial results in a charge of \$0.60 to your account.

Question definition

The most common way to define tasks in Mechanical Turk is using the `HTMLQuestion` data structure, which is defined as XML that encapsulates the HTML that is displayed to the worker. For this task, we use the following definition.

```
<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">
  <HTMLContent><![CDATA[
    <!DOCTYPE html>
    <script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
    <crowd-form>
      <p>Describe the current weather where you live.</p>
      <p><textarea name="weather" cols="80" rows="3"></textarea></p>
    </crowd-form>
  ]]>
```

```
</HTMLContent>
  <FrameHeight>0</FrameHeight>
</HTMLQuestion>
```

Note that the HTML includes a reference to the `crowd-html-elements.js` library which includes the `crowd-form` element. We use the `crowd-form` element in place of the standard `form` element because it removes the need to specify the endpoint for the form to submit results. It also automatically appends a **Submit** button if one isn't present. More information about this library can be found in [Crowd HTML Elements](#).

We've also set the value of `FrameHeight` to zero, which directs the marketplace website to render the task interface using the full browser window.

Task attributes

Next, we can define the attributes for our task. We'll use the following attributes:

Attribute	Value
Title	Describe the weather
Description	Describe the current weather where you live
Reward	0.1
MaxAssignments	5
LifetimeInSeconds	14,400
AssignmentDurationInSeconds	300
AutoApprovalDelayInSeconds	259,200

Here, we give an accurate description of our task and indicate that we will reward each worker with 10 cents for each successful completion. In addition, we set the `MaxAssignments` to 5 to indicate that we would like to get five responses from different workers. Finally, we set the lifetime and assignment duration to four hours and five minutes respectively. Workers have five minutes to complete the assignment before it expires and becomes available to other workers. If, after four hours, we haven't yet gotten a response, the task is automatically removed from the Mechanical Turk marketplace. We've also set the `AutoApprovalDelay` at three days which means that an assignment is automatically approved after three days if we don't take any action to approve or reject it before then.

For more detail on the attributes that can be specified for a HIT, visit the [CreateHIT](#) documentation.

Post the task

You can post a task using the AWS CLI or a language-specific AWS SDK. Select a tab in the following table to see an example of how you can post a task using the AWS CLI and the AWS SDK for Python (Boto3).

AWS CLI

The following AWS CLI example creates a new task using `create-hit`.

```
$ aws mturk create-hit \
  --title "Describe the weather" \
```

```
--description "Describe the current weather where you live" \  
--reward "0.10" \  
--max-assignments 5 \  
--lifetime-in-seconds 14400 \  
--assignment-duration-in-seconds 300 \  
--auto-approval-delay-in-seconds 259200 \  
--question '<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/  
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">  
  <HTMLContent><![CDATA[  
    <!DOCTYPE html>  
    <script src="https://assets.crowd.aws/crowd-html-elements.js"></script>  
    <crowd-form>  
      <p>Describe the current weather where you live</p>  
      <p><textarea name="weather" cols="80" rows="3"></textarea></p>  
    </crowd-form>  
  ]]>  
</HTMLContent>  
<FrameHeight>0</FrameHeight>  
</HTMLQuestion>'
```

Using `create-hit` returns the following sample result.

```
{  
  "HIT": {  
    "HITId": "3TL87M08CLOFYXKXNRLMZ01MOK4FL5",  
    "HITTypeId": "3AKE04YHPN13791QQA6BU4GMS7CHZJ",  
    "HITGroupId": "367GCHJ5533R84AG2MXJ0OFCCJ7M9Q",  
    "CreationTime": "2020-09-29T14:30:03-07:00",  
    "Title": "Describe the weather",  
    "Description": "Describe the current weather where you live",  
    "Question": "<HTMLQuestion xmlns=\"http://mechanicalturk.amazonaws.com/  
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd\">\n<HTMLContent><![CDATA[\n<!DOCTYPE html>\n<script src=\"https://assets.crowd.aws/crowd-html-elements.js\"></  
script>\n<crowd-form>\n<p>Describe the current weather where you live</p>\n<p><textarea  
name=\"weather\" cols=\"80\" rows=\"3\"></textarea></p>\n</crowd-form>\n]]>\n</  
HTMLContent>\n<FrameHeight>0</FrameHeight>\n</HTMLQuestion>",  
    "HITStatus": "Assignable",  
    "MaxAssignments": 5,  
    "Reward": "0.10",  
    "AutoApprovalDelayInSeconds": 259200,  
    "Expiration": "2020-09-29T18:30:03-07:00",  
    "AssignmentDurationInSeconds": 300,  
    "QualificationRequirements": [],  
    "HITReviewStatus": "NotReviewed",  
    "NumberOfAssignmentsPending": 0,  
    "NumberOfAssignmentsAvailable": 5,  
    "NumberOfAssignmentsCompleted": 0  
  }  
}
```

SDK for Python (Boto3)

The following Python code creates a new task using `create_hit`. This code can be run within a Jupyter Notebook or IPython as is, or can be incorporated into a Python script and executed.

```
import boto3  
  
mturk = boto3.client('mturk')  
  
question = ""
```

```
<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/  
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">  
<HTMLContent><![CDATA[  
<!DOCTYPE html>  
<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>  
<crowd-form>  
<p>Describe the current weather where you live</p>  
<p><textarea name="weather" cols="80" rows="3"></textarea></p>  
</crowd-form>  
<]>  
</HTMLContent>  
<FrameHeight>0</FrameHeight>  
</HTMLQuestion>"""  
  
response = mturk.create_hit(  
    Title='Describe the weather',  
    Description='Describe the current weather where you live',  
    Reward='0.10',  
    MaxAssignments=5,  
    LifetimeInSeconds=14400,  
    AssignmentDurationInSeconds=300,  
    AutoApprovalDelayInSeconds=259200,  
    Question=question)  
  
hit_id = response['HIT']['HITId']  
print('Created HIT: {}'.format(hit_id))
```

This creates the HIT in the marketplace and displays the following.

```
Created HIT: 3QQUBC64ZDDMMJE3ZX577RS5PMNNXJ
```

After creating the HIT, capture the `HITId` that was created and proceed to Step 2: Check task status.

Step 2: Check task status

In this step, we check the status of a task we created in Mechanical Turk. We poll the API for the status of our HIT using the `HITId` from the previous step. You need to capture that identifier and insert it in the following appropriate location to retrieve your results.

The code block examples in this section have been spaced out for readability.

AWS CLI

The following AWS CLI command retrieves the current state of a HIT using `get-hit`.

```
$ aws mturk get-hit --hit-id 3QQUBC64ZDDMMJE3ZX577RS5PMNNXJ
```

Using `get-hit` immediately after creating the HIT returns the following result. The CLI returns the same information that was returned when you first created the HIT. This includes all of the attributes for the HIT and its current state. As you can see in the highlighted lines below, all five assignments are available to be requested by workers and the `HITStatus` is `Assignnable`.

```
{  
  "HIT": {  
    "HITId": "3TL87MO8CLOFYXKXNRLMZO1MOK4FL5",
```

Amazon Mechanical Turk Developer Guide

Step 2: Check task status

```
"HITTypeId": "3AKE04YHPN13791QQA6BU4GMS7CHZJ",
"HITGroupId": "367GCHJ5533R84AG2MXJ00FCCJ7M9Q",
"CreationTime": "2020-09-29T14:30:03-07:00",
"Title": "Describe the weather",
>Description": "Describe the current weather where you live",
"Question": "<HTMLQuestion xmlns=\"http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd\"
  \n<HTMLContent><![CDATA[\n<!DOCTYPE html>\n<script src=\"https://
assets.crowd.aws/crowd-html-elements.js\"></script>
  \n<crowd-form>\n<p>Describe the current weather where you live</p>
  \n<p><textarea name=\"weather\" cols=\"80\" rows=\"3\"></
textarea></p>\n</crowd-form>\n]]>
  \n</HTMLContent>\n<FrameHeight>0</FrameHeight>\n</HTMLQuestion>",
"HITStatus": "Assignable",
"MaxAssignments": 5,
"Reward": "0.10",
"AutoApprovalDelayInSeconds": 259200,
"Expiration": "2020-09-29T18:30:03-07:00",
"AssignmentDurationInSeconds": 300,
"QualificationRequirements": [],
"HITReviewStatus": "NotReviewed",
"NumberOfAssignmentsPending": 0,
"NumberOfAssignmentsAvailable": 5,
"NumberOfAssignmentsCompleted": 0
}
}
```

If you run the same `get-hit` command again after about five to ten minutes, the results will likely match the following example. As you can see in the highlighted sections, the `HITStatus` is now `Reviewable` and there are no longer any assignments available or pending with workers. However, the completed count is still zero because none of the work has been approved yet.

```
{
  "HIT": {
    "HITId": "3TL87M08CLOFYKXNRLMZ01MOK4FL5",
    "HITTypeId": "3AKE04YHPN13791QQA6BU4GMS7CHZJ",
    "HITGroupId": "367GCHJ5533R84AG2MXJ00FCCJ7M9Q",
    "CreationTime": "2020-09-29T14:30:03-07:00",
    "Title": "Describe the weather",
    "Description": "Describe the current weather where you live",
    "Question": "<HTMLQuestion xmlns=\"http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd\"
  \n<HTMLContent><![CDATA[\n<!DOCTYPE html>\n<script src=\"https://
assets.crowd.aws/crowd-html-elements.js\"></script>
  \n<crowd-form>
  \n<p>Describe the current weather where you live</p>
  \n<p><textarea name=\"weather\" cols=\"80\" rows=\"3\"></
textarea></p>
  \n</crowd-form>\n]]>
  \n</HTMLContent>\n<FrameHeight>0</FrameHeight>
  \n</HTMLQuestion>",
    "HITStatus": "Reviewable",
    "MaxAssignments": 5,
    "Reward": "0.10",
    "AutoApprovalDelayInSeconds": 259200,
    "Expiration": "2020-09-29T18:30:03-07:00",
    "AssignmentDurationInSeconds": 300,
    "QualificationRequirements": [],
    "HITReviewStatus": "NotReviewed",
    "NumberOfAssignmentsPending": 0,
    "NumberOfAssignmentsAvailable": 0,
    "NumberOfAssignmentsCompleted": 0
  }
}
```

```
}
```

Python SDK (Boto3)

The following Python code retrieves the state of a HIT using `get_hit`. This code can be run within a Jupyter Notebook or IPython as is, or can be incorporated into a Python script and executed.

```
import boto3

mturk = boto3.client('mturk')

response = mturk.get_hit(HITId='3QQUBC64ZDDMMJE3ZX577RS5PMNNXJ')
response['HIT']
```

Using `get_hit` immediately after creating the HIT returns the following result. The command returns the same information that was returned when you first created the HIT. That includes all of the attributes for the HIT and its current state. As you can see in the highlighted lines below, all five assignments are available to be requested by workers and the `HITStatus` is `Assignable`.

```
{
  'HITId': '3TL87MO8CLOFYXKXNRLMZO1MOK4FL5',
  'HITTypeId': '3AKE04YHPN13791QQA6BU4GMS7CHZJ',
  'HITGroupId': '367GCHJ5533R84AG2MXJ00FCCJ7M9Q',
  'CreationTime': datetime.datetime(2020, 9, 29, 14, 30, 3, tzinfo=tzlocal()),
  'Title': 'Describe the weather',
  'Description': 'Describe the current weather where you live',
  'Question': '<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">
    \n<HTMLContent><![CDATA[\n<!DOCTYPE html>\n<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
    \n<crowd-form>\n<p>Describe the current weather where you live</p>
    \n<p><textarea name="weather" cols="80" rows="3"></textarea></p>
    \n</crowd-form>\n]]>\n</HTMLContent>\n<FrameHeight>0</FrameHeight>
    \n</HTMLQuestion>',
  'HITStatus': 'Assignable',
  'MaxAssignments': 5,
  'Reward': '0.10',
  'AutoApprovalDelayInSeconds': 259200,
  'Expiration': datetime.datetime(2020, 9, 29, 18, 30, 3, tzinfo=tzlocal()),
  'AssignmentDurationInSeconds': 300,
  'QualificationRequirements': [],
  'HITReviewStatus': 'NotReviewed',
  'NumberOfAssignmentsPending': 0,
  'NumberOfAssignmentsAvailable': 5,
  'NumberOfAssignmentsCompleted': 0
}
```

If you run the same `get_hit` command again after about five minutes, the results will likely appear as shown below. As you can see in the highlighted sections, the `HITStatus` is now `Reviewable` and there are no longer any assignments available or pending with workers. However, the completed count is still zero because none of the work has been approved yet.

```
{
  'HITId': '3TL87MO8CLOFYXKXNRLMZO1MOK4FL5',
  'HITTypeId': '3AKE04YHPN13791QQA6BU4GMS7CHZJ',
  'HITGroupId': '367GCHJ5533R84AG2MXJ00FCCJ7M9Q',
  'CreationTime': datetime.datetime(2020, 9, 29, 14, 30, 3, tzinfo=tzlocal()),
  'Title': 'Describe the weather',
```

```
'Description': 'Describe the current weather where you live',
'Question': '<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">
  \n<HTMLContent><![CDATA[\n<!DOCTYPE html>\n<script src="https://
assets.crowd.aws/crowd-html-elements.js"></script>
  \n<crowd-form>\n<p>Describe the current weather where you live</p>
\n<p><textarea name="weather" cols="80" rows="3"></textarea></p>
  \n</crowd-form>\n]]>\n</HTMLContent>\n<FrameHeight>0</FrameHeight>
\n</HTMLQuestion>',
'HITStatus': 'Reviewable',
'MaxAssignments': 5,
'Reward': '0.10',
'AutoApprovalDelayInSeconds': 259200,
'Expiration': datetime.datetime(2020, 9, 29, 18, 30, 3, tzinfo=tzlocal()),
'AssignmentDurationInSeconds': 300,
'QualificationRequirements': [],
'HITReviewStatus': 'NotReviewed',
'NumberOfAssignmentsPending': 0,
'NumberOfAssignmentsAvailable': 0,
'NumberOfAssignmentsCompleted': 0
}
```

Now that we've confirmed that the HIT is in the `Reviewable` state, we can proceed to Step 3: Retrieve Results.

Step 3: Retrieve results

In this step, we retrieve the results of a task we created in Mechanical Turk. We use the `HITId` from that was generated in Step 1. You need to capture that identifier and insert it in the appropriate location below to retrieve your results. Before running the commands in this step, you should confirm the `HITStatus` is `Reviewable` as shown in Step 2, or you may get incomplete results.

AWS CLI

The following AWS CLI command retrieves all of the submitted assignments for your HIT using `list-assignments-for-hit`.

```
$ aws mturk list-assignments-for-hit --hit-id 3TL87MO8CLOFYXXKNRLMZ01MOK4FL5
```

Using `list-assignments-for-hit` returns an array of results similar to those shown below. Each assignment has an `AssignmentId` and includes information about the Worker who submitted it, when they first accepted it, and when they submitted their answer. The answer information is captured in a `QuestionFormAnswers` XML data structure that you can read in to extract the results. For example, the first assignment below was submitted by the worker with `WorkerId` `A1KYPXUBSBWJBY`, it took them 25 seconds to complete it, and their answer was "Its currently raining lightly". Scrolling through the other assignments, you can see how other workers answered this question.

```
{
  "Assignments": [
    {
      "AssignmentId": "3IOEN3P9S7I9CGLBJQ50ANAQJXB16B",
      "WorkerId": "AIDACKCEVSQ6C2EXAMPLE",
      "HITId": "3TL87MO8CLOFYXXKNRLMZ01MOK4FL5",
      "AssignmentStatus": "Submitted",
      "AutoApprovalTime": "2020-10-02T14:39:42-07:00",
```

Amazon Mechanical Turk Developer Guide
Step 3: Retrieve results

```
    "AcceptTime": "2020-09-29T14:39:17-07:00",
    "SubmitTime": "2020-09-29T14:39:42-07:00",
    "Answer": "<?xml version='1.0' encoding='ASCII'><QuestionFormAnswers
xmlns='http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd'><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{\\"weather\\":\\"Its currently raining lightly\\"}]</
FreeText></Answer></QuestionFormAnswers>"
  },
  {
    "AssignmentId": "32AT8R96GL8U8BA6SRINMUBFCJ5USP",
    "WorkerId": "DAAICKCEVSO6C2EXAMPLE",
    "HITId": "3TL87MO8CLOFYXKNRMLMZO1MOK4FL5",
    "AssignmentStatus": "Submitted",
    "AutoApprovalTime": "2020-10-02T14:44:47-07:00",
    "AcceptTime": "2020-09-29T14:42:45-07:00",
    "SubmitTime": "2020-09-29T14:44:47-07:00",
    "Answer": "<?xml version='1.0' encoding='ASCII'><QuestionFormAnswers
xmlns='http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd'><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{\\"weather\\":\\"It is currently chilly and cloudy
outside and looks like it may rain in a little bit.\\"}</FreeText></Answer></
QuestionFormAnswers>"
  },
  {
    "AssignmentId": "3FTOP5WARFNLTMF07QVP5MWL0BPJ0W",
    "WorkerId": "DDAEBCKCEVSO6C2EXAMPLE",
    "HITId": "3TL87MO8CLOFYXKNRMLMZO1MOK4FL5",
    "AssignmentStatus": "Submitted",
    "AutoApprovalTime": "2020-10-02T14:48:12-07:00",
    "AcceptTime": "2020-09-29T14:43:28-07:00",
    "SubmitTime": "2020-09-29T14:48:12-07:00",
    "Answer": "<?xml version='1.0' encoding='ASCII'><QuestionFormAnswers
xmlns='http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd'><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{\\"weather\\":\\"Currently indianapolis weather very cool,
because northwest area are most of this situation landing. In currently approximately
cloud is 15 to 18'C.\\"}</FreeText></Answer></QuestionFormAnswers>"
  },
  {
    "AssignmentId": "3DI28L7YXADDPVEQP8OYMB230VZE19",
    "WorkerId": "AICACKCEVSO6C2EXAMPLE",
    "HITId": "3TL87MO8CLOFYXKNRMLMZO1MOK4FL5",
    "AssignmentStatus": "Submitted",
    "AutoApprovalTime": "2020-10-02T14:54:52-07:00",
    "AcceptTime": "2020-09-29T14:53:05-07:00",
    "SubmitTime": "2020-09-29T14:54:52-07:00",
    "Answer": "<?xml version='1.0' encoding='ASCII'><QuestionFormAnswers
xmlns='http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd'><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{\\"weather\\":\\"32\\"}</FreeText></Answer></
QuestionFormAnswers>"
  },
  {
    "AssignmentId": "3LO69W1SU3COZGELODW56TWTBHJLG3",
    "WorkerId": "SDLKSDEVSQ6C2EXAMPLE",
    "HITId": "3TL87MO8CLOFYXKNRMLMZO1MOK4FL5",
    "AssignmentStatus": "Submitted",
    "AutoApprovalTime": "2020-10-02T14:58:14-07:00",
    "AcceptTime": "2020-09-29T14:57:39-07:00",
    "SubmitTime": "2020-09-29T14:58:14-07:00",
    "Answer": "<?xml version='1.0' encoding='ASCII'><QuestionFormAnswers
xmlns='http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd'><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{\\"weather\\":\\"Hot and dry, with poor air quality.\\"}</
FreeText></Answer></QuestionFormAnswers>"
  }
}
```

```
]
}
```

Note that the AWS CLI supports various output formats as well as a option `--query` that can be used to filter your results. For example, the following only returns a list of the `WorkerIds` that completed the HIT.

```
$ aws mturk list-assignments-for-hit \
  --hit-id 3TL87MO8CLOFYXXKNRLMZ01MOK4FL5 \
  --query 'Assignments[*].WorkerId'
```

SDK for Python (Boto3)

The following Python code retrieves all of the submitted assignments for your HIT using `list_assignments_for_hit`. This code can be run within a Jupyter Notebook or IPython as is, or can be incorporated into a Python script and executed.

```
import boto3

mturk = boto3.client('mturk')

response = mturk.list_assignments_for_hit(HITId='3TL87MO8CLOFYXXKNRLMZ01MOK4FL5')
response['Assignments']
```

Using `list_assignments_for_hit` returns an array of results similar to those shown below. Each assignment has an `AssignmentId` and includes information about the worker who submitted it, when they first accepted it, and when they submitted their answer. The answer information is captured in a `QuestionFormAnswers` XML data structure that you can read in to extract the results. For example, the first assignment below was submitted by the Worker with `workerId` `AIDACKCEVSQ6C2EXAMPLE`, it took them 25 seconds to complete it, and their answer was "Its currently raining lightly". Scrolling through the other assignments, you can see how other workers answered this question.

```
[
  {
    'AssignmentId': '3IOEN3P9S7I9CGLBJQ50ANAQJXB16B',
    'WorkerId': 'AIDACKCEVSQ6C2EXAMPLE',
    'HITId': '3TL87MO8CLOFYXXKNRLMZ01MOK4FL5',
    'AssignmentStatus': 'Submitted',
    'AutoApprovalTime': datetime.datetime(2020, 10, 2, 14, 39, 42, tzinfo=tzlocal()),
    'AcceptTime': datetime.datetime(2020, 9, 29, 14, 39, 17, tzinfo=tzlocal()),
    'SubmitTime': datetime.datetime(2020, 9, 29, 14, 39, 42, tzinfo=tzlocal()),
    'Answer': '<?xml version="1.0" encoding="ASCII"?><QuestionFormAnswers
      xmlns="http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
      QuestionFormAnswers.xsd"><Answer><QuestionIdentifier>taskAnswers</
      QuestionIdentifier><FreeText>[{"weather": "Its currently raining lightly"}]</FreeText></
      Answer></QuestionFormAnswers>'
  },
  {
    'AssignmentId': '32AT8R96GL8U8BA6SRINMUBFCJ5USP',
    'WorkerId': 'DAAICKCEVSQ6C2EXAMPLE',
    'HITId': '3TL87MO8CLOFYXXKNRLMZ01MOK4FL5',
    'AssignmentStatus': 'Submitted',
    'AutoApprovalTime': datetime.datetime(2020, 10, 2, 14, 44, 47, tzinfo=tzlocal()),
    'AcceptTime': datetime.datetime(2020, 9, 29, 14, 42, 45, tzinfo=tzlocal()),
    'SubmitTime': datetime.datetime(2020, 9, 29, 14, 44, 47, tzinfo=tzlocal()),
```

```
'Answer': '<?xml version="1.0" encoding="ASCII"?><QuestionFormAnswers
xmlns="http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd"><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{"weather": "It is currently chilly and cloudy outside and
looks like it may rain in a little bit."}]</FreeText></Answer></QuestionFormAnswers>'
},
{'AssignmentId': '3FTOP5WARFNLTMF07QVP5MWL0BPJOW',
'WorkerId': 'DDAEBCKCEVSQ6C2EXAMPLE',
'HITId': '3TL87MO8CLOFYKXNRLMZ01MOK4FL5',
'AssignmentStatus': 'Submitted',
'AutoApprovalTime': datetime.datetime(2020, 10, 2, 14, 48, 12, tzinfo=tzlocal()),
'AcceptTime': datetime.datetime(2020, 9, 29, 14, 43, 28, tzinfo=tzlocal()),
'SubmitTime': datetime.datetime(2020, 9, 29, 14, 48, 12, tzinfo=tzlocal()),
'Answer': '<?xml version="1.0" encoding="ASCII"?><QuestionFormAnswers
xmlns="http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd"><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{"weather": "Currently indianapolis weather very cool,
because northwest area are most of this situation landing. In currently approximately
cloud is 15 to 18\C."}]</FreeText></Answer></QuestionFormAnswers>'
},
{'AssignmentId': '3DI28L7YXADDPVEQP8OYMB230VZE19',
'WorkerId': 'AICACKCEVSQ6C2EXAMPLE',
'HITId': '3TL87MO8CLOFYKXNRLMZ01MOK4FL5',
'AssignmentStatus': 'Submitted',
'AutoApprovalTime': datetime.datetime(2020, 10, 2, 14, 54, 52, tzinfo=tzlocal()),
'AcceptTime': datetime.datetime(2020, 9, 29, 14, 53, 5, tzinfo=tzlocal()),
'SubmitTime': datetime.datetime(2020, 9, 29, 14, 54, 52, tzinfo=tzlocal()),
'Answer': '<?xml version="1.0" encoding="ASCII"?><QuestionFormAnswers
xmlns="http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd"><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{"weather": "32"}]</FreeText></Answer></
QuestionFormAnswers>'
},
{'AssignmentId': '3L069W1SU3COZGELODW56TWTBHJLG3',
'WorkerId': 'SDLKSDEVSQ6C2EXAMPLE',
'HITId': '3TL87MO8CLOFYKXNRLMZ01MOK4FL5',
'AssignmentStatus': 'Submitted',
'AutoApprovalTime': datetime.datetime(2020, 10, 2, 14, 58, 14, tzinfo=tzlocal()),
'AcceptTime': datetime.datetime(2020, 9, 29, 14, 57, 39, tzinfo=tzlocal()),
'SubmitTime': datetime.datetime(2020, 9, 29, 14, 58, 14, tzinfo=tzlocal()),
'Answer': '<?xml version="1.0" encoding="ASCII"?><QuestionFormAnswers
xmlns="http://mechanicalturk.amazonaws.com/AWSMechanicalTurkDataSchemas/2005-10-01/
QuestionFormAnswers.xsd"><Answer><QuestionIdentifier>taskAnswers</
QuestionIdentifier><FreeText>[{"weather": "Hot and dry, with poor air quality."}]</
FreeText></Answer></QuestionFormAnswers>'
}
]
```

This information is returned in a Python dict that can be used to access any of the relevant values for each assignment. In addition, you can use the [ElementTree](#) library as shown below to parse the Answer XML into a format that can be more easily viewed. This code works well with tasks that use the `crowd-form` element as we did in this in this tutorial, but may need to be modified if you use a standard `form` element.

```
import xml.etree.ElementTree as ET
import json

answers = []
namespace = {'mt': 'http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2005-10-01/QuestionFormAnswers.xsd'}

for assignment in response['Assignments']:
```

```
assignment_answer = ET.fromstring(assignment['Answer']).find('mt:Answer',
namespace)
answers.append(json.loads(assignment_answer.find('mt:FreeText', namespace).text))

print(answers)
```

Now that we've reviewed the results of our HIT, we can proceed to Step 4: Approve Assignments.

Step 4: Approve Assignments

In this step, we approve the assignments submitted by workers so that the reward is transferred to their account. In Step 1, we set the Auto Approval Delay so that if we do nothing, workers are paid automatically after three days. However, it is always a best practice to approve work quickly if at all possible so that workers don't have to wait for the time to expire. Alternatively, if you plan to approve all of the assignments that are submitted, the Auto Approval Delay can be set to 0 and you can skip this step.

We use the same `HITId` that was generated in Step 1. You will need to capture that identifier and insert it in the appropriate location below to approve the results.

AWS CLI

To approve the assignments, we need to start by retrieving a list of the `AssignmentIds`. We begin by getting a list of the `AssignmentIds` using a variation of the query we used in Step 3 to retrieve results.

```
$ aws mturk list-assignments-for-hit \
--hit-id 3TL87M08CLOFYKXNRLMZ01MOK4FL5 \
--query 'Assignments[*].AssignmentId'
```

This returns a list of the `AssignmentIds` for our HIT as shown below.

```
[
  "3IOEN3P9S7I9CGLBJQ50ANAQJXB16B",
  "32AT8R96GL8U8BA6SRINMUBFCJ5USP",
  "3FTOP5WARFNLTMF07QVP5MWL0BPJ0W",
  "3DI28L7YXADDPVEQP8OYMB230VZE19",
  "3LO69W1SU3COZGELODW56TWTBHJLG3"
]
```

Now we can use the `approve-assignment` operation to approve each of the assignments.

```
$ aws mturk approve-assignment --assignment-id 3IOEN3P9S7I9CGLBJQ50ANAQJXB16B
$ aws mturk approve-assignment --assignment-id 32AT8R96GL8U8BA6SRINMUBFCJ5USP
$ aws mturk approve-assignment --assignment-id 3FTOP5WARFNLTMF07QVP5MWL0BPJ0W
$ aws mturk approve-assignment --assignment-id 3DI28L7YXADDPVEQP8OYMB230VZE19
$ aws mturk approve-assignment --assignment-id 3LO69W1SU3COZGELODW56TWTBHJLG3
```

SDK for Python (Boto3)

To approve the assignments, retrieve the list of assignments using the `HITId` and iterating through the results to call the `approve_assignment` operation.

```
import boto3

mturk = boto3.client('mturk')

response = mturk.list_assignments_for_hit(HITId='3TL87M08CLOFYXKXNRLMZ01MOK4FL5')

for assignment in response['Assignments']:
    mturk.approve_assignment(AssignmentId=assignment['AssignmentId'])
```

Creating and managing tasks (HITs)

This section contains information about how to create tasks (HITs) in Amazon Mechanical Turk. This includes the components of a task, how to create a task, and how to modify an existing task.

To learn how to retrieve HIT status and worker results, see [Retrieving HIT status \(p. 44\)](#) and [Retrieving results \(p. 41\)](#) respectively.

Topics

- [Define questions \(p. 29\)](#)
- [HIT attributes \(p. 35\)](#)
- [Creating HITs \(p. 37\)](#)
- [Modifying HITs \(p. 38\)](#)
- [Using the sandbox \(p. 39\)](#)
- [HIT references using RequesterAnnotation \(p. 40\)](#)

Define questions

The central component of a task in Amazon Mechanical Turk (Mechanical Turk) is the interface you provide for workers to give you the data you need. In Mechanical Turk tasks are presented via a web interface with which workers can interact and submit completed tasks. The following sections outline how to define a *question* in Mechanical Turk.

Topics

- [Use HTML to define questions \(p. 29\)](#)
- [Question definitions \(p. 33\)](#)
- [Requester website layouts \(p. 35\)](#)

Use HTML to define questions

Because workers access tasks (HITs) on Amazon Mechanical Turk (Mechanical Turk) using a web browser, the most common approach for creating a task interface for Mechanical Turk is to use HTML, CSS, and JavaScript. The following describes how the Mechanical Turk worker website works and the necessary steps to define an HTML task interface.

Topics

- [How Mechanical Turk tasks are rendered \(p. 29\)](#)
- [Defining an HTML form \(p. 30\)](#)
- [Crowd HTML Elements \(p. 31\)](#)
- [Templating \(p. 31\)](#)
- [Submitting from JavaScript \(p. 32\)](#)

How Mechanical Turk tasks are rendered

When a worker accepts a task in the Mechanical Turk marketplace, they are directed to a page for the assignment they've accepted. This page contains an `iframe` HTML element that has a URL for your task interface. If you define your question using the `HTMLQuestion` schema, the URL is for a Mechanical Turk

location containing the HTML you provided. If you use the `ExternalQuestion` schema, this is the URL you provide, which points to your own resources.

Mechanical Turk appends four query parameters to the URL that provide your task interface with information about the assignment, as shown in the following example.

```
https://tictactoe.amazon.com/gamesurvey.cgi?gameid=01523
&hitId=123RVWYBAZW00EXAMPLE
&assignmentId=123RVWYBAZW00EXAMPLE456RVWYBAZW00EXAMPLE
&turkSubmitTo=https://www.mturk.com/
&workerId=AZ3456EXAMPLE
```

In this example, the first line is the URL that was provided in an `ExternalQuestion` definition, and the additional lines contain query parameters that are appended by Mechanical Turk. These are described below.

- `hitId`: The ID of the HIT
- `assignmentId`: The ID of the assignment that the worker has accepted for this HIT
- `turkSubmitTo`: The Mechanical Turk server to which your form should submit a response
- `workerId`: The ID of the worker

In most cases, you won't need to directly interact with these values, but it can be useful information in various situations. For example, you can have your question check the `assignmentId` value to see if a worker is currently previewing your task or if they've accepted it. If the `assignmentId` is `ASSIGNMENT_ID_NOT_AVAILABLE`, you can disable input fields so that workers don't start working on it before they've accepted it.

Defining an HTML form

The example below describes a simple HTML form that prompts workers to answer a single question. The centerpiece of this task is the `form` element that directs the form to submit the worker response to <https://www.mturk.com/mturk/externalSubmit> with an HTTP POST call. This element contains a `textarea` input element with a `name` attribute that is associated with the worker's response. You can include as many input elements as you wish in your form, but each element must have a unique `name` attribute.

```
<!DOCTYPE html>
<script type='text/javascript' src='https://s3.amazonaws.com/mturk-public/
externalHIT_v1.js'></script>
<form method='post' id='mturk_form' action='https://www.mturk.com/mturk/externalSubmit'>
  <p>Describe the current weather where you live</p>
  <p><textarea name="weather" cols="80" rows="3"></textarea></p>
  <p><input type="submit" id="submitButton" class="btn btn-primary" value="Submit"/></p>
</form>
<script language='Javascript'>turkSetAssignmentID(</script>
```

To be able to successfully process the worker's response, Mechanical Turk needs the form response to include the `assignmentId` that was provided in the query parameters for this task. The two `script` elements in the example above use a Mechanical Turk library that will automatically populate the `assignmentId` value in the form.

Note that the form `action` attribute in the preceding example explicitly specifies the URL to which the form should be submitted. If you are testing in the sandbox environment, you need to modify this value to <https://workersandbox.mturk.com/mturk/externalSubmit>, or construct it using the value provided in the `turkSubmitTo` query parameter.

The HTML you include in your task can be as simple or complex as necessary to allow workers to provide the data you need. It's not uncommon for developers to include CSS and JavaScript code to provide a rich interface to workers, and in many cases, developers have leveraged React or other frontend libraries. Regardless of how you present the task to workers, you must use of a form element to submit a POST to Mechanical Turk when the task is complete. This directs the browser to advance and notify Mechanical Turk that the task is complete. For Mechanical Turk to accept the POST operation, it must include the `assignmentId` field so that Mechanical Turk can associate it with the worker's submission.

Crowd HTML Elements

[Crowd HTML Elements](#) are web components, a web standard that abstracts HTML markup, CSS, and JavaScript functionality into an HTML tag or set of tags. This allows you to build powerful task interfaces more easily.

To use Crowd HTML Elements you must include the following in your task HTML.

```
<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
```

This imports the library and gives you access to all of the elements. You can then replace the `form` element in your HTML with the `crowd-form` element as shown below. In this example, we have the same task interface as the one described in the previous section, but the HTML is simplified because the `crowd-form` element encapsulates the necessary code to direct where the response will be submitted, includes the `assignmentId`, and automatically adds a **Submit** button if it's not present.

```
<!DOCTYPE html>
<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
<crowd-form>
  <p>Describe the current weather where you live</p>
  <p><textarea name="weather" cols="80" rows="3"></textarea></p>
</crowd-form>
```

Crowd HTML Elements includes a range of elements ranging from wrappers for common form elements such as text fields and check boxes (`crowd-input` and `crowd-checkbox`), to full task interfaces for tasks such as drawing bounding boxes on images (`crowd-bounding-box`) and text entity annotation (`crowd-entity-annotation`).

For a full list of the available elements, see the [Crowd HTML Elements Reference](#).

Templating

When using Mechanical Turk tasks it's common to use a template approach to reuse the same question definition on multiple tasks. For example, the following can be used to gather keywords for a single image, but wouldn't be as useful when collecting keywords for thousands of images.

```
<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
<crowd-form>
  
  <p>Provide three keywords for the image above:</p>
  <p><input type="text" name="keyword1"></p>
  <p><input type="text" name="keyword2"></p>
  <p><input type="text" name="keyword3"></p>
</crowd-form>
```

To enable reuse of this HTML, you can replace the image URL with a template variable surrounded by curly braces as shown below. We've also added the URL as a *hidden* form element to make it easier to process the results by including the source URL alongside the keywords workers provide.

```
<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
<crowd-form>
  
  <input type="hidden" name="url" value="${url}">
  <p>Provide three keywords for the image above:</p>
  <p><input type="text" name="keyword1"></p>
  <p><input type="text" name="keyword2"></p>
  <p><input type="text" name="keyword3"></p>
</crowd-form>
```

When you iterate through the list of images, you can then simply perform a find/replace on the `${url}` value in the HTML so that each task has a unique question corresponding to an image URL.

Note that the `{ }` syntax used above is also used for tasks created using the Mechanical Turk requester website. There are a variety of different templating languages and libraries you can use to render your task interface. It's recommended you choose a library that works best for you in your desired programming language.

Submitting from JavaScript

In some cases, it's necessary to collect data using mechanisms other than form fields. For example, a task interface that prompts workers to draw a bounding box on an image would capture the coordinates of the box workers drew as a JavaScript variable. To allow workers to submit the data, you must place it in a form element that can be submitted to Mechanical Turk.

There are two common ways to do this. The first is using a hidden form value within your task as shown in the following code excerpt. The form includes hidden values for the `assignmentId` and `coordinates` that we want to return. When workers choose the **Submit** button, the `handleFormSubmit` function is called. The function populates the values in the hidden form elements and then submits the form.

```
<html>
...
<button onclick="handleFormSubmit()">Submit</button>
<form method='post' id='mturk_form' action='https://www.mturk.com/mturk/externalSubmit'>
  <input type="hidden" id="inputCoordinates" name="coordinates">
  <input type="hidden" id="inputAssignmentId" name="assignmentId">
</form>

<script>
function handleFormSubmit() {
  const urlParams = new URLSearchParams(window.location.search)
  document.getElementById('inputAssignmentId').value = urlParams.get('assignmentId')
  document.getElementById('inputCoordinates').value = JSON.stringify(coordinates)
  document.getElementById('mturk_form').submit()
}
</script>
```

You must include a value for `assignmentId` in your form submission so that Mechanical Turk can correctly associate the response with the correct worker and HIT.

In cases where you don't want to include the form and hidden inputs in the HTML, you can instead create and populate the elements dynamically from your JavaScript code as shown in the following example.

Note that in the following code, we're assigning the value for action based on the `turkSubmitTo` value from the URL search parameters for the task. This sets the correct value based on whether or not you are working in the production or sandbox environment.

```
const handleClick = () => {
  const urlParams = new URLSearchParams(window.location.search)

  // create the form element and point it to the correct endpoint
  const form = document.createElement('form')
  form.action = (new URL('mturk/externalSubmit', urlParams.get('turkSubmitTo'))).href
  form.method = 'post'

  // attach the assignmentId
  const inputAssignmentId = document.createElement('input')
  inputAssignmentId.name = 'assignmentId'
  inputAssignmentId.value = urlParams.get('assignmentId')
  inputAssignmentId.hidden = true
  form.appendChild(inputAssignmentId)

  // attach data I want to send back
  const inputCoordinates = document.createElement('input')
  inputCoordinates.name = 'coordinates'
  inputCoordinates.value = JSON.stringify(coordinates)
  inputCoordinates.hidden = true
  form.appendChild(inputCoordinates)

  // attach the form to the HTML document and trigger submission
  document.body.appendChild(form)
  form.submit()
}
```

Question definitions

Mechanical Turk provides three XML schemas that you can use to define your questions:

- [HTMLQuestion](#)
- [ExternalQuestion](#)
- [QuestionForm](#)

Use the following topics to learn more about these schemas.

Topics

- [HTMLQuestion \(p. 33\)](#)
- [ExternalQuestion \(p. 34\)](#)
- [QuestionForm \(p. 35\)](#)

HTMLQuestion

Most developers use the [HTMLQuestion](#) schema to create HITs. `HTMLQuestion` wraps an HTML form that is displayed to workers. This typically takes the following form:

```
<HTMLQuestion xmlns="http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2011-11-11/HTMLQuestion.xsd">
  <HTMLContent><![CDATA[
    <!DOCTYPE html>
```

```
<script src="https://assets.crowd.aws/crowd-html-elements.js"></script>
<crowd-form>
  <p>Describe the current weather where you live</p>
  <p><textarea name="weather" cols="80" rows="3"></textarea></p>
</crowd-form>
]]></HTMLContent>
<FrameHeight>0</FrameHeight>
</HTMLQuestion>
```

As you can see in the preceding XML, the HTML content is enclosed in a CDATA element within the XML and includes the elements that define how your question appears to workers. More information on defining your HTML can be found in [Use HTML to define questions \(p. 29\)](#).

ExternalQuestion

When you create a task using the `HTMLQuestion` format, your HTML is hosted by Mechanical Turk. Using Mechanical Turk to host your task helps ensure that it is in a highly available location. If, however, you want to host the task interface on your own servers or cloud resources, you can use the `ExternalQuestion` format.

To use `ExternalQuestion`, your URL must meet the following criteria:

- The location specified by the URL must support HTTPS.
- The location must be highly available and able to respond quickly when workers accept or preview your task.
- The URL cannot include any of the reserved query parameters (`assignmentId`, `hitId`, `turkSubmitTo`, and `workerId`) that is appended by Mechanical Turk, as described in [How Mechanical Turk tasks are rendered \(p. 29\)](#).
- The location specified by the URL must perform a form `POST` operation with the `assignmentId` when the task is completed, as described in [Use HTML to define questions \(p. 29\)](#).

Amazon S3 is a great option for hosting your task HTML. When you use an S3-hosted layout, you need to add a `ContentType` header and set the ACL to public read, as shown in the following Python SDK `PutObject` call:

```
s3_client.put_object(
    ACL='public-read',
    Body=HTML layout,
    Bucket=S3 bucket name,
    Key=Object name,
    ContentType="text/html"
)
```

When using this schema, you simply need to indicate the URL you want to use.

```
<?xml version="1.0" encoding="UTF-8"?>
<ExternalQuestion xmlns="http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2006-07-14/ExternalQuestion.xsd">
  <ExternalURL>https://tictactoe.amazon.com/gamesurvey.cgi?gameid=01523</ExternalURL>
  <FrameHeight>0</FrameHeight>
</ExternalQuestion>
```

When you use `ExternalQuestion`, you can make the served HTML as complex or simple as necessary for your particular task. However, you must use of a form element to submit a `POST` to Mechanical Turk

when the task is complete. This directs the browser to advance and notify Mechanical Turk that the task is complete.

QuestionForm

[QuestionForm](#) is a legacy XML format that can be used to define Mechanical Turk tasks using an XML schema. While it is still supported, we recommend building your task using [HTMLQuestion](#) or [ExternalQuestion](#).

Requester website layouts

If you have an existing task you created using the Mechanical Turk requester website, you can create HITs by referencing the `LayoutId` for that task and avoiding the need to provide a question value. To find the `LayoutId`, navigate to your [Project List](#) on the requester website and select the name of the project you want to use. A box pops up showing the **HIT Type ID**, **Layout ID**, and **Layout Parameters** that you can reference.

As shown in the following code example, the definition also includes the template parameters that are defined within the HTML. When you create a HIT, you need to pass both the `HITLayoutId` and values for the parameters that are defined.

```
{
  ...,
  HITLayoutId: "3O2UWD6SNTXSG9Z4I77HMZAX0U6499",
  HITLayoutParameters: [
    {
      "Name": "image_url",
      "Value": " https://my-bucket.s3.amazonaws.com/img1234.jpg "
    }
  ]
}
```

More detail on creating HITs with these values can be found in [Creating HITs \(p. 37\)](#).

HIT attributes

When you create a task (HIT) with Amazon Mechanical Turk (Mechanical Turk), you provide a number of attributes about the task that tell Mechanical Turk how to display it in the marketplace. These are separate from the content and question of the task itself, and include the title, description, reward amount, and attributes describing how long the task will remain active.

The attributes can be grouped as follows.

Descriptive attributes

- Title
- Description
- Keywords

Reward time-related attributes

- `AssignmentDurationInSeconds`
- `AutoApprovalDelayInSeconds`

- `LifetimeInSeconds`

MaxAssignments QualificationRequirements

The first six attributes (`Title`, `Description`, `Keywords`, `Reward`, `AssignmentDurationInSeconds`, and `AutoApprovalDelayInSeconds`) and `QualificationRequirements` comprise the *HIT Type* for the task. The HIT Type attributes are common across a group of HITs and allow the Mechanical Turk marketplace to group tasks together that have the same HIT Type. When workers complete a task, they can immediately start on the next task with the same HIT Type.

When you call the `CreateHIT` API and provide these values, Mechanical Turk first attempts to find an existing HIT Type that has the same values for these attributes. If one doesn't exist, then a new HIT Type is created. Alternatively, you can call `CreateHITType` with these attributes to directly create a HIT Type and then use the resulting HIT Type ID to create HITs using the `CreateHITWithHITType` API.

Title, Description, and Keywords

These should provide a clear description of the task to be completed, as well as any relevant search keywords that help workers find your task.

Reward

The amount that is transferred to the worker when the task is approved. This is in US dollars and should be provided to the API as a string. For example, to set the reward for a HIT at 10 cents, the value would be "0.10".

AssignmentDurationInSeconds

`AssignmentDurationInSeconds` is the amount of time that workers have to complete the task. For example, if you create a HIT that has an `AssignmentDurationInSeconds` of 300 it expires after 5 minutes. If a worker does not complete it in that time, it is assigned to a new worker.

You should specify a duration that is long enough that a worker can still complete it even if they run into difficulties. However, don't set it to be overly long, as this can result in an assignment becoming orphaned for the duration of the time. Setting a long assignment duration can also discourage workers from accepting it if workers believe the amount of time required is significantly out of line with the reward amount.

Auto-approval delay

After a worker submits an `Assignment`, you have the option to *approve* or *reject* their work. However, if the amount of time specified by the `AutoApprovalDelayInSeconds` elapses, then it is automatically approved. Timely approval or rejection of worker submissions is important to workers as it directly influences how quickly they get paid for the tasks they complete. As such, you should strive to review work shortly after it is submitted and set the auto-approval delay as low as possible so that workers don't have to wait inordinately long if you fail to approve some assignments. To automatically approve all work that is submitted, you can set the auto-approval delay to 0.

Lifetime

`LifetimeInSeconds` is the maximum amount of time that a HIT is available to be accepted in the Mechanical Turk marketplace. If the lifetime expires before all of the available assignments have been accepted, it is removed from the marketplace. Note that if the lifetime expires after a worker accepts an assignment, the worker can still submit a response up until the assignment duration expires.

Max Assignments

`MaxAssignments` specifies the maximum number of workers that can submit responses for a HIT.

Qualification Requirements

`QualificationRequirements` can be used to manage which workers can view and accept a HIT. More information on using qualification requirements can be found in [Managing workers \(p. 45\)](#).

Creating HITs

A human intelligence task, or HIT, is a question your application asks and a worker answers. There are two primary ways to create tasks (HITs) in Amazon Mechanical Turk (Mechanical Turk): directly or using a HIT Type ID. In both cases, you need to provide a valid question and attributes for your HIT.

Topics

- [Create a HIT Directly \(p. 37\)](#)
- [Create a HIT using a HIT Type \(p. 37\)](#)

Create a HIT Directly

The most common way to create HITs in Mechanical Turk is via the `CreateHIT` operation. This API can be called with a JSON object containing the following values.

```
{
  "Title": String,
  "Description": String,
  "Question": String,
  "HITLayoutId": String,
  "HITLayoutParameters": HITLayoutParameterList,
  "Reward": String,
  "AssignmentDurationInSeconds": Integer,
  "LifetimeInSeconds": Integer,
  "Keywords": String,
  "MaxAssignments": Integer,
  "AutoApprovalDelayInSeconds": Integer,
  "QualificationRequirements": QualificationRequirementList,
  "AssignmentReviewPolicy": ReviewPolicy,
  "HITReviewPolicy": ReviewPolicy,
  "RequesterAnnotation": String,
  "UniqueRequestToken": String
}
```

The response includes the `HITId` that was generated for the task as well as the `HITTypeId` that has been assigned to the HIT. It also includes all of the attributes of the HIT that were provided in the `CreateHIT` call.

Create a HIT using a HIT Type

Creating a HIT with a HIT Type allows you to be explicit about which HITs ought to be the same type and is a best practice for customers creating large numbers of HITs. It is also valuable when connecting notifications to your HITs as described in [Use Mechanical Turk notifications \(p. 56\)](#).

When creating a HIT using a HIT Type, you can use a HIT Type ID generated via a previous `CreateHIT` call or generate one by calling the `CreateHITType` operation. This API can be called with a JSON object containing the following values.

```
{
  "Title": String,
```

```
"Description": String,  
"Reward": String,  
"AssignmentDurationInSeconds": Integer,  
"Keywords": String,  
"AutoApprovalDelayInSeconds": Integer,  
"QualificationRequirements": QualificationRequirementList  
}
```

The response includes a generated `HITTypeId`, or the ID of an existing HIT Type in your account that has the same attributes. The `HITTypeId` can then be used to call `CreateHITWithHITType` with a JSON object containing the following values.

```
{  
  "HITTypeId": String,  
  "Question": String,  
  "HITLayoutId": String,  
  "HITLayoutParameters": HITLayoutParameterList,  
  "LifetimeInSeconds": Integer,  
  "MaxAssignments": Integer,  
  "AssignmentReviewPolicy": ReviewPolicy,  
  "HITReviewPolicy": ReviewPolicy,  
  "RequesterAnnotation": String,  
  "UniqueRequestToken": String  
}
```

The response includes the `HITId` that was generated for the HIT and all of the attributes of the HIT.

Modifying HITs

Once a HIT has been created in the Amazon Mechanical Turk (Mechanical Turk) marketplace, it's possible to modify it; however, there is an important caveat. Because your HIT is now available to workers, Mechanical Turk will not allow you to make changes that will negatively impact them. For example, Mechanical Turk won't let you delete a HIT on which a worker is actively working. The following are the common operations you can perform to modify HITs.

Topics

- [Modify expiration time \(p. 38\)](#)
- [Add additional assignments \(p. 39\)](#)
- [Modify the HIT Type \(p. 39\)](#)
- [Delete \(p. 39\)](#)

Modify expiration time

When a HIT is created, the `LifetimeInSeconds` is used to calculate an `ExpiresAt` value that tells Mechanical Turk when to remove a HIT from the marketplace if it hasn't been completed yet. You can use the `UpdateExpirationForHIT` operation to either extend this time further into the future to allow workers more time to complete it or shorten the time if responses are no longer valuable.

A common use of `UpdateExpirationForHIT` is to call it with a value of 0 or a time in the past to tell Mechanical Turk to immediately expire a HIT. This is useful when you make a mistake in your HIT definition and immediately need to remove the task from the marketplace. Note that this won't prevent workers who have already accepted your HIT from completing and submitting it.

Add additional assignments

When a HIT is created, the `MaxAssignments` value is provided and informs Mechanical Turk of how many workers can submit responses for the task. If you need to allow additional workers to provide responses, you can call [CreateAdditionalAssignmentsForHIT](#) to add additional available assignments.

A common pattern for managing quality in Mechanical Turk is to ask multiple workers to provide responses for a given task and then compare the results. One approach for doing this is to start with a minimal number of assignments and then add additional assignments if there is disagreement among workers. You can use this approach to limit the number of assignments needed for tasks where there is general agreement, while gathering additional data points for tasks that are more ambiguous.

Note that HITs created with fewer than 10 assignments cannot be extended to have 10 or more assignments.

Modify the HIT Type

The HIT Type attributes of a HIT, such as the reward amount, title, or assignment duration, can be modified by changing the HIT Type assigned using the [UpdateHITTypeOfHIT](#) operation. While this can allow you to modify any of the HIT Type attribute values on an active HIT, Mechanical Turk prevents you from changing the HIT Type on HITs with assignments that have been accepted or submitted by workers.

Delete

The `DeleteHIT` operation disposes of HITs that are no longer needed. Calling this operation deletes your HIT from Mechanical Turk. Once deleted, it is no longer available if you call [ListHITs](#), [GetHIT](#), or [ListAssignmentsForHIT](#). HITs must be in a `Reviewable` state to be disposed and all of the completed assignments must be either approved or rejected. As a result, this operation isn't a solution for removing HITs that you published by mistake, you should instead use the [UpdateExpirationForHIT](#) operation to immediately expire your task. HITs are automatically disposed after 120 days.

Using the sandbox

Many requesters choose to test their Amazon Mechanical Turk (Mechanical Turk) tasks in the sandbox environment. The sandbox is a mirror image of the *production* marketplace and is a useful way to test task interfaces and processes without spending money on worker rewards or fees. In the sandbox environment, you can perform all of the same operations you can perform in the production environment, such as creating HITs and retrieving results. This can be a great way to test your task interface and confirm that the results you receive meet your needs.

Because the sandbox is a mirror of production environment, you need to complete account setup a second time, as described in [Set up Amazon Mechanical Turk \(p. 11\)](#). To enable easy switching between testing and production, you can link to the same AWS account and use the same AWS credentials. To switch to the sandbox environment, simply specify that you want the CLI or SDK to use the sandbox endpoint (<https://mturk-requester-sandbox.us-east-1.amazonaws.com>).

API calls using the CLI

When using the sandbox from the AWS CLI, you need to specify the sandbox endpoint with each operation as shown in the following `GetAccountBalance` request.

```
aws mturk get-account-balance --endpoint https://mturk-requester-sandbox.us-east-1.amazonaws.com
```

This returns an available balance of *10000.00*, the default balance in the sandbox.

API calls using the Python SDK (boto3)

The Python AWS SDK, like all of the AWS SDKs, allows you to specify the endpoint when you instantiate a Mechanical Turk client. The following example shows how you would create a client and make a request using the `GetAccountBalance` operation.

```
import boto3

client = boto3.client(
    'mturk',
    endpoint_url='https://mturk-requester-sandbox.us-east-1.amazonaws.com'
)
print(client.get_account_balance()['AvailableBalance'])
```

This returns an available balance of *10000.00*, the default balance in the sandbox.

Testing HITs

To test your HITs in the sandbox, you can use the same operations described in [Creating HITs \(p. 37\)](#), provided you've configured the SDK or CLI to use the sandbox endpoint. When you create HITs in the sandbox, they are published to <https://workersandbox.mturk.com> instead of the production marketplace. Since this is a separate location and workers won't receive a reward for completing your tasks there, you need to create an account and complete the tasks yourself, or have members of your team assist, to fully test your HITs.

When HITs are completed in the sandbox environment, you can retrieve results using the same operations as those described in the following sections of this guide.

HIT references using `RequesterAnnotation`

When building processes that leverage Amazon Mechanical Turk (Mechanical Turk), it's often valuable to keep track of identifiers associated with the data in each HIT, particularly when handling HIT responses via notifications. For example, you might want to associate your HITs with a record in a database such as Amazon DynamoDB, and want your HIT to reference the primary key of the record.

The `RequesterAnnotation` attribute is a useful option for tracking these references. When you create a HIT using `CreateHIT` or `CreateHITWithHITType`, you can provide a `RequesterAnnotation` field that contains arbitrary data about each HIT. Although it is limited to 255 ASCII characters, this is generally adequate to capture identifiers that denote the origin of your data. The data provided here is only visible to the requester who created the HIT.

When you receive a notification that a HIT has been completed, you can use the `GetHIT` operation to retrieve the `RequesterAnnotation`. The identifier captured in the `RequesterAnnotation` can then be used to make updates in your database or other systems.

Retrieving results

Retrieving the results of a HIT involves gathering the responses from each of the assignment submissions provided by workers. At any point in a HIT's lifecycle (before it is disposed) you can call the [ListAssignmentsForHIT](#) operation to retrieve all of assignments that have been submitted and the answers provided by workers. If the some of the assignments are still awaiting submission by workers, you receive partial results containing the assignments have been submitted so far.

Topics

- [Assignment attributes \(p. 41\)](#)
- [Assignment answer \(p. 41\)](#)

Assignment attributes

The `ListAssignmentsForHIT` operation returns an array of zero or more assignments, each captured in an `Assignment` data structure. The following example shows the data structure that is returned. Each assignment includes the ID of the HIT with which it's associated, the worker who submitted it, and the ID of the assignment itself. The `AssignmentStatus` is one of `Submitted`, `Approved`, or `Rejected`. When a worker first submits an assignment, the status is `Submitted`, and changes based on your decision to approve or reject the assignment.

```
{
  AssignmentId: "123RVWYBAZW00EXAMPLE456RVWYBAZW00EXAMPLE",
  WorkerId: "AZ3456EXAMPLE",
  HITId: "123RVWYBAZW00EXAMPLE",
  AssignmentStatus: "Submitted",
  AcceptTime: "2019-12-01T12:00:00Z",
  SubmitTime: "2019-12-01T13:04:59Z",
  AutoApprovalTime: "2019-12-04T13:04:59Z",
  Answer: '...'
}
```

The `AcceptTime` and `SubmitTime` indicate when a worker first accepted the task and when they submitted it. This can be used to infer a rough approximation of how long it took the worker to complete the task; however, keep in mind that a worker may step away from their computer for a period of time after accepting a task so the time delta may be inflated.

The `AutoApprovalTime` indicates when the assignment will be automatically approved if you don't approve or reject it yourself. This is computed by adding the auto-approval delay you set when creating the HIT to the `SubmitTime`.

Assignment answer

The `Answer` value is returned as a string containing a `QuestionFormAnswers` XML data structure. The layout of this data structure corresponds to the HTML form fields you provided in your question HTML. For example, consider the following task form using a standard HTML form element and form field.

```
<form method='post' id='mturk_form' action='https://www.mturk.com/mturk/externalSubmit'>
  <p>What country do you live in?</p>
  <p><input type="text" name="country"></p>
  <p>Describe the current weather where you live</p>
  <p><textarea name="weather" cols="80" rows="3"></textarea></p>
  <p><input type="submit" id="submitButton" class="btn btn-primary" value="Submit"/></p>
</form>
```

The `QuestionFormAnswers` data structure contains an `Answer` value for each form field in your HTML. The `QuestionIdentifier` is the name supplied for the form field in your HTML and the `FreeText` attribute is the value that was entered in the field by the worker. Here the `QuestionIdentifier` values `country` and `weather` match the names specified in the form definition. The `FreeText` values contain the worker's responses.

```
<?xml version="1.0" encoding="ASCII"?>
<QuestionFormAnswers xmlns="http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2005-10-01/QuestionFormAnswers.xsd">
  <Answer>
    <QuestionIdentifier>country</QuestionIdentifier>
    <FreeText>United States</FreeText>
  </Answer>
  <Answer>
    <QuestionIdentifier>weather</QuestionIdentifier>
    <FreeText>It's currently raining lightly</FreeText>
  </Answer>
</QuestionFormAnswers>
```

Assignment Answer when using crowd-form

The structure of the answers returned in the `QuestionFormAnswer` data structure is slightly different when using the `crowd-form` element from the `Crowd HTML Elements` library. Consider the following example that mirrors the preceding one, but replaces `form` with `crowd-form`.

```
<crowd-form>
  <p>What country do you live in?</p>
  <p><input type="text" name="country"></p>
  <p>Describe the current weather where you live</p>
  <p><textarea name="weather" cols="80" rows="3"></textarea></p>
</ crowd-form>
```

The `crowd-form` element collapses all of the form field values into a single `Answer` value in the `QuestionFormAnswer` data structure as shown in the following example. This answer always has the identifier `taskAnswers` and the `FreeText` value is an array of JSON key-value pairs containing the worker answers.

```
<?xml version="1.0" encoding="ASCII"?>
<QuestionFormAnswers xmlns="http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2005-10-01/QuestionFormAnswers.xsd">
  <Answer>
    <QuestionIdentifier>taskAnswers</QuestionIdentifier>
    <FreeText>
      [
        {
          "country": "United States",
          "weather": "It's currently raining lightly"
```

```
    }  
  ]  
</FreeText>  
</Answer>  
</QuestionFormAnswers>
```

Retrieving HIT status

At any point after a HIT is created in Amazon Mechanical Turk (Mechanical Turk) up until it is disposed, you can retrieve the status of a HIT using the `GetHIT` operation. In addition, you can call `ListHITs`, `ListHITsForQualificationType`, or `ListReviewableHITs` to retrieve HIT information. The `HIT data structure` returned by these operations includes the attributes and question used to create the HIT, as well as attributes that describe the current status of the HIT.

The most useful attribute to monitor is the `HITstatus` value, which can be used to evaluate if a HIT is complete. When a HIT is created, its `HITstatus` initially has a value of `Assignable`, which indicates that it's possible for a worker to accept it and begin working on it. After the maximum number of assignments have been accepted, the HIT moves to a status of `Unassignable` because no additional workers can accept it. Finally, a HIT moves to the `Reviewable` state when all of the assignments have been submitted or the HIT has expired. Keep in mind that you can retrieve interim results at any time, regardless of whether or not all of the HITs have been submitted.

You can use the values `NumberOfAssignmentsAvailable` and `NumberOfAssignmentsPending` to monitor assignment acceptance and submission. When a HIT is first created, the number of assignments available is equal to the `MaxAssignments` value. As HITs are accepted by workers, this number is reduced and the number of assignments pending increases. When workers submit assignments, they are no longer reflected in `NumberOfAssignmentsPending`. Similarly, if workers *return* an assignment or the assignment duration elapses, the assignment is no longer reflected in pending but returns to available. When all of the available assignments have been submitted or the HIT expires, the number of pending and available are both 0.

HIT Lifecycle

The following example shows the HIT status values through the lifecycle of a HIT with `MaxAssignments` set to 5.

Step	HITstatus	Number Of assignments available	Number Of assignments pending
Initial state	Assignable	5	0
Three assignments have been accepted	Assignable	2	3
All assignments accepted and two submitted	Unassignable	0	3
One assignment is returned	Assignable	1	2
All assignments submitted	Reviewable	0	0
Three of the submitted assignments have been approved	Reviewable	0	0
All of the assignments have been approved	Reviewable	0	0

Managing workers

This section contains information about how you can improve the quality of the results you get from Amazon Mechanical Turk (Mechanical Turk) by rewarding your workers for successful completion of your tasks and specifying which workers you want working on them.

Topics

- [Approving and rejecting work \(p. 45\)](#)
- [Awarding a bonus \(p. 45\)](#)
- [Blocking workers \(p. 46\)](#)
- [Selecting eligible workers \(p. 46\)](#)
- [Working with custom qualification types \(p. 51\)](#)
- [Communicate with workers \(p. 55\)](#)

Approving and rejecting work

Before workers can receive payment for tasks you post to Amazon Mechanical Turk (Mechanical Turk), the work must be approved. You can review a worker's submission prior to transferring funds to their account.

When you create tasks (HITs) in Mechanical Turk, one of the attributes you can set is the auto-approval delay. This specifies how long you have to review a task before it's automatically approved. As a general rule, it's best to keep this delay as short as possible so that workers don't have to wait overly long to receive payment for the work they did for you. Being prompt and fair in payment contributes to a positive relationship with the worker community.

Between the time a task is submitted and when the auto-approval delay is reached, you can use `ListAssignmentsForHIT` to review a worker's submission and validate that they've made a reasonable effort to complete the task successfully. If so, you can call `ApproveAssignment` to authorize payment to the worker. If you identify a worker that is putting in no effort at all on your task (spamming), you can use the `RejectAssignment` operation to invalidate their submission.

We recommend you only reject work when workers are clearly putting in no effort to submit an accurate response to your task. It's inappropriate to penalize a worker for submitting data incorrectly because you provided unclear instructions or they simply made a mistake in interpreting what you wanted them to do. Most workers zealously guard their approval rating and avoid doing work for requesters that they believe are in unfair in how they reject work.

In the event that you reject an assignment but then discover that the issue was not the worker's fault, you can call `ApproveAssignment` to reverse the rejection, but only for assignments submitted in the last 30 days that haven't been deleted.

Awarding a bonus

You can send bonus payments to workers who have completed an assignment for you in Amazon Mechanical Turk (Mechanical Turk) in the past six months. Requesters commonly use bonus payments to recognize workers that perform tasks particularly well, or go above and beyond in helping to resolve problems with a task interface.

To send a bonus, you can use the [SendBonus](#) operation. You need to provide the ID of the worker and a past assignment that they've completed for you. The operation also requires that you specify the bonus amount in US Dollars and provide a reason for the award.

Note that your account is charged for the bonus payment as well as Mechanical Turk fees.

Blocking workers

When you identify workers that are not putting in the requested effort on your tasks on Amazon Mechanical Turk (Mechanical Turk), you have the option to block them. This prevents them from doing any future work for you so long as the block is in place. Workers are notified that they've been blocked from your tasks.

The [CreateWorkerBlock](#) operation can be used to block a worker by simply providing the ID of the worker and a reason for the block. Similarly, you can remove a block by using the [DeleteWorkerBlock](#) operation. At any point you can retrieve all of the workers that have been blocked by using the [ListWorkerBlocks](#) operation.

Note that we recommend you be judicious in your use of worker blocks and only block those workers that are clearly not making an attempt to correctly respond to your task (spamming). If a worker is simply misreading instructions or lacks the requisite skills to complete your task successfully, we advise you to use a custom qualification requirement to exclude them from future tasks, rather than a block. Because the blocks a worker receives are a component of Mechanical Turk worker review policies and frequent blocks may result in account suspension, workers are sensitive to being blocked by requesters. If the worker community believes that you are blocking workers unfairly, they may choose to avoid accepting your tasks in the future.

Selecting eligible workers

By default, all tasks (HITs) posted to Amazon Mechanical Turk (Mechanical Turk) are available to all active workers in the Mechanical Turk marketplace. You can restrict the audience that's eligible for your HITs by adding qualification requirements. Qualification requirements can be used to both restrict the audience to workers that meet certain criteria, or exclude those that have certain attributes. These requirements operate on attributes that are assigned to workers as *qualifications*. Qualifications can be assigned by the Mechanical Turk system or requesters, and are visible to workers in their account.

Topics

- [Qualifications and qualification types \(p. 46\)](#)
- [Qualification requirements \(p. 47\)](#)
- [System qualification types \(p. 48\)](#)

Qualifications and qualification types

Qualification types are system- or requester-defined descriptions of an attribute that can be associated with a worker. One example is the system-generated qualification type `NumberHITSApproved`, which measures the number of HITs a worker has submitted and had approved. Another would be a requester-defined qualification type that tracks how accurate a worker has been on previous tasks that the requester has posted.

When a qualification type is assigned to a worker, it is applied as a *qualification* for that worker. In the case of the system-generated `NumberHITSApproved` qualification type, a qualification is automatically

created for a worker as the work they submit is approved. For custom qualification types, a requester can assign the qualification to a worker using the [AssociateQualificationWithWorker](#) operation and optionally providing an integer value to associate with it.

Qualification requirements

A requirement is defined when calling either `CreateHIT` or `CreateHITType`. Either operation accepts an array of one or more [QualificationRequirement](#) data structures to specify the qualifications workers must have to be eligible for your HIT.

The `QualificationRequirement` data structure comprises four attributes: `QualificationTypeId`, `Comparator`, `value` (either `IntegerValues` or `LocaleValues`), and `ActionsGuarded`. The `QualificationTypeId` specifies the qualification type that should be applied and can be either the ID of an Mechanical Turk system qualification type, or one you create in your account. The `Comparator` and `value` are then used to evaluate if the worker has the required qualification attributes to be eligible for the HIT. Finally, the `ActionsGuarded` indicates the level of visibility that a HIT has to workers that aren't eligible to accept it.

The `Comparator` attribute specifies how the qualification type is evaluated and is typically used with a `value`. The following table illustrates the values that can be used for `Comparator` and the required `value` attribute, if any. The existence `Comparators` don't require a `value` attribute since they are only used to evaluate if the qualification type has been assigned, not the value that has been associated with it.

Type	Values	Required value attribute	Example
Existence	Exists	None	Only include workers that have been assigned the qualification type, regardless of value.
	DoesNotExist	None	Exclude workers that have been assigned the qualification type.
Numeric	LessThan	IntegerValues	Only include workers that have been assigned the qualification type where the assigned value is less than 50.
	LessThanOrEqualTo	IntegerValues	Only include workers that have been assigned the qualification type where the assigned value is less than or equal to 50.
	GreaterThan	IntegerValues	Only include Workers that have been assigned the Qualification Type where the assigned value is greater than 50.
	GreaterThanOrEqualTo	IntegerValues	Only include workers that have been assigned the qualification type where the assigned

Type	Values	Required value attribute	Example
			value is greater than or equal to 50.
Equivalence	EqualTo	IntegerValues or LocaleValues	Only include workers located in Spain.
	NotEqualTo	IntegerValues or LocaleValues	Only include workers who have been assigned the qualification type where the assigned value is not 42.
Set	In	IntegerValues or LocaleValues	Only include workers who have been assigned the qualification type where the assigned value is 1, 2, 3 or 8.
	NotIn	IntegerValues or LocaleValues	Only include workers who are not located in the US States of Florida and Georgia.

The `ActionsGuarded` attribute indicates the level of visibility that your HIT has to workers who aren't eligible for it. This defaults to `Accept`, which indicates that ineligible workers can't accept it, but they can see it in the Mechanical Turk marketplace and preview the task if they wish so that they can request a qualification to work on the HIT if they wish to. If you want to prevent them from previewing it, you can set the `ActionsGuarded` to `PreviewAndAccept`; they can then see it in their list of available tasks and request any custom qualifications. Finally, `DiscoverPreviewAndAccept` hides the HIT from ineligible workers.

The `CreateHIT` and `CreateHITType` operations accept an array of qualification requirements which can include one or more qualification requirement data structures so you can apply multiple requirements to workers to be eligible for your task. Because workers must meet *all* of the requirements, be careful to ensure that the requirements do not conflict. For example, if you had a qualification requirement that specified workers must be located in the US and a second requirement that they be located in Canada, it would be impossible for any worker to meet both criteria. If your goal is to include workers in either the US or Canada, you would need to replace the two requirements with a single requirement using the `In` `Comparator` to restrict workers to those in an array containing both the US and Canada.

System qualification types

The most commonly used qualification types are those provided by Mechanical Turk. These include the following:

- A *HITs approved* qualification for the number of HITs that workers have successfully completed in the past, which can be used to identify workers with more or less experience.
- An *Approval Percentage* qualification to specify workers that you have approved at a specified rate on previous tasks.
- A *Locale* qualification to select workers in specific countries or US states.

- A *Masters* qualification that is awarded to workers that have demonstrated superior performance over a period of time across thousands of HITs.
- An *Adult* qualification that selects workers who have indicated they are over 18 years of age and are willing to work on potentially offensive content.

Each of these qualification types has an associated `QualificationTypeId` which can be found in the documentation for [QualificationRequirement](#).

Using the HITs Approved qualification type

The `NumberHITsApproved` qualification type restricts tasks to workers with more or less experience based on their past work on Mechanical Turk. For example, if you only wanted to use workers who were relatively new to the platform and had successfully submitted fewer than 500 HITs, you would use the following value for `QualificationRequirements`.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '00000000000000000040',  
    Comparator: 'LessThan',  
    IntegerValues: [500]  
  }  
]
```

If, instead, you wanted more experienced workers who had successfully completed 100 HITs, you would use the following.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '00000000000000000040',  
    Comparator: 'GreaterThanOrEqualTo',  
    IntegerValues: [100]  
  }  
]
```

Using the Masters qualification type

The `Masters` qualification type is a Mechanical Turk–managed qualification type that is assigned to workers when they have demonstrated superior performance over a period of time across thousands of HITs. There is no value associated with it, so you can simply use the `Exists` comparator to apply it to your tasks.

Note that there is an additional fee for using the `Masters` qualification on your task as described in [Mechanical Turk Pricing](#).

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '2F1QJWKUDD8XADTFD2Q0G6UTO95ALH',  
    Comparator: 'Exists'  
  }  
]
```

Using the Percentage Approved type

The `PercentAssignmentsApproved` qualification type restricts tasks based on how often you have approved or rejected past work a worker has done for you. For example, to only accept workers that

have an approval rate of greater than or equal to 95%, the following qualification requirement would be included in your [CreateHIT](#) calls.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '000000000000000000L0',  
    Comparator: 'GreaterThanOrEqualTo',  
    IntegerValues: [95]  
  }  
]
```

Note that a worker's approval rate is statistically meaningless for small numbers of assignments, since a single rejection can reduce the approval rate by many percentage points. To ensure that a new worker's approval rate is unaffected by these statistically meaningless changes, if a worker has submitted fewer than 100 assignments for you, the worker's approval rate is 100%.

Using the Locale qualification type

Locale is a Mechanical Turk qualification type that specifies the workers that are eligible for your task based on where they are located. To use Locale, you must specify one or more `LocaleValues` using a JSON data structure that includes a `Country` attribute and can optionally include a `Subdivision` attribute. The `Country` attribute should specify the two-character country code of the country. The `Subdivision` attribute is only supported when the `Country` is "US" and should specify the two-character state code for the US state. The following example would restrict workers to those in the US state of Minnesota.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '0000000000000000000071',  
    Comparator: 'EqualTo',  
    LocaleValues: [  
      {  
        Country: "US",  
        Subdivision: "MN"  
      }  
    ]  
  }  
]
```

To select multiple locations, you should use the `In` comparator and a list of locales as shown in the following example, which restricts the task to workers in the US and Canada.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '0000000000000000000071',  
    Comparator: 'In',  
    LocaleValues: [  
      { Country: "US" },  
      { Country: "CA" }  
    ]  
  }  
]
```

Use caution when using multiple Locale qualification requirements in the same HIT. If the requirements above were split into two requirements, one for the US and one for CA, no workers could accept the task.

However, using two Locale requirements is a good way to restrict workers to the US but exclude selected states. The following would include all workers in the US with the exception of workers in Florida and Georgia.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '00000000000000000071',  
    Comparator: 'Equals',  
    LocaleValues: [  
      { Country: "US" }  
    ]  
  },  
  {  
    QualificationTypeId: '00000000000000000071',  
    Comparator: 'NotIn',  
    LocaleValues: [  
      { Country: "US", Subdivision: 'FL'},  
      { Country: "US", Subdivision: 'GA'}  
    ]  
  }  
]
```

Objectionable content

Some tasks, such as image moderation, involve handling content that some workers might find objectionable, typically because it involves imagery that contains violence or nudity. If there is the potential that some of your HITs may contain objectionable content, you should make use of the Adult qualification type. This restricts the task to workers who have confirmed they are over 18 years of age and are willing to view potentially objectionable content. The qualification requirement should also specify an `ActionsGuarded` value of `PreviewAndAccept` or `DiscoverPreviewAndAccept`.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '00000000000000000060',  
    Comparator: 'Equals',  
    IntegerValues: [1],  
    ActionsGuarded: 'PreviewAndAccept'  
  }  
]
```

In addition, you should include "(WARNING: This HIT may contain adult content. Worker discretion is advised.)" in the title of your HIT.

Custom qualification type

Requester-defined qualification types can also be created to handle a range of needs in managing who can work on your tasks. Information on how to create and use custom qualification types can be found in [Working with custom qualification types \(p. 51\)](#).

Working with custom qualification types

When using Amazon Mechanical Turk (Mechanical Turk), you can create qualification types that you can then assign to workers as qualifications. Qualifications can be used for a range of worker management

approaches, such as identifying workers that have met certain criteria in past tasks (HITs) or assigning a score based on performance over time. The following discusses how to create and assign qualification types to workers, as well as how to modify or revoke them.

Mechanical Turk also provides the option to create qualification tests that allow workers to take a test to be assigned a qualification automatically. That topic isn't addressed here, but more information can be found in the [API Documentation](#).

Topics

- [Create a qualification type](#) (p. 52)
- [Assign or remove a worker qualification](#) (p. 52)
- [Qualification requests](#) (p. 52)
- [Tutorial: Creating a qualification requirement that requires workers be in a group](#) (p. 53)
- [Tutorial: Create a qualification requirement that workers have achieved at least 80% accuracy on previous tasks](#) (p. 53)
- [Tutorial: Creating a qualification type to exclude workers from selected tasks](#) (p. 54)

Create a qualification type

The `CreateQualificationType` operation can be used to register a new qualification type in your account. Simply specify the name, provide a brief description, and specify `Active` as the status. Note that the qualification type name and description are visible to workers. You can update these values using the `UpdateQualificationType` operation.

Assign or remove a worker qualification

To assign a qualification type to a worker, use the `AssociateQualificationWithWorker` operation, specifying the ID of the qualification type and the worker it should be applied to. You can also assign an integer value such as a score. To modify the integer value, call the `AssociateQualificationWithWorker` operation again with the new value.

You can remove a qualification using the `DisassociateQualificationFromWorker` operation.

Qualification requests

When workers don't have one of the custom qualification types required to do your task, they have the option to request it from the Mechanical Turk marketplace. This is most commonly associated with tasks that have qualification tests but all custom qualification types can be requested.

These requests can be queried using the `ListQualificationRequests` operation and can be approved or rejected using the `AcceptQualificationRequest` or `RejectQualificationRequest` operations respectively.

Additional operations

The following operations can be used when working with qualifications.

Additional Operations

- `ListQualificationTypes`: Retrieves a list of your existing qualification types.
- `GetQualificationType`: Retrieves the details of a qualification type.
- `ListWorkersWithQualificationType`: Retrieves a list of workers that have been assigned a qualification type.
- `ListHITsForQualificationType`: Retrieves a list of HITs that include a qualification type in their requirements.

- [GetQualificationScore](#): Retrieves the qualification assigned to a worker for a qualification type.

Tutorial: Creating a qualification requirement that requires workers be in a group

In the following example, we create a qualification type that describes a group of workers that have demonstrated expertise at a task and add it to our qualification requirements. To start, we use the [CreateQualificationType](#) operation to create the type with which we're working.

```
{
  Name: 'Experts',
  Description: 'Demonstrated expertise at my task',
  QualificationTypeStatus: 'Active'
}
```

The [CreateQualificationType](#) operation will return an ID, 3TL87MO8CLOFYXKXNRLM00EXAMPLE, that we can assign to workers. For each worker, we call the [AssociateQualificationWithWorker](#) operation to add them to our group.

```
{
  WorkerId: 'AZ3456EXAMPLE',
  QualificationTypeId: '3TL87MO8CLOFYXKXNRLM00EXAMPLE'
}
```

Now that we've built our group, we can reference it in the [QualificationRequirements](#) for our HITs as shown in the following example.

```
QualificationRequirements: [
  {
    QualificationTypeId: '3TL87MO8CLOFYXKXNRLM00EXAMPLE',
    Comparator: 'Exists',
    ActionsGuarded: 'DiscoverPreviewAndAccept'
  }
]
```

Because the `ActionsGuarded` is set to `DiscoverPreviewAndAccept`, it is only visible to workers who've been assigned the qualification type.

Tutorial: Create a qualification requirement that workers have achieved at least 80% accuracy on previous tasks

In the following example, we create a qualification type that we can use to record how well workers did on a previous set of tasks and then build a qualification requirement that requires them to have achieved at least 80% accuracy. In this approach, we start by posting a set of HITs to which we already know the answer. When workers respond to these HITs, we can track their responses against the known answers and assign them a score as a qualification.

To start, we use the [CreateQualificationType](#) operation to create the type with which we want to work.

```
{
  Name: 'Task Scores',
  Description: 'Score on previous tasks',
  QualificationTypeStatus: 'Active'
}
```

The [CreateQualificationType](#) operation returns an ID, 3TL87MO8CLOFYXKXNRLM00EXAMPLE, that we can assign to workers. For each worker, we call the [AssociateQualificationWithWorker](#) operation to record the score they achieved on the earlier tasks. In the example below, we record that the worker was 93% accurate on the test HITs.

```
{
  WorkerId: 'AZ3456EXAMPLE',
  QualificationTypeId: '3TL87MO8CLOFYXKXNRLM00EXAMPLE',
  IntegerValue: 93
}
```

Now that we've built our group, we can reference it in the [QualificationRequirements](#) for our HITs as shown below.

```
QualificationRequirements: [
  {
    QualificationTypeId: '3TL87MO8CLOFYXKXNRLM00EXAMPLE',
    Comparator: 'GreaterThanOrEqualTo',
    IntegerValues: [80]
  }
]
```

Tutorial: Creating a qualification type to exclude workers from selected tasks

In the following example, we create a qualification type that describes a group of workers that have demonstrated they don't perform well at our tasks and excludes them in our qualification requirements. To start, we use the [CreateQualificationType](#) operation to create the type with which we want to work.

```
{
  Name: 'Excluded',
  Description: 'Excluded from this task',
  QualificationTypeStatus: 'Active'
}
```

The [CreateQualificationType](#) operation returns an ID, 3TL87MO8CLOFYXKXNRLM00EXAMPLE, that we can assign to workers. For each worker, we call the [AssociateQualificationWithWorker](#) operation to add them to the excluded group.

```
{  
  WorkerId: 'AZ3456EXAMPLE',  
  QualificationTypeId: '3TL87MO8CLOFYXKXNRLM00EXAMPLE'  
}
```

Now that we've built our group, we can reference it in the `QualificationRequirements` for our HITs as shown in the following example.

```
QualificationRequirements: [  
  {  
    QualificationTypeId: '3TL87MO8CLOFYXKXNRLM00EXAMPLE',  
    Comparator: 'DoesNotExist',  
    ActionsGuarded: 'DiscoverPreviewAndAccept'  
  }  
]
```

Because the `ActionsGuarded` has been set to `DiscoverPreviewAndAccept`, it is not visible to workers who've been assigned the qualification type.

Communicate with workers

You can send messages to workers in Amazon Mechanical Turk (Mechanical Turk) if you've previously accepted or rejected an assignment from that worker using the [NotifyWorkers](#) operation. It's common to use this operation when you want to notify workers that you've posted new tasks for them to work on, or alert them to changes in your task interface.

Similarly, workers can send messages to you via email if they have questions about your task, rejections, or other comments. You are welcome to engage with workers if you wish. We encourage you to maintain a positive relationship with the worker community.

Worker forums

There are a number of forums where workers congregate to discuss tasks, requesters, and the Mechanical Turk platform in general. If you wish to engage the worker community to get their input on proposed tasks or other topics, you can post to one of the forums listed below.

- [Mechanical Turk Subreddit](#)
- [Mechanical Turk Crowd](#)
- [Mechanical Turk Forum](#)
- [TurkerView](#)

Use Mechanical Turk notifications

Amazon Mechanical Turk (Mechanical Turk) has a notifications capability that can be used to trigger actions when your HITs reach various stages. Using notifications, you can process the results of assignments and HITs immediately after they are submitted to evaluate worker submissions or begin downstream processing of data. This allows you to integrate Mechanical Turk more easily into your processes and can support moving from batch processing of data to a real-time approach.

Common use cases include:

- Immediately add more assignments to a HIT when there is disagreement between workers that have responded.
- Update a database with the results of the HIT.
- Evaluate worker submissions as soon as they are submitted so you can take action if workers are consistently making errors.
- Chain multiple Mechanical Turk steps together by triggering the next step when a HIT is completed.

Notification event types

Notifications are associated with a HIT Type. You can request that Mechanical Turk send a notification when any of the following events occur for HITs using a given HIT Type.

- `AssignmentAccepted`: A worker has accepted a HIT and has an assignment.
- `AssignmentAbandoned`: An assignment has been abandoned because the assignment duration has elapsed.
- `AssignmentReturned`: A worker has chosen to return an assignment rather than complete the task.
- `AssignmentSubmitted`: A worker has submitted an assignment.
- `AssignmentRejected`: You have rejected an assignment.
- `AssignmentApproved`: An assignment has been approved.
- `HITCreated`: A HIT has been created using the HIT Type.
- `HITExtended`: Assignments have been added to a HIT.
- `HITDisposed`: A HIT has been disposed.
- `HITReviewable`: A HIT has reached the `Reviewable` state, either because all of the assignments have been submitted or the HIT has expired.
- `HITExpired`: The HIT has expired (the lifetime has elapsed) before all of the available assignments have been submitted.
- `Ping`: Is only be sent when using the `SendTestEventNotification` operation.

The `AssignmentSubmitted` and `HITReviewable` events are the most commonly used event notifications because they allow you to setup processes that can be triggered as soon as an assignment or HIT is complete. Using the `HITReviewable` notification, you can immediately update a database or other system with the values returned by the HIT. The `HITExpired` event is also useful because it can be treated as a dead letter queue, letting you act on tasks that weren't completed by workers. Using the `HITExpired` notification, you might sideline tasks to be completed by members of your team, or attempt to repost them at a higher reward amount.

Notification destination

Notifications can be sent to either an Amazon Simple Queue Service (Amazon SQS) queue or an Amazon Simple Notification Service (Amazon SNS) topic. In both cases, multiple events may be batched into a single message if appropriate.

To set up an Amazon SQS queue for use with Mechanical Turk, follow the setup instructions found in [Notification Handling Using Amazon SQS](#). To set up an Amazon SNS topic, follow the steps in [Notification Handling Using Amazon SNS](#). Note that in both cases, you need to configure permissions properly to allow Mechanical Turk to send messages to your topic or queue.

Enabling Notifications

Notifications can be enabled for a HIT Type using the `UpdateNotificationSettings` operation. The following is an example of a request to receive a message on an Amazon SNS topic when HITs become *reviewable*.

```
{
  'HITTypeId': '3AKE04YHPN13791QQA6EXAMPLE',
  'Notification': {
    'Destination': 'arn:aws:sns:us-east-1:7429088EXAMPLE:my_mturk_topic',
    'Transport': 'SNS',
    'Version': '2014-08-15',
    'EventTypes': ['HITReviewable']
  },
  'Active': True
}
```

The notification data structure specifies the ARN of the destination you want to receive the messages, specifies that the Transport is either Amazon SNS or Amazon SQS, and includes a list of the event types about which you want to be notified. The only valid value for `Version` is '2014-08-15'. The `Active` value indicates if the notification should be enabled.

Because only one notification configuration can be assigned to a HIT Type, calling `UpdateNotificationSettings` with a new value for notification replaces any existing notifications. You can call `UpdateNotificationSettings` with just the `HITTypeId` and a value for `Active` if you want to enable or disable notifications on the HIT Type.

Handling notifications using AWS Lambda

There are a wide range of approaches for using Mechanical Turk notifications, but the most common one is to use Amazon SNS with AWS Lambda. AWS Lambda provides a straightforward way to set up code that processes the event and attaches a trigger to kick off processing when new Amazon SNS messages are received.

Start by creating an Amazon SNS topic you can use for Mechanical Turk notifications. The instructions in [Notification Handling Using Amazon SNS](#) can be used to set it up and configure permissions to allow Mechanical Turk to send messages to the topic. Then, use the following procedure to create a lambda function to process your Amazon SNS messages.

To create a Lambda function to handle your Mechanical Turk events:

1. Navigate to Lambda in the AWS Console: <https://console.aws.amazon.com/lambda/>.
2. Select **Create Function**.
3. Select the **Author from scratch** option.
4. Under **Function name**, provide a name for your function.
5. You can select the language runtime you wish to use; in this example we use a Python 3 runtime.
6. You can keep the default permission configuration, which creates a new role for this function. If you want to use an existing role, or create a role from AWS policy templates, select the arrow next to **Change default execution role** to expand that section. Select one of the options. If you select

another option, make sure that role you use has required permissions described in [AWS Lambda execution role](#) in the AWS Lambda Developer Guide.

Note down the name of your execution role – you will use it in the following procedure.

7. Select **Create function** to create the lambda function.

Next, you must add the `AmazonMechanicalTurkFullAccess` policy to the execution role you created or used in step 6 of the preceding procedure so your Lambda function can retrieve the results of the HIT. The following procedure assumes that you are already in the Lambda console. If you are on the summary page for your new lambda function, you can skip the first step.

To add required permissions to your Lambda execution role:

1. Select the lambda function you want to use to process Amazon SNS requests. This brings you to the summary page for that function. You should see the function name at the top of the page.
2. Select the **Permissions** tab.
3. Select the **Role name**. This redirects you to summary page for that role in the IAM console.
4. Choose **Attach Policies**.
5. In the search field, enter `AmazonMechanicalTurkFullAccess` and select the check box next to that policy.
6. Select **Attach policy**.

Now that you have created a lambda function with permission to process Amazon SNS notifications send from Amazon Mechanical Turk, you can create a trigger for your Amazon SNS topic. A trigger is used to configure the conditions under which your function is called.

The following procedure assumes you are in the Lambda console.

To create a lambda trigger for your Amazon SNS topic:

1. On the lambda function summary page, select the **Configuration** tab.
2. Choose **Add trigger** to add your Amazon SNS topic as a trigger for your Lambda function.
3. When prompted to select a trigger service, search for and select **SNS**, then select the topic you created for notifications.
4. Enter the code that you want to use to process Amazon SNS messages. The following Python code provides a template for getting started.

```
import json
import boto3
import xml.etree.ElementTree as ET

def lambda_handler(event, context):
    for record in event['Records']:
        notification = json.loads(record['Sns']['Message'])

        for mturk_event in notification['Events']:
            mturk = boto3.client('mturk', region_name='us-east-1')

            if mturk_event['EventType'] == 'HITReviewable':
                # Retrieve the answers that were provided by Workers
                response = mturk.list_assignments_for_hit(HITId=mturk_event['HITId'])
                assignments = response['Assignments']
                answers = []
                for assignment in assignments:
                    answers.append(parse_answers(assignment))
```

```
        # Do something with the answers
        # ...

# Function to parse the Answer XML object
def parse_answers(assignment):
    result = {
        'WorkerId': assignment['WorkerId'],
        'Answer': []
    }

    ns = {'mt': 'http://mechanicalturk.amazonaws.com/
AWSMechanicalTurkDataSchemas/2005-10-01/QuestionFormAnswers.xsd'}
    root = ET.fromstring(assignment['Answer'])

    for a in root.findall('mt:Answer', ns):
        name = a.find('mt:QuestionIdentifier', ns).text
        value = a.find('mt:FreeText', ns).text
        result['Answer'].append({'name': value})
    return result
```

5. After you've added the code to your Lambda function and edited it to meet your needs, you can select **Deploy** and begin using it.

It is recommended that you test your trigger to make sure it works as expected.

To test your Amazon SNS trigger on your Lambda function summary page:

1. In the **Function code** section, select the arrow next to **Test**.
2. Select **Configure test event**.
3. Select the **Create new test event** radio button.
4. Enter a name for your event in the text box.
5. Enter the test event. The following can be used for your test parameters after replacing the `HITID` with a completed HIT in your account.

```
{
  "Records": [
    {
      "Sns": {
        "Message": "{\"Events\": [{\"EventType\": \"HITReviewable\", \"HITID\": \"31ANT7FQ71LA48IEXAMPLE\"}]}"
      }
    }
  ]
}
```

Sending test events

To test the configuration of your Amazon SNS topic or Amazon SQS queue and any handlers you have in place, you can use the [SendTestEventNotification](#) operation. Provide the notification configuration you want to use and the event you would like to test.

```
{
  'Notification': {
```

```
'Destination': 'arn:aws:sns:us-east-1:7429088EXAMPLE:my_mturk_topic',  
'Transport': 'SNS',  
'Version': '2014-08-15',  
'EventTypes': ['HITReviewable']  
},  
'TestEventType': 'HITReviewable'  
}
```

HIT references using requester annotation

When building processes that leverage Amazon Mechanical Turk (Mechanical Turk), it's often valuable to keep track of identifiers associated with the data in each HIT, particularly when handling HIT responses via notifications. For example, you might want to associate your HITs with a record in a database such as Amazon DynamoDB, and would like your HIT to reference the primary key of the record.

The `RequesterAnnotation` attribute is a useful option for tracking these references. When you create a HIT using [CreateHIT](#) or [CreateHITWithHITType](#), you can provide a `RequesterAnnotation` field that contains arbitrary data about each HIT. Although it is limited to 255 ASCII characters, this is generally adequate to capture identifiers that denote the origin of your data. The data provided here is only visible to the requester who created the HIT.

When you receive a notification that a HIT has been completed, you can use the [GetHIT](#) operation to retrieve the `RequesterAnnotation`. The identifier captured in the `RequesterAnnotation` can then be used to make updates in your database or other systems.

Use request tokens

Many of the API operations that have an impact on HITs or the money rewarded to workers, such as `CreateHIT` and `SendBonus`, can include a `UniqueRequestToken` attribute. This denotes a unique identifier for the request that can be used in scenarios where you want to gracefully handle and retry errors without creating duplicate HITs or payments. This is useful in cases such as network timeouts, where it is unclear whether or not the call succeeded on the server. If the operation has already been performed using the same `UniqueRequestToken`, subsequent calls return an error with a message containing the request ID.

Note that the token must not be longer than 64 characters in length and it is your responsibility to ensure uniqueness of the token. The unique token expires after 24 hours.

CORS configuration requirement

Earlier in 2020, widely used browsers like Chrome and Firefox changed their default behavior for rotating images based on image metadata, referred to as [EXIF data](#). Previously, images would always display in browsers exactly how they are stored on disk, which is typically unrotated. After the change, images now rotate according to a piece of image metadata called *orientation value*. This has important implications for the entire machine learning (ML) community. For example, if the EXIF orientation is not considered, applications that are used to annotate images may display images in unexpected orientations and result in incorrect labels.

Starting with Chrome 89, AWS can no longer automatically prevent the rotation of images because the web standards group W3C has decided that the ability to control rotation of images violates the web's Same Origin Policy. Therefore, to ensure human workers annotate your input images in a predictable orientation when you submit requests to label an image, you must add a CORS header policy to the Amazon S3 buckets that contain your input images.

Important

If you do not add a CORS configuration to the S3 buckets that contains your input data, tasks for those input data objects will fail.

You can add a CORS policy to an S3 bucket that contains input data in the S3 console. To set the required CORS headers on the S3 bucket that contain your input images in the S3 console, follow the directions detailed in [How do I add cross-domain resource sharing with CORS?](#). Use the following CORS configuration code for the buckets that hosts your images. You must use the JSON format if you add a CORS configuration using the console.

JSON

```
[{
  "AllowedHeaders": [],
  "AllowedMethods": ["GET"],
  "AllowedOrigins": ["*"],
  "ExposeHeaders": []
}]
```

XML

```
<CORSConfiguration>
  <CORSRule>
    <AllowedOrigin>*</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
  </CORSRule>
</CORSConfiguration>
```