



Developer Guide

Amazon Comprehend Medical



Amazon Comprehend Medical: Developer Guide

Copyright © 2026 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 Comprehend Medical?	1
Important notice	1
Amazon Comprehend Medical use cases	1
Benefits of Amazon Comprehend Medical	2
HIPAA compliance	3
Accessing Amazon Comprehend Medical	3
How to get started with Amazon Comprehend Medical	4
How it works	5
Synchronous entity detection	5
Asynchronous batch analysis	5
Ontology linking	6
Linking to concepts in the ICD-10-CM knowledge base of medical conditions	6
Linking to concepts in the RxNorm knowledge base of medications	6
Linking to concepts in the SNOMED CT knowledge base of medical concepts	7
Getting Started	8
Step 1: Set Up an Account	8
Sign up for AWS	8
Create an IAM User	9
Next step	10
Step 2: Set Up the AWS CLI	10
Next step	10
Step 3: Getting Started Using the Console	11
Analyzing clinical text using the console	11
Step 4: Getting Started Using the API	13
Detecting medical entities using the AWS Command Line Interface	14
Detecting medical entities using the AWS SDK for Java	16
Detecting medical entities using the AWS SDK for Python (Boto)	17
VPC endpoints (AWS PrivateLink)	19
Considerations for Amazon Comprehend Medical VPC endpoints	19
Creating an interface VPC endpoint for Amazon Comprehend Medical	19
Creating a VPC endpoint policy for Amazon Comprehend Medical	20
Text analysis APIs	21
Detect entities (Version 2)	21
Anatomy category	23

Behavioral, environmental, and social health category	23
Medical condition category	24
Medication category	25
Protected health information category	26
Test, treatment, and procedure category	26
Time expression category	27
Detect PHI	28
Text analysis batch APIs	33
Important notice	1
Performing batch analysis using the APIs	33
Performing batch analysis using the console	34
IAM policies	34
Batch analysis output files	36
Ontology Linking APIs	39
InferICD10CM	39
ICD-10-CM category	39
ICD-10-CM types	39
ICD-10-CM traits	40
ICD-10-CM attributes	40
Time expression category	40
Types	41
Relationship type	41
Input and response examples	41
InferRxNorm	44
Important notice	1
RxNorm category	44
RxNorm types	45
RxNorm attributes	45
RxNorm traits	45
Input and response examples	45
InferSNOMEDCT	48
Anatomy category	49
Medical condition category	49
Test, treatment, and procedure category	50
SNOMED CT details	51
Input and response examples	52

Ontology linking batch analysis	59
Performing batch analysis	59
IAM Policies	60
Batch analysis output files	62
Security	65
Data Protection	66
Authentication and access control	67
Authentication	67
Access Control	68
Overview of managing access	68
Using Identity-Based policies (IAM policies) for Amazon Comprehend Medical	71
Amazon Comprehend Medical API Permissions Reference	79
AWS managed policies	79
CloudTrail logs	81
Amazon Comprehend Medical information in CloudTrail	82
Understanding Amazon Comprehend Medical log file entries	83
Compliance Validation	84
Resilience	85
Infrastructure Security	85
Guidelines and quotas	86
Important notice	86
Supported regions	86
Throttling	86
Overall quotas	86
Document history	92
Amazon Comprehend Medical change log	102
Updated API version for DetectEntitiesV2	102
Updated API version for DetectEntities	102
Updated API version for InferICD10CM	103
Updated API version for InferRxNorm	103
Updated API version for InferSNOMEDCT	103
Updated API version for InferICD10CM	103
Updated API version for InferRxNorm	103
Updated API version for InferSNOMEDCT	104
Updated API version for InferICD10CM	104
Updated API version for InferRxNorm	104

Updated API version for InferSNOMEDCT	104
Updated API version for DetectEntities	104
Updated API version for DetectEntitiesV2	105
Updated API version for InferRxNorm	105
Updated API version for InferSNOMEDCT	105
Updated API version for InferICD10CM	105
Changes to the Amazon Comprehend Medical API operations	105
Changes to the Amazon Comprehend Medical API operations	106
Changes to the Amazon Comprehend Medical API operation	106
Changes to the Amazon Comprehend Medical API operations	106
Changes to the Amazon Comprehend Medical API operations	107
Changes to the Amazon Comprehend Medical API operations	107
Changes to the Amazon Comprehend Medical API operations	107
Changes to Amazon Comprehend Medical API operations	107
Changes to the Amazon Comprehend Medical InferSNOMEDCT API operations	108
Changes to Amazon Comprehend Medical API operations	108

What is Amazon Comprehend Medical?

Amazon Comprehend Medical detects and returns useful information in unstructured clinical text such as physician's notes, discharge summaries, test results, and case notes. Amazon Comprehend Medical uses natural language processing (NLP) models to detect entities, which are textual references to medical information such as medical conditions, medications, or Protected Health Information (PHI). For a full list of detected entities, see [Detect entities \(Version 2\)](#). Amazon Comprehend Medical also enables users to link these detected entities to standardized medical knowledge bases such as RxNorm and ICD-10-CM through ontology linking operations.

The information in this developer guide is intended for application developers. This guide includes information about using Amazon Comprehend Medical programmatically through either the AWS CLI or the Amazon Comprehend Medical APIs.

Pricing for Amazon Comprehend Medical differs from Amazon Comprehend pricing. For more information, see [Amazon Comprehend Medical Pricing](#).

Supported Languages

Amazon Comprehend Medical only detects medical entities in English language (US-EN) texts.

Important notice

Amazon Comprehend Medical is not a substitute for professional medical advice, diagnosis, or treatment. Amazon Comprehend Medical provides confidence scores that indicate the level of confidence in the accuracy of the detected entities. Identify the right confidence threshold for your use case, and use high confidence thresholds in situations that require high accuracy. In certain use cases, results should be reviewed and verified by appropriately trained human reviewers. For example, Amazon Comprehend Medical should only be used in patient care scenarios after review for accuracy and sound medical judgment by trained medical professionals.

Amazon Comprehend Medical use cases

You can use Amazon Comprehend Medical for the following healthcare applications:

- **Patient case management and outcome**— Doctors and healthcare providers can manage and easily access medical information that doesn't fit into traditional forms. Patients can report their health concerns in a narrative with more information than standard formats. By analyzing case

notes, providers can identify candidates for early screening of medical conditions before the condition becomes more difficult and expensive to treat.

- **Clinical research**—Life sciences and research organizations can optimize the matching process for enrolling patients into clinical trials. By using Amazon Comprehend Medical to detect pertinent information in clinical text, researchers can improve pharmacovigilance, perform post-market surveillance to monitor adverse drug events, and assess therapeutic effectiveness by easily detecting vital information in follow-up notes and other clinical texts. For instance, it can be easier and more effective to monitor how patients respond to certain therapies by analyzing their narratives.
- **Medical billing and healthcare revenue cycle management**—Payors can expand their analytics to include unstructured documents such as clinical notes. More information about a diagnosis can be analyzed and used to help determine appropriate billing codes from unstructured documents. Natural language processing (NLP) is the most critical component of computer-assisted coding (CAC). Amazon Comprehend Medical uses the latest advances in NLP to analyze clinical text, helping to decrease time to revenue and improve reimbursement accuracy.
- **Ontology linking** —Use the ontology linking features to detect entities from clinical text and link those entities to standardized concepts in common medical ontologies. **InferICD10CM** identifies possible medical conditions as entities. **InferICD10CM** links those entities to unique codes from the 2021 version of the [International Classification of Diseases, 10th Revision, Clinical Modification \(ICD-10-CM\)](#). **InferRxNorm** identifies medications listed in clinical text as entities and links those entities to normalized concept identifiers from the [RxNorm database from the US National Library of Medicine](#). **InferSNOMEDCT** detects medical concepts such as medical conditions and anatomy, medical tests, or treatments and procedures, as entities and links them to codes from the [Systematized Nomenclature of Medicine, Clinical Terms \(SNOMED CT\)](#) ontology.

Benefits of Amazon Comprehend Medical

Some of the benefits of using Amazon Comprehend Medical include:

- **Easy, powerful natural language processing integration into your applications**— Use APIs to build text analysis capabilities into your applications for powerful and accurate natural language processing.
- **Accuracy**— Use deep learning technology to accurately analyze text. Our models are constantly trained with new data across multiple domains to improve accuracy.

- **Scalability**— Detect information from multiple documents, making rapid insights into patient health and care possible.
- **Integrate with other AWS services**—Amazon Comprehend Medical is designed to work seamlessly with other AWS services like Amazon S3 and AWS Lambda. Store your documents in Amazon S3, analyze real-time data with Firehose, or use Amazon Transcribe to transcribe patient narratives into text that can be analyzed by Amazon Comprehend Medical. Support for AWS Identity and Access Management (IAM) makes it easy to securely control access to Amazon Comprehend Medical operations. Using IAM, you can create and manage AWS users and groups to grant the appropriate access to your developers and end users.
- **Low cost**— Only pay for the documents you analyze. There are no minimum fees or upfront commitments.

HIPAA compliance

This is a HIPAA Eligible Service. For more information about AWS, U.S. Health Insurance Portability and Accountability Act of 1996 (HIPAA), and using AWS services to process, store, and transmit protected health information (PHI), see [HIPAA Overview](#).

Connections to Amazon Comprehend Medical containing PHI must be encrypted. By default, all connections to Amazon Comprehend Medical use HTTPS over TLS. Amazon Comprehend Medical does not persistently store customer content. Therefore, you do not need to configure encryption at-rest within the service.

Accessing Amazon Comprehend Medical

1. **AWS Management Console**— Provides a web interface that you can use to access Amazon Comprehend Medical.
2. **AWS Command Line Interface (AWS CLI)** – Provides commands for a broad set of AWS services, including Amazon Comprehend Medical, and is supported on Windows, macOS, and Linux. For more information about installing the AWS CLI, see [AWS Command Line Interface](#).
3. **AWS SDKs** – AWS provides SDKs (software development kits) that consist of libraries and sample code for various programming languages and platforms (Java, Python, Ruby, .NET, iOS, Android, etc.). The SDKs provide a convenient way to create programmatic access to Amazon Comprehend Medical and AWS. For more information, see [AWS SDKs](#).

How to get started with Amazon Comprehend Medical

If you are a first-time user of Amazon Comprehend Medical, we recommend that you read the following sections in order:

1. [How Amazon Comprehend Medical works](#) – This section introduces Amazon Comprehend Medical concepts.
2. [Getting started with Amazon Comprehend Medical](#) – This section explains how to set up your account and test Amazon Comprehend Medical.

How Amazon Comprehend Medical works

Amazon Comprehend Medical uses a pretrained natural language processing (NLP) model to analyze unstructured clinical text through entity detection. An entity is a textual reference to medical information such as medical conditions, medications, or Protected Health Information (PHI). Some operations go one step further by detecting entities and then linking those entities to standardized ontologies. The model is continuously trained on a large body of medical texts, so you don't need to provide training data. All results include a confidence score, which indicates the confidence that Amazon Comprehend Medical has in the accuracy of the detected entities.

Both entity detection and ontology linking can be performed as either synchronous or asynchronous operations:

- **Synchronous operations**— Enables analysis on single documents which return the results of the analysis directly to your applications. Use the single-document operations when you are creating an interactive application that works on one document at a time.
- **Asynchronous operations** — Enables analysis on a collection or batch of documents stored in an Amazon S3 bucket. The results of the analysis are returned in an S3 bucket.

Note

Amazon Comprehend Medical can analyze only text in English (US-EN).

Synchronous entity detection

The **DetectEntitiesV2** and **DetectPHI** operations detect entities in unstructured clinical text from individual documents. You send a document to the Amazon Comprehend Medical service and receive the results of the analysis in the response.

Asynchronous batch analysis

The **StartEntitiesDetectionV2Job** and **StartPHIDetectionJob** operations start asynchronous jobs to detect references to medical information such as medical condition, treatment, tests and results or protected health information that is stored in an Amazon S3 bucket. The output of the detection

job is written to a separate Amazon S3 bucket from which it can be used for further processing or downstream analysis.

The **StartICD10CMInferenceJob**, and **StartRxNormInferenceJob** operations start ontology linking batch operations which detect entities and link those entities to standardized codes in the RxNorm and ICD-10-CM knowledge bases.

Ontology linking

The **InferICD10CM**, **InferSNOMEDCT**, and **InferRxNorm** operations detect potential medical conditions and medications and link them to codes in the ICD-10-CM, SNOMED CT, or RxNorm knowledge bases, respectively. You can use ontology linking batch analysis to analyze either a collection of documents or a single large document. By using either the console or the ontology linking batch APIs, you can perform operations to start, stop, list, and describe ongoing batch analysis jobs.

Linking to concepts in the ICD-10-CM knowledge base of medical conditions

The **InferICD10CM** operation detects potential medical conditions and links them to codes from the 2019 version of the International Classification of Diseases, 10th Revision, Clinical Modification (ICD-10-CM). For each potential medical condition detected, Amazon Comprehend Medical lists the matching ICD-10-CM codes and descriptions. Listed medical conditions in the results include a confidence score, which indicates the confidence that Amazon Comprehend Medical has in the accuracy of the entities to the matched concepts in the results.

Linking to concepts in the RxNorm knowledge base of medications

The **InferRxNorm** operation identifies medications that are listed in a patient record as entities. It links entities to concept identifiers (RxCUI) from the RxNorm database from the National Library of Medicine. Each RxCUI is unique for different strengths and dose forms. Listed medications in the results include a confidence score, which indicates the confidence that Amazon Comprehend Medical has in the accuracy of the entities matched to the concepts from the RxNorm knowledge base. Amazon Comprehend Medical lists the top RxCUIs that potentially match for each medication that it detects in descending order based on confidence score.

Linking to concepts in the SNOMED CT knowledge base of medical concepts

The InferSNOMEDCT operation identifies possible medical concepts as entities and links them to codes from the 2021-03 version of the Systematized Nomenclature of Medicine, Clinical Terms (SNOMED CT). SNOMED CT provides a comprehensive vocabulary of medical concepts, including medical conditions and anatomy, as well as medical tests, treatments, and procedures. For each matched concept ID, Amazon Comprehend Medical returns the top five medical concepts, each with a confidence score and contextual information such as traits and attributes. The SNOMED CT concept IDs can then be used to structure patient clinical data for medical coding, reporting, or clinical analytics when used with the SNOMED CT polyhierarchy.

Getting started with Amazon Comprehend Medical

To get started using Amazon Comprehend Medical, set up an AWS account and create an AWS Identity and Access Management (IAM) user. To use the Amazon Comprehend Medical CLI, download and configure it.

Topics

- [Step 1: Set up an AWS account and create an administrator user](#)
- [Step 2: Set up the AWS Command Line Interface \(AWS CLI\)](#)
- [Step 3: Getting started using the Amazon Comprehend Medical console](#)
- [Step 4: Getting started using the Amazon Comprehend Medical APIs](#)

Step 1: Set up an AWS account and create an administrator user

Before you use Amazon Comprehend Medical for the first time, complete the following tasks:

1. [Sign up for AWS](#)
2. [Create an IAM User](#)

Sign up for AWS

When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all AWS services, including Amazon Comprehend Medical. You are charged only for the services that you use.

With Amazon Comprehend Medical, you pay only for the resources that you use. If you are a new AWS customer, you can get started with Amazon Comprehend Medical for free. For more information, see [AWS Free Usage Tier](#).

If you already have an AWS account, skip to the next section.

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 or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

Record your AWS account ID because you'll need it for the next task.

Create an IAM User

Services in AWS, such as Amazon Comprehend Medical, require that you provide credentials when you access them. This allows the service to determine whether you have permissions to access the service's resources.

We strongly recommend that you access AWS using AWS Identity and Access Management (IAM), not the credentials for your AWS account. To use IAM to access AWS, create an IAM user, add the user to an IAM group with administrative permissions, and then grant administrative permissions to the IAM user. You can then access AWS using a special URL and the IAM user's credentials.

The Getting Started exercises in this guide assume that you have a user with administrator privileges, `adminuser`.

To create an administrator and sign in to the console

1. Create a user named `adminuser` in your AWS account. For instructions, see [Creating Your First IAM User and Administrators Group](#) in the *IAM User Guide*.
2. Sign in to the AWS Management Console using a special URL. For more information, see [How Users Sign In to Your Account](#) in the *IAM User Guide*.

For more information about IAM, see the following:

- [AWS Identity and Access Management \(IAM\)](#)
- [Getting started with IAM](#)
- [IAM User Guide](#)

Next step

[Step 2: Set up the AWS Command Line Interface \(AWS CLI\)](#)

Step 2: Set up the AWS Command Line Interface (AWS CLI)

You don't need the AWS CLI to perform the steps in the Getting Started exercises. However, some of the other exercises in this guide do require it. If you prefer, you can skip this step and go to [Step 3: Getting started using the Amazon Comprehend Medical console](#), and set up the AWS CLI later.

To set up the AWS CLI

1. Download and configure the AWS CLI. For instructions, see the following topics in the *AWS Command Line Interface User Guide*:
 - [Getting Set Up with the AWS Command Line Interface](#)
 - [Configuring the AWS Command Line Interface](#)
2. In the AWS CLI config file, add a named profile for the administrator:

```
[profile adminuser]
aws_access_key_id = adminuser access key ID
aws_secret_access_key = adminuser secret access key
region = aws-region
```

You use this profile when executing the AWS CLI commands. For more information about named profiles, see [Named Profiles](#) in the *AWS Command Line Interface User Guide*. For a list of AWS Regions, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

3. Verify the setup by typing the following help command at the command prompt:

```
aws help
```

Next step

[Step 3: Getting started using the Amazon Comprehend Medical console](#)

Step 3: Getting started using the Amazon Comprehend Medical console

The easiest way to get started using the Comprehend Medical console is to analyze a short text file. If you haven't reviewed the concepts and terminology in [How Amazon Comprehend Medical works](#), we recommend that you do that before proceeding.

Topics

- [Analyzing clinical text using the console](#)

Analyzing clinical text using the console

The Comprehend Medical console enables you to analyze the contents of clinical text, up to 20,000 characters long. The results are shown in the console so that you can review the analysis.

To start analyzing documents, sign in to the AWS Management Console and open the Comprehend Medical console.

Under **Comprehend Medical**, choose **Real-time analysis**.

The console displays sample text and the analysis of that text:

The screenshot shows the Amazon Comprehend Medical console interface. At the top, it says "Amazon Comprehend Medical > Real-time analysis". Below that is the heading "Real-time analysis" with an "Info" link. A sub-heading "Input text" is followed by a link for "Supported languages". A large text area contains the following sample text: "Pt is 87 yo woman, highschool teacher with past medical history that includes - status post cardiac catheterization in April 2019. She presents today with palpitations and chest pressure. HPI : Sleeping trouble on present dosage of Clonidine. Severe Rash on face and leg, slightly itchy. Meds : Vyvanse 50 mgs po at breakfast daily, Clonidine 0.2 mgs -- 1 and 1 / 2 tabs po qhs. HEENT : Boggy inferior turbinates, No oropharyngeal lesion". Below the text area, it indicates "415 of 20000 characters used." At the bottom right, there are two buttons: "Clear text" and "Analyze".

You can replace the sample text with your own text in English and then choose **Analyze** to get an analysis of your text.

The screenshot displays the 'Insights' section of the Amazon Comprehend Medical console. It features a navigation bar with 'Entities' selected, and tabs for 'RxNorm concepts', 'ICD-10-CM concepts', and 'SNOMED CT concepts'. The 'Analyzed text' pane shows the following text with color-coded entity tags and relationships:

Pt is **87** yo woman, **highschool teacher** with past medical history that includes

- Age (87) (Orange tag)
- Dx name (highschool teacher) (Green tag)
- Profession (highschool teacher) (Orange tag)

Relationships shown:

- 'Overlap' between '87' and 'highschool teacher'.
- 'Overlap' between 'highschool teacher' and 'cardiac catheterization in'.
- 'Overlap' between 'cardiac catheterization in' and 'April 2019'.

Other entities and relationships:

- '- status post' (Blue tag) with 'Time to procedure name (status post)' (Pink tag).
- 'cardiac catheterization in' (Blue tag) with 'Procedure name (cardiac catheterization)' (Blue tag).
- 'April 2019.' (Pink tag) with 'Time to procedure name (April 2019)' (Pink tag) and 'Date (April 2019)' (Orange tag).
- 'Sign' relationship between 'palpitations' and 'chest pressure'.
- 'palpitations' (Green tag) with 'Dx name (palpitations)' (Green tag).
- 'chest pressure.' (Blue tag) with 'System organ site (chest)' (Purple tag).

Below the input text, the analyzed text is color-coded to indicate the entity category:


- Orange tags identify PHI data.
- Red tags identify Medication.
- Green tags identify Medical Condition.
- Blue tags identify Test, Treatment, or Procedure (TTP).
- Purple tags identify Anatomy.
- Pink tags identify Time Expressions.

For more information, see [How Amazon Comprehend Medical works](#).

In the console, below the input box, the **Analyzed Text** pane shows more information about the text.

The **Entity** section displays cards for the entities found in the text:

▼ Results (27)

Entity	Type	Category	Traits
87 0.9997 score	● Age	Protected health information	-
highschool teacher 0.7305 score	● Dx name	Medical condition	-
highschool teacher 0.2063 score	● Profession	Protected health information	-
 status post 0.9525 score	● Time to procedure name	Time expression	-
cardiac catheterization 0.8982 score	● Procedure name	Test treatment procedure	-

Each card shows the text and its entity type.

Next to each of the entities, a score represents the confidence that Comprehend Medical has in the identification of the text as the type of entity shown.

To see the JSON structure of both the request and the results, choose **Application integration**. The JSON structure is the same as the structure returned by the operation.

Next Step

[Step 4: Getting started using the Amazon Comprehend Medical APIs](#)

Step 4: Getting started using the Amazon Comprehend Medical APIs

The following examples demonstrate how to use Amazon Comprehend Medical operations using the AWS CLI, Java, and Python. Use them to learn about Amazon Comprehend Medical operations and as building blocks for your own applications.

To run the AWS CLI and Python examples, install the AWS CLI. For more information, see [Step 2: Set up the AWS Command Line Interface \(AWS CLI\)](#).

To run the Java examples, install the AWS SDK for Java. For instructions for installing the AWS SDK for Java, see [Set up the AWS SDK for Java](#).

Topics

- [Detecting medical entities using the AWS Command Line Interface](#)
- [Detecting medical entities using the AWS SDK for Java](#)
- [Detecting medical entities using the AWS SDK for Python \(Boto\)](#)

Detecting medical entities using the AWS Command Line Interface

The following example demonstrates using the DetectEntitiesV2 operation using the AWS CLI to return the medical entities detected in text. To run the example, you must install the AWS CLI. For more information, see [the section called "Step 2: Set Up the AWS CLI"](#).

The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```
aws comprehendmedical detect-entities-v2 \  
  --endpoint endpoint \  
  --region region \  
  --text "aspirin is required 20 mg po daily for 2 times as tab"
```

The response is the following:

```
{  
  "Entities": [  
    {  
      "Category": "MEDICATION",  
      "BeginOffset": 0,  
      "EndOffset": 7,  
      "Text": "aspirin",  
      "Traits": [],  
      "Score": 0.9988090991973877,  
      "Attributes": [  
        {  
          "BeginOffset": 20,  
          "EndOffset": 25,  
          "Text": "20 mg",  
          "Traits": [],  
          "Score": 0.9559056162834167,  
        }  
      ]  
    }  
  ]  
}
```

```
    "Type": "DOSAGE",
    "Id": 1,
    "RelationshipScore": 0.9981593489646912
  },
  {
    "BeginOffset": 26,
    "EndOffset": 28,
    "Text": "po",
    "Traits": [],
    "Score": 0.9995359182357788,
    "Type": "ROUTE_OR_MODE",
    "Id": 2,
    "RelationshipScore": 0.9969323873519897
  },
  {
    "BeginOffset": 29,
    "EndOffset": 34,
    "Text": "daily",
    "Traits": [],
    "Score": 0.9803128838539124,
    "Type": "FREQUENCY",
    "Id": 3,
    "RelationshipScore": 0.9990783929824829
  },
  {
    "BeginOffset": 39,
    "EndOffset": 46,
    "Text": "2 times",
    "Traits": [],
    "Score": 0.8623972535133362,
    "Type": "DURATION",
    "Id": 4,
    "RelationshipScore": 0.9996501207351685
  },
  {
    "BeginOffset": 50,
    "EndOffset": 53,
    "Text": "tab",
    "Traits": [],
    "Score": 0.784785270690918,
    "Type": "FORM",
    "Id": 5,
    "RelationshipScore": 0.9986748695373535
  }
}
```

```
        ],
        "Type": "GENERIC_NAME",
        "Id": 0
    }
],
"UnmappedAttributes": []
}
```

Detecting medical entities using the AWS SDK for Java

The following example uses the DetectEntitiesV2 operation with Java. To run the example, install the AWS SDK for Java. For instructions on installing the AWS SDK for Java, see [Set up the AWS SDK for Java](#).

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.auth.AWSStaticCredentialsProvider;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.client.builder.AwsClientBuilder;
import com.amazonaws.services.comprehendmedical.AWSComprehendMedical;
import com.amazonaws.services.comprehendmedical.AWSComprehendMedicalClient;
import com.amazonaws.services.comprehendmedical.model.DetectEntitiesRequest;
import com.amazonaws.services.comprehendmedical.model.DetectEntitiesResult;

public class SampleAPICall {

    public static void main() {

        AWSCredentialsProvider credentials
            = new AWSStaticCredentialsProvider(new BasicAWSCredentials("YOUR AWS
ACCESS KEY", "YOUR AWS SECRET"));

        AWSComprehendMedical client = AWSComprehendMedicalClient.builder()

.withCredentials(credentials)

.withRegion("YOUR
REGION")

.build();

        DetectEntitiesV2Request request = new DetectEntitiesV2Request();
        request.setText("cerealx 84 mg daily");
```

```

        DetectEntitiesV2Result result = client.detectEntitiesV2(request);
        result.getEntities().forEach(System.out::println);
    }
}

```

The output contains the three entities found in the input text, their location in the input text. A confidence level that the entity was correctly identified is also listed with each entity. The following output shows the `Generic_Name`, `Dosage`, and `Frequency` entities from the preceding example.

```

{Id: 0,BeginOffset: 0,EndOffset: 3,Score: 0.9940211,Text: Bob,Category:
PROTECTED_HEALTH_INFORMATION,Type: NAME,Traits: [],}
{Id: 2,BeginOffset: 23,EndOffset: 30,Score: 0.99914634,Text: aspirin,Category:
MEDICATION,Type: GENERIC_NAME,Traits: [],Attributes:
[{Type: DOSAGE,Score: 0.9630807,RelationshipScore: 0.99969745,Id: 1,BeginOffset:
14,EndOffset: 19,Text: 50 mg,Traits: []}]}

```

Detecting medical entities using the AWS SDK for Python (Boto)

The following example uses the `DetectEntitiesV2` operation with Python. To run the sample, install the AWS CLI. For more information, see [the section called “Step 2: Set Up the AWS CLI”](#).

```

import boto3
client = boto3.client(service_name='comprehendmedical', region_name='YOUR REGION')
result = client.detect_entities(Text= 'cerealx 84 mg daily')
entities = result['Entities']
for entity in entities:
    print('Entity', entity)

```

The output contains the three entities found in the input text, their location in the input text. A confidence level that the entity was correctly identified is also listed with each entity. The following output shows the `Generic_Name`, `Dosage`, and `Frequency` entities from the preceding example.

```

('Entity', {u'Category': u'MEDICATION', u'BeginOffset': 0, u'EndOffset': 7,
            u'Text': u'cerealx', u'Traits': [], u'Score': 0.8877691626548767,
            u'Attributes': [{u'BeginOffset': 8, u'EndOffset': 13,
                            u'Text': u'84 mg', u'Traits': [], u'Score': 0.9337134957313538, u'Type':
u'DOSAGE', u'Id': 1, u'RelationshipScore': 0.9995118379592896},
                            {u'BeginOffset': 14, u'EndOffset': 19, u'Text': u'daily', u'Traits': [],
u'Score': 0.990627646446228, u'Type': u'FREQUENCY',

```

```
u'Id': 2, u'RelationshipScore': 0.9987651109695435}], u'Type':  
u'BRAND_NAME', u'Id': 0})
```

Amazon Comprehend Medical and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your VPC and Amazon Comprehend Medical by creating an *interface VPC endpoint*. Interface VPC endpoints are powered by [AWS PrivateLink](#), a technology that you can use to privately access Amazon Comprehend Medical APIs without an internet gateway, NAT device, VPN connection, or Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with Amazon Comprehend Medical APIs. Traffic between your VPC and Amazon Comprehend Medical does not leave the Amazon network.

Each interface endpoint is represented by one or more [Elastic Network Interfaces](#) in your subnets.

For more information, see [Interface VPC endpoints \(AWS PrivateLink\)](#) in the *Amazon VPC User Guide*.

Considerations for Amazon Comprehend Medical VPC endpoints

Before you set up an interface VPC endpoint for Amazon Comprehend Medical, ensure that you review [Interface endpoint properties and limitations](#) in the *Amazon VPC User Guide*.

Amazon Comprehend Medical supports making calls to all of its API actions from your VPC.

Creating an interface VPC endpoint for Amazon Comprehend Medical

You can create a VPC endpoint for the Amazon Comprehend Medical service using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

Create a VPC endpoint for Amazon Comprehend Medical using the following service name:

- `com.amazonaws.region.comprehendmedical`

If you turn on private DNS for the endpoint, you can make API requests to Amazon Comprehend Medical using its default DNS name for the Region. For example, *comprehendmedical.us-east-1.amazonaws.com*.

For more information, see [Accessing a service through an interface endpoint](#) in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for Amazon Comprehend Medical

You can attach an endpoint policy to your VPC endpoint that controls access to Amazon Comprehend Medical. The policy specifies the following information:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

For more information, see [Controlling access to services with VPC endpoints](#) in the *Amazon VPC User Guide*.

Example: VPC endpoint policy for Amazon Comprehend Medical actions

The following is an example of an endpoint policy for Amazon Comprehend Medical. When attached to an endpoint, this policy grants access to the Amazon Comprehend Medical `DetectEntitiesV2` action for all principals on all resources.

```
{
  "Statement": [
    {
      "Principal": "*",
      "Effect": "Allow",
      "Action": [
        "comprehendmedical:DetectEntitiesV2"
      ],
      "Resource": "*"
    }
  ]
}
```

Text analysis API operations

Use Amazon Comprehend Medical to examine clinical documents and to gain various insights about their content using pre-trained natural language processing (NLP) models. You can perform analysis both on single files or as a batch analysis on multiple files stored in an Amazon Simple Storage Service (S3) bucket.

With Amazon Comprehend Medical, you can perform the following on your documents:

- [Detect entities \(Version 2\)](#) — Examine unstructured clinical text to detect textual references to medical information such as medical condition, treatment, tests and results, and medications. This version uses a different model than the original Detect entities API, and there are a few changes in the output.
- [Detect PHI](#) — Examine unstructured clinical text to detect textual references to protected health information (PHI) such as names and addresses.

Amazon Comprehend Medical also includes multiple API operations that you can use to perform batch text analysis on clinical documents. To learn more about how to use these API operations, see [the section called “Text analysis batch APIs”](#).

Topics

- [Detect entities \(Version 2\)](#)
- [Detect PHI](#)
- [Text analysis batch APIs](#)

Detect entities (Version 2)

Use the **DetectEntitiesV2** to detect entities in single files or **StartEntitiesDetectionV2Job** for batch analysis on multiple files. You can detect entities in the following categories:

- **ANATOMY**: Detects references to the parts of the body or body systems and the locations of those parts or systems.
- **BEHAVIORAL_ENVIRONMENTAL_SOCIAL**: Detects the behaviors and conditions in the environment that impact a person's health. This includes tobacco usage, alcohol consumption, recreational drug usage, allergies, gender, and race/ethnicity.

- **MEDICAL_CONDITION**: Detects the signs, symptoms, and diagnoses of medical conditions.
- **MEDICATION**: Detects medication and dosage information on the patient.
- **PROTECTED_HEALTH_INFORMATION**: Detects the patient's personal information.
- **TEST_TREATMENT_PROCEDURE**: Detects the procedures that are used to determine a medical condition.
- **TIME_EXPRESSION**: Detects entities related to time when they are associated with a detected entity.

All six categories are detected by the **DetectEntitiesV2** operation. For analysis specific to detecting PHI, use **DetectPHI** on single files and **StartPHIDetectionJob** for batch analysis.

Amazon Comprehend Medical detects information in the following classes:

- *Entity*: A text reference to the name of relevant objects, such as people, treatments, medications, and medical conditions. For example, *ibuprofen*.
- *Category*: The generalized grouping to which an entity belongs. For example, *ibuprofen* is part of the **MEDICATION** category.
- *Type*: The type of entity detected within a single category. For example, *ibuprofen* is in the **GENERIC_NAME** type in the **MEDICATION** category.
- *Attribute*: Information related to an entity, such as the dosage of a medication. For example, *200 mg* is an attribute of the *ibuprofen* entity.
- *Trait*: Something that Amazon Comprehend Medical understands about an entity, based on context. For example, a medication has the **NEGATION** trait if a patient is not taking it.
- *Relationship Type*: The relationship between an entity and an attribute.

Amazon Comprehend Medical provides you the location of an entity in the input text. In the Amazon Comprehend console, it shows you the location graphically. When you use the API, it shows you the location by numerical offset.

Each entity and attribute includes a score that indicates the confidence level that Amazon Comprehend Medical has in the accuracy of the detection. Each attribute also has a relationship score. The score indicates the confidence level that Amazon Comprehend Medical has in the accuracy of the relationship between the attribute and its parent entity. Identify the appropriate confidence threshold for your use case. Use high-confidence thresholds in situations that require great accuracy. Filter out data that doesn't meet the threshold.

Anatomy category

The ANATOMY category detects references to the parts of the body or body systems and the locations of those parts or systems.

Types

- SYSTEM_ORGAN_SITE: Body systems, anatomic locations or regions, and body sites.

Attributes

- DIRECTION: Directional terms. For example, left, right, medial, lateral, upper, lower, posterior, anterior, distal, proximal, contralateral, bilateral, ipsilateral, dorsal, ventral, and so on.

Behavioral, environmental, and social health category

The BEHAVIORAL_ENVIRONMENTAL_SOCIAL category detects references to behaviors and conditions in the environment that impact a person's health.

Type

- ALCOHOL_CONSUMPTION: Defines the patient's alcohol consumption in terms of use status, frequency, amount, and duration.
- ALLERGIES: Defines the patient's allergies and responses to allergens.
- GENDER: An identification of the characteristics of gender identity.
- RACE_ETHNICITY: A social-political construct of a patient's identification with particular racial and ethnic groups.
- REC_DRUG_USE: Defines the patient's use of recreational drugs in terms of use status, frequency, amount, and duration.
- TOBACCO_USE: Defines the patient's tobacco usage in terms of use status, frequency, amount, and duration.

Attributes

The following detected attributes only apply to the types ALCOHOL_CONSUMPTION, TOBACCO_USE, and REC_DRUG_USE:

- **AMOUNT:** The amount of alcohol, tobacco, or recreational drug used.
- **DURATION:** How long the alcohol, tobacco, or recreational drug has been used.
- **FREQUENCY:** How often the alcohol, tobacco, or recreational drug is used.

Traits

The following detected traits only apply to the types `ALCOHOL_CONSUMPTION`, `ALLERGIES`, `TOBACCO_USE`, and `REC_DRUG_USE`:

- **NEGATION:** An indication that a result or action is negative or not being performed.
- **PAST_HISTORY:** An indication that use of alcohol, tobacco, or recreational drugs is from the patient's past (prior to the current encounter).

Medical condition category

The `MEDICAL_CONDITION` category detects the signs, symptoms, and diagnoses of medical conditions. The category has one entity type, four attributes, and four traits. One or more traits can be associated with a type. Contextual information about attributes and their relationship to the diagnosis is detected and mapped to the `DX_NAME` through `RELATIONSHIP_EXTRACTION`. For instance, from the text "chronic pain in left leg", "chronic" is detected as the attribute `ACUITY`, "left" is detected as the attribute `DIRECTION`, and "leg" is detected as the attribute `SYSTEM_ORGAN_SITE`. The relationships of each of these attributes are mapped to the medical condition entity "pain," along with a confidence score.

Types

- **DX_NAME:** All medical conditions listed. The `DX_NAME` type includes present illness, reason for visit, and medical history.

Attributes

- **ACUITY:** Determination of disease instance, such as chronic, acute, sudden, persistent, or gradual.
- **DIRECTION:** Directional terms. For example, left, right, medial, lateral, upper, lower, posterior, anterior, distal, proximal, contralateral, bilateral, ipsilateral, dorsal, or ventral.
- **SYSTEM_ORGAN_SITE:** Anatomical location.

- **QUALITY:** Any descriptive term of the medical condition, such as stage or grade.

Traits

- **DIAGNOSIS:** A medical condition that is determined as the cause or result of the symptoms. Symptoms can be found through physical findings, laboratory or radiological reports, or any other means.
- **HYPOTHETICAL:** An indication that a medical condition is expressed as a hypothesis.
- **LOW_CONFIDENCE:** An indication that a medical condition is expressed as having high uncertainty. This is not directly related to the confidence scores provided.
- **NEGATION:** An indication that a result or action is negative or not being performed.
- **PERTAINS_TO_FAMILY:** An indication that a medical condition is relevant to the patient's family, not the patient.
- **SIGN:** A medical condition that the physician reported.
- **SYMPTOM:** A medical condition that the patient reported.

Medication category

The **MEDICATION** category detects medication and dosage information for the patient. One or more attributes can apply to a type.

Types

- **BRAND_NAME:** The copyrighted brand name of the medication or therapeutic agent.
- **GENERIC_NAME:** The non-brand name, ingredient name, or formula mixture of the medication or therapeutic agent.

Attributes

- **DOSAGE:** The amount of medication ordered.
- **DURATION:** How long the medication should be administered.
- **FORM:** The form of the medication.
- **FREQUENCY:** How often to administer the medication.
- **RATE:** The administration rate of the medication (primarily for medication infusions or IVs).

- **ROUTE_OR_MODE:** The administration method of the medication.
- **STRENGTH:** The medication strength.

Traits

- **NEGATION:** Any indication that the patient is not taking a medication.
- **PAST_HISTORY:** An indication that a medication detected is from the patient's past (prior to current encounter).

Protected health information category

The `PROTECTED_HEALTH_INFORMATION` category detects the patient's personal information. See [Detect PHI](#) to learn more about this operation.

Types

- **ADDRESS:** All geographical subdivisions of an address of any facility, units, or wards within a facility.
- **AGE:** All components of age, spans of age, or any age mentioned. This includes those of a patient, family members, or others. The default is in years, unless otherwise noted.
- **EMAIL:** Any email address.
- **ID:** Social Security number, medical record number, facility identification number, clinical trial number, certificate or license number, vehicle or device number, the place of care, or provider. This also includes any biometric number of the patient, such as height, weight, or a lab value.
- **NAME:** All names. Typically, names of the patient, family, or provider.
- **PHONE_OR_FAX:** Any phone, fax, or pager number. Excludes named phone numbers, such as 1-800-QUIT-NOW and 911.
- **PROFESSION:** Any profession or employer that pertains to the patient or the patient's family. It does not include the profession of the clinician mentioned in the note.

Test, treatment, and procedure category

The `TEST_TREATMENT_PROCEDURE` category detects the procedures that are used to determine a medical condition. One or more attributes can be related to an entity of the `TEST_NAME` type.

Types

- **PROCEDURE_NAME**: Interventions as a one-time action performed on the patient to treat a medical condition or to provide patient care.
- **TEST_NAME**: Procedures performed on a patient for diagnostic, measurement, screening, or rating that might have a resulting value. This includes any procedure, process, evaluation, or rating to determine a diagnosis, to rule out or find a condition, or to scale or score a patient.
- **TREATMENT_NAME**: Interventions performed over a span of time for combating a disease or disorder. This includes groupings of medications, such as antivirals and vaccinations.

Attributes

- **TEST_VALUE**: The result of a test. Applies only to the **TEST_NAME** entity type.
- **TEST_UNIT**: The unit of measure that might accompany the value of the test. Applies only to the **TEST_NAME** entity type.

Traits

- **FUTURE**: An indication that a test, treatment, or procedure refers to an action or event that will occur after the subject of the notes.
- **HYPOTHETICAL**: An indication that a test, treatment, or procedure is expressed as a hypothesis.
- **NEGATION**: An indication that a result or action is negative or not being performed.
- **PAST_HISTORY**: An indication that a test, treatment, or procedure is from the patient's past (prior to current encounter).

Time expression category

The **TIME_EXPRESSION** category detects entities related to time. This includes entities such as dates and time expressions such as "three days ago," "today," "currently," "day of admission," "last month," or "16 days." Results in this category are only returned if they are associated with an entity. For example, "Yesterday, the patient took 200 mg of ibuprofen" would return `Yesterday` as a **TIME_EXPRESSION** entity that overlaps with **GENERIC_NAME** entity "ibuprofen." However, it would not be recognized as an entity in "yesterday, the patient walked their dog."

Types

- **TIME_TO_MEDICATION_NAME**: The date a medication was taken. The attributes specific to this type are **BRAND_NAME** and **GENERIC_NAME**.
- **TIME_TO_DX_NAME**: The date a medical condition occurred. The attribute for this type is **DX_NAME**.
- **TIME_TO_TEST_NAME**: The date a test was performed. The attribute for this type is **TEST_NAME**.
- **TIME_TO_PROCEDURE_NAME**: The date a procedure was performed. The attribute for this type is **PROCEDURE_NAME**.
- **TIME_TO_TREATMENT_NAME**: The date a treatment was administered. The attribute for this type is **TREATMENT_NAME**.

Relationship type

- The relationship between an entity and an attribute. The recognized **Relationship_type** is the following:

Overlap – The **TIME_EXPRESSION** concurs with the entity detected.

Detect PHI

Use the **DetectPHI** operation when you only want to detect Protected Health Information (PHI) data when scanning the clinical text. To detect all available entities in the clinical text use **DetectEntitiesV2**.

This API is best for a use case where only detecting PHI entities is required. For information about information in the non-PHI categories, see [Detect entities \(Version 2\)](#).

Important

Amazon Comprehend Medical provides confidence scores that indicate the level of confidence in the accuracy of the detected entities. Evaluate these confidence scores and identify the right confidence threshold for your use case. For specific compliance use cases, we recommend that you use additional human review or other methods to confirm the accuracy of detected PHI.

Under the HIPAA act, PHI that is based on a list of 18 identifiers must be treated with special care. Amazon Comprehend Medical detects entities associated with these identifiers but these entities don't map 1:1 to the list specified by the Safe Harbor method. Not all identifiers are contained in unstructured clinical text, but Amazon Comprehend Medical does cover all the relevant identifiers. These identifiers consist of data that can be used to identify an individual patient, including the following list. For more information, see [Health Information Privacy](#) on the *U.S. Government Health and Human Services* website.

Each PHI-related entity includes a score (`Score` in the response) that indicates the level of confidence Amazon Comprehend Medical has in the accuracy of the detection. Identify the right confidence threshold for your use case and filter out entities that do not meet it. When identifying occurrences of PHI, it may be better to use a low confidence threshold for filtering to capture more potential detected entities. This is especially true when not using the values of the detected entities in compliance use cases.

The following PHI-related entities can be detected by running the **DetectPHI** or **DetectEntitiesV2** operations:

Detected PHI Entities

Entity	Description	HIPAA Category
AGE	All components of age, spans of age, and any age mentioned, be it patient or family member or others involved in the note. Default is in years unless otherwise noted.	3. Dates related to an individual
DATE	Any date related to patient or patient care.	3. Dates related to an individual
NAME	All names mentioned in the clinical note, typically belonging to patient, family, or provider.	1. Name

Entity	Description	HIPAA Category
PHONE_OR_FAX	Any phone, fax, pager; excludes named phone numbers such as 1-800-QUIT-NOW as well as 911.	4. Phone number 5. FAX number
EMAIL	Any email address.	6. Email addresses
ID	Any sort of number associated with the identity of a patient. This includes their social security number, medical record number, facility identification number, clinical trial number, certificate or license number, vehicle or device number. It also includes biometric numbers, and numbers identifying the place of care or provider.	7. Social Security Number 8. Medical Record number 9. Health Plan number 10. Account numbers 11. Certificate/License numbers 12. Vehicle identifiers 13. Device numbers 16. Biometric information 18. Any other identifying characteristics
URL	Any web URL.	14. URLs
ADDRESS	This includes all geographical subdivisions of an address of any facility, named medical facilities, or wards within a facility.	2. Geographic location

Entity	Description	HIPAA Category
PROFESSION	Includes any profession or employer mentioned in a note as it pertains to the patient or the patient's family.	18. Any other identifying characteristics

Example

The text "Patient is John Smith, a 48-year-old teacher and resident of Seattle, Washington." returns:

- "John Smith" as an *entity* of type NAME in the PROTECTED_HEALTH_INFORMATION category.
- "48" as an *entity* of type AGE in the PROTECTED_HEALTH_INFORMATION category.
- "teacher" as an *entity* of type PROFESSION (identifying characteristic) in the PROTECTED_HEALTH_INFORMATION category.
- "Seattle, Washington" as an ADDRESS *entity* in the PROTECTED_HEALTH_INFORMATION category.

In the Amazon Comprehend Medical console, this is shown like this:

Patient is John Smith, a 48 year old teacher and resident
 ● Name (John Smith) ● Age (48) ● Profession (teacher)

of Seattle, Washington.
 ● Address (Seattle , Washington)

When using the **DetectPHI** operation, the response appears like this. When you use the **StartPHIDetectionJob** operation, Amazon Comprehend Medical creates a file in the output location with this structure.

```
{
  "Entities": [
    {
      "Id": 0,
      "BeginOffset": 11,
      "EndOffset": 21,
```

```
    "Score": 0.997368335723877,
    "Text": "John Smith",
    "Category": "PROTECTED_HEALTH_INFORMATION",
    "Type": "NAME",
    "Traits": []
  },
  {
    "Id": 1,
    "BeginOffset": 25,
    "EndOffset": 27,
    "Score": 0.9998362064361572,
    "Text": "48",
    "Category": "PROTECTED_HEALTH_INFORMATION",
    "Type": "AGE",
    "Traits": []
  },
  {
    "Id": 2,
    "BeginOffset": 37,
    "EndOffset": 44,
    "Score": 0.8661606311798096,
    "Text": "teacher",
    "Category": "PROTECTED_HEALTH_INFORMATION",
    "Type": "PROFESSION",
    "Traits": []
  },
  {
    "Id": 3,
    "BeginOffset": 61,
    "EndOffset": 68,
    "Score": 0.9629441499710083,
    "Text": "Seattle",
    "Category": "PROTECTED_HEALTH_INFORMATION",
    "Type": "ADDRESS",
    "Traits": []
  },
  {
    "Id": 4,
    "BeginOffset": 78,
    "EndOffset": 88,
    "Score": 0.38217034935951233,
    "Text": "Washington",
    "Category": "PROTECTED_HEALTH_INFORMATION",
    "Type": "ADDRESS",
```

```
        "Traits": []
    }
],
"UnmappedAttributes": []
}
```

Text analysis batch APIs

Use Amazon Comprehend Medical to analyze medical text stored in an Amazon S3 bucket. Analyze up to 10 GB of documents in one batch. You use the console to create and manage batch analysis jobs, or use batch APIs to detect medical entities, including protected health information (PHI). The APIs start, stop, list, and describe ongoing batch analysis jobs.

Pricing information for batch analysis and other Amazon Comprehend Medical operations can be found [here](#).

Important notice

The batch analysis operations of Amazon Comprehend Medical are not a substitute for professional medical advice, diagnosis, or treatment. Identify the right confidence threshold for your use case, and use high confidence thresholds in situations that require high accuracy. For certain use cases, results should be reviewed and verified by appropriately trained human reviewers. All operations of Amazon Comprehend Medical should only be used in patient care scenarios after review for accuracy and sound medical judgment by trained medical professionals.

Performing batch analysis using the APIs

You can run a batch analysis job using either the Amazon Comprehend Medical console or the Amazon Comprehend Medical Batch APIs.

Prerequisites

When you are using the Amazon Comprehend Medical API, create an AWS Identity Access and Management (IAM) policy and attach it to an IAM role. To learn more about IAM roles and trust policies, see [IAM Policies and Permissions](#).

1. Upload your data into an S3 bucket.
2. To start a new analysis job, use either the `StartEntitiesDetectionV2Job` operation or the `StartPHIDetectionJob` operation. When you start the job, tell Amazon Comprehend Medical

the name of the input S3 bucket that contains the input files and designate the output S3 bucket to write the files after batch analysis.

3. Monitor the progress of the job by using the console or the `DescribeEntitiesDetectionV2Job` operation or the `DescribePHIDetectionJob` operation. Additionally, `ListEntitiesDetectionV2Jobs` and `ListPHIDetectionJobs` enable you to see the status of all ontology linking batch analysis jobs.
4. If you need to stop a job in progress, use `StopEntitiesDetectionV2Job` or `StopPHIDetectionJob` to stop analysis.
5. To view the results of your analysis job, see the output S3 bucket that you configured when you started the job.

Performing batch analysis using the console

1. Upload your data into an S3 bucket.
2. To start a new analysis job, select the type of analysis you will be performing. Then provide the name of the S3 bucket that contains the input files and the name of the S3 bucket where you want to send the output files.
3. Monitor the status of your job while it is ongoing. From the console, you can view all batch analysis operations and their status, including when analysis was started and ended.
4. To see the results of your analysis job, see the output S3 bucket that you configured when you started the job.

IAM Policies for batch operations

The IAM role that calls the Amazon Comprehend Medical batch APIs must have a policy that grants access to the S3 buckets that contain the input and output files. It must also be assigned a trust relationship that enables the Amazon Comprehend Medical service to assume the role. To learn more about IAM roles and trust policies, see [IAM Roles](#).

The role must have the following policy.

JSON

```
{
```

```
"Version": "2012-10-17",
"Statement": [
  {
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::input-bucket/*"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::input-bucket",
      "arn:aws:s3:::output-bucket"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::output-bucket/*"
    ],
    "Effect": "Allow"
  }
]
```

The role must have the following trust relationship. It is recommended that you use the `aws:SourceAccount` and `aws:SourceArn` condition keys to prevent the confused deputy security issue. To learn more about the confused deputy problem and how to protect your AWS account, see [The confused deputy problem](#) in the IAM documentation.

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Principal":{"
        "Service":["
          "comprehendmedical.amazonaws.com"
        ]
      },
      "Action":"sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "account_id"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:comprehendmedical:us-
east-1:account_id:*"
        }
      }
    }
  ]
}
```

Batch analysis output files

Amazon Comprehend Medical creates one output file for each input file in the batch. The file has the extension `.out`. Amazon Comprehend Medical first creates a directory in the output S3 bucket using the *AwsAccountId-JobType-JobId* as the name, and then writes all of the output files for the batch to this directory. Amazon Comprehend Medical creates this new directory so that output from one job does not overwrite the output of another.

The output from a batch operation produces the same output as a synchronous operation. For examples of the output generated by Amazon Comprehend Medical, see [Detect entities \(Version 2\)](#).

Each batch operation produces three manifest files that contain information about the job.

- **Manifest** – Summarizes the job. Provides information about the parameters used for the job, the total size of the job, and the number of files processed.

- **success** – Provides information about the files that were successfully processed. Includes the input and output file name and the size of the input file.
- **unprocessed** – Lists files the batch job did not process, including error codes and error messages per file.

Amazon Comprehend Medical writes the files to the output directory that you specified for the batch job. The summary manifest file will be written to the output folder, along with a folder titled `Manifest_AccountId-Operation-JobId`. Within the manifest folder is a `success` folder, which contains the success manifest. Also included is a `failed` folder, which contains the unprocessed file manifest. The following sections show the structure of the manifest files.

Batch manifest file

The following is the JSON structure of the batch manifest file.

```
{
  "Summary" : {
    "Status" : "COMPLETED | FAILED | PARTIAL_SUCCESS | STOPPED",
    "JobType" : "EntitiesDetection | PHIDetection",
    "InputDataConfiguration" : {
      "Bucket" : "input bucket",
      "Path" : "path to files/account ID-job type-job ID"
    },
    "OutputDataConfiguration" : {
      "Bucket" : "output bucket",
      "Path" : "path to files"
    },
    "InputFileCount" : number of files in input bucket,
    "TotalMeteredCharacters" : total characters processed from all files,
    "UnprocessedFilesCount" : number of files not processed,
    "SuccessFilesCount" : total number of files processed,
    "TotalDurationSeconds" : time required for processing,
    "SuccessfulFilesListLocation" : "path to file",
    "UnprocessedFilesListLocation" : "path to file",
    "FailedJobErrorMessage": "error message or if not applicable,
      The status of the job is completed"
  }
}
```

Success manifest file

The following is the JSON structure of the file that contains information about successfully processed files.

```
{
  "Files": [{
    "Input": "input path/input file name",
    "Output": "output path/output file name",
    "InputSize": size in bytes of input file
  }, {
    "Input": "input path/input file name",
    "Output": "output path/output file name",
    "InputSize": size in bytes of input file
  }]
}
```

Unprocessed manifest file

The following is the JSON structure of the manifest file that contains information about unprocessed files.

```
{
  "Files" : [ {
    "Input": "file_name_that_failed",
    "ErrorCode": "error code for exception",
    "ErrorMessage": "explanation of the error code and suggestions"
  },
  { ...}
 ]
}
```

Ontology linking

Use Amazon Comprehend Medical to detect entities in clinical text and link those entities to concepts in standardized medical ontologies, including the RxNorm, ICD-10-CM, and SNOMED CT knowledge bases. You can perform analysis both on single files or as a batch analysis on large documents or multiple files stored in an Amazon Simple Storage Service (S3).

ICD-10-CM linking

Use `InferICD10CM` to detect possible medical conditions as entities and link them to codes from the 2026 version of the [International Classification of Diseases, 10th Revision, Clinical Modification \(ICD-10-CM\)](#). The ICD-10-CM is provided by the US Centers for Disease Control and Prevention (CDC).

When medical conditions are detected, `InferICD10CM` returns the matching ICD-10-CM codes and descriptions. The detected conditions are listed in descending order of confidence. The scores indicate the confidence in the accuracy of the entities matched to the concepts found in the text. Related information such as family history, signs, symptoms, and negation are recognized as traits. Additional information such as anatomical designations and acuity are listed as attributes.

`InferICD10CM` is well-suited for the following scenarios:

- Assisting the professional medical coding of patient records
- Clinical studies and trials
- Integration with a medical software system
- Early detection and diagnosis
- Population health management

ICD-10-CM category

`InferICD10CM` detects entities in the `MEDICAL_CONDITION` category. Additional related information is also detected and linked as attributes or traits.

ICD-10-CM types

`InferICD10CM` detects entities of the types `DX_NAME` and `TIME_EXPRESSION`.

ICD-10-CM traits

InferICD10CM detects the following contextual information as traits:

- **DIAGNOSIS**: An identification of a medical condition that is determined by evaluation of the symptoms.
- **HYPOTHETICAL**: An indication that a medical condition is expressed as a hypothesis.
- **LOW_CONFIDENCE**: An indication that a medical condition is expressed as having high uncertainty. This is not directly related to the confidence scores provided.
- **NEGATION**: An indication that a medical condition is not present.
- **PERTAINS_TO_FAMILY**: An indication that a medical condition is relevant to the patient's family, not the patient.
- **SIGN**: A medical condition that is reported by the physician.
- **SYMPTOM**: A medical condition that is reported by the patient.

ICD-10-CM attributes

InferICD10CM detects the following contextual information as attributes:

- **DIRECTION**: Directional terms. For example, left, right, medial, lateral, upper, lower, posterior, anterior, distal, proximal, contralateral, bilateral, ipsilateral, dorsal, or ventral.
- **SYSTEM_ORGAN_SITE**: Anatomical location.
- **ACUITY**: Determination of a disease instance, such as chronic, acute, sudden, persistent, or gradual. This only applies to the **MEDICAL_CONDITION** type.
- **QUALITY**: Any descriptive term of the medical condition, such as stage or grade.

Time expression category

The **TIME_EXPRESSION** category detects entities related to time. This includes entities such as dates, and time expressions such as "three days ago," "today," "currently," "day of admission," "last month," or "16 days." Results in this category are only returned if they are associated with an entity. For example, the expression, "Yesterday, the patient was diagnosed with influenza" would return `Yesterday` as a **TIME_EXPRESSION** entity that overlaps with the **DX_NAME** entity, "influenza." However, "yesterday" would not be recognized as an entity in the expression, "yesterday, the patient walked their dog."

Types

The recognized type of `TIME_EXPRESSION` is `TIME_TO_DX_NAME`: the date that a medical condition occurred. The attribute for this type is `DX_NAME`.

Relationship type

The `RELATIONSHIP_TYPE` refers to the relationship between an entity and an attribute. The recognized `RELATIONSHIP_TYPE` is `OVERLAP`– the `TIME_EXPRESSION` concurs with the entity detected.

Input and response examples

Note

For specific API input and response syntax, see [InferICD10CM](#) in the *Amazon Comprehend Medical API Reference*.

The following example input text shows how the `InferICD10CM` operation works. To view all input text, scroll over the **Copy** button.

```
"The patient is a 71-year-old female patient of Dr. X. The patient presented to the emergency room last evening with approximately 7 to 8 day history of abdominal pain which has been persistent. She has had no nausea and vomiting, but has had persistent associated anorexia. She is passing flatus, but had some obstipation symptoms with the last bowel movement two days ago. She denies any bright red blood per rectum and no history of recent melena. Her last colonoscopy was approximately 5 years ago with Dr. Y. She has had no definite fevers or chills and no history of jaundice. The patient denies any significant recent weight loss."
```

The `InferICD10CM` operation returns the following output in JSON format (abbreviated for brevity).

```
{
  "Entities": [
    {
      "Id": 1,
      "Text": "abdominal pain",
      "Category": "MEDICAL_CONDITION",
      "Type": "DX_NAME",
```

```
"Score": Float,
"BeginOffset": 153,
"EndOffset": 167,
"Attributes": [
  {
    "Type": "ACUITY",
    "Score": Float,
    "RelationshipScore": Float,
    "Id": 2,
    "BeginOffset": 183,
    "EndOffset": 193,
    "Text": "persistent",
    "Traits": []
  }
],
"Traits": [
  {
    "Name": "SYMPTOM",
    "Score": Float
  }
],
"ICD10CMConcepts": [
  {
    "Description": "Unspecified abdominal pain",
    "Code": "R10.9",
    "Score": Float
  },
  {
    "Description": "Epigastric pain",
    "Code": "R10.13",
    "Score": Float
  },
  {
    "Description": "Lower abdominal pain, unspecified",
    "Code": "R10.30",
    "Score": Float
  },
  {
    "Description": "Generalized abdominal pain",
    "Code": "R10.84",
    "Score": Float
  },
  {
    "Description": "Upper abdominal pain, unspecified",
```

```

                "Code": "R10.10",
                "Score": Float
            }
        ]
    }
    ...
    "ModelVersion": "3.3.0.20251001"
}

```

InferICD10CM also recognizes when an entity is negated in text. For instance, if a patient is not experiencing a symptom, both the symptom and negation are identified as traits and listed with a confidence score. Based on the input for the previous example, the symptom Nausea will be listed under NEGATION because the patient isn't experiencing nausea.

```

{
  "Id": 3,
  "Text": "nausea",
  "Category": "MEDICAL_CONDITION",
  "Type": "DX_NAME",
  "Score": Float,
  "BeginOffset": 210,
  "EndOffset": 216,
  "Attributes": [],
  "Traits": [
    {
      "Name": "SYMPTOM",
      "Score": Float
    },
    {
      "Name": "NEGATION",
      "Score": Float
    }
  ],
  "ICD10CMConcepts": [
    {
      "Description": "Nausea with vomiting, unspecified",
      "Code": "R11.2",
      "Score": Float
    },
    {
      "Description": "Nausea",
      "Code": "R11.0",
      "Score": Float
    }
  ]
}

```

```
    }  
  ]  
}
```

RxNorm linking

Use the **InferRxNorm** operation to identify medications that are listed in a patient record as entities. The operation also links those entities to concept identifiers (RxCUI) from [the RxNorm database from the National Library of Medicine](#). The source for each RxCUI is the 2022-11-07 RxNorm and RxTerms Release. Each RxCUI is unique for different strengths and dose forms. Amazon Comprehend Medical lists the top potentially matching RxCUIs for each medication that it detects in descending order by confidence score. Use the RxCUI codes for downstream analysis that is not possible with unstructured text. Related information such as strength, frequency, dose, dose form, and route of administration are listed as attributes in JSON format.

You can use **InferRxNorm** for the following scenarios:

- Screening for medications the patient has taken.
- Preventing potentially negative reactions between newly prescribed drugs and drugs the patient is currently taking.
- Screening for inclusion in clinical trials based on drug history using the RxCUI.
- Checking whether the dosage and frequency of a drug is appropriate.
- Screening for uses, indications, and side effects of drugs.
- Managing population health.

Important notice

The **InferRxNorm** operation of Amazon Comprehend Medical is not a substitute for professional medical advice, diagnosis, or treatment. Identify the right confidence threshold for your use case, and use high confidence thresholds in situations that require high accuracy. Only use Amazon Comprehend Medical operations in patient care scenarios *after* reviewing for accuracy and receiving sound judgment by trained medical professionals.

RxNorm category

InferRxNorm detects entities in the MEDICATION category. It also detects additional related information that is linked as attributes or traits.

RxNorm types

The types of entities in the Medication category are

- **BRAND_NAME**: The copyrighted brand name of the medication or therapeutic agent.
- **GENERIC_NAME**: Non-brand name, ingredient name, or formula mixture of the medication or therapeutic agent.

RxNorm attributes

- **DOSAGE**: The amount of medication ordered.
- **DURATION**: How long the medication should be administered.
- **FORM**: The form of the medication.
- **FREQUENCY**: How often to administer the medication.
- **RATE**: The administration rate of the medication (primarily for medication infusions or IVs).
- **ROUTE_OR_MODE**: The administration method of a medication.
- **STRENGTH**: The medication strength.

RxNorm traits

- **NEGATION**: Any indication that the patient is *not* taking a medication.
- **PAST_HISTORY**: An indication that a medication detected is from the patient's past (prior to current encounter).

Input and response examples

Note

For specific API input and response syntax, see [InferRxNorm](#) in the *Amazon Comprehend Medical API Reference*.

The following example input text shows how the `InferRxNorm` operation works. To view all input text, scroll over the **Copy** button.

```
"fluoride topical ( fluoride 1.1 % topical gel ) 1 application Topically daily Brush
onto teeth before bed time , spit , do not rinse, eat or drink for 20-30 minutes"
```

The InferRxNorm operation returns the following output in JSON format:

```
{
  "Entities": [
    {
      "Id": 1,
      "Text": "fluoride",
      "Category": "MEDICATION",
      "Type": "GENERIC_NAME",
      "Score": Float,
      "BeginOffset": 19,
      "EndOffset": 27,
      "Attributes": [],
      "Traits": [],
      "RxNormConcepts": [
        {
          "Description": "fluorine",
          "Code": "1310123",
          "Score": Float
        },
        {
          "Description": "sodium fluoride",
          "Code": "9873",
          "Score": Float
        },
        {
          "Description": "magnesium fluoride",
          "Code": "1435860",
          "Score": Float
        },
        {
          "Description": "sulfuryl fluoride",
          "Code": "2289224",
          "Score": Float
        },
        {
          "Description": "acidulated phosphate fluoride",
          "Code": "236",
          "Score": Float
        }
      ]
    }
  ]
}
```

```
    ]
  }
],
"ModelVersion": "3.3.0.20221107"
}
```

Using the following input text, the InferRxNorm operation recognizes the negation trait, too.

```
'patient is not on warfarin'
```

The InferRxNorm operation returns the following output in JSON format:

```
{
  "Entities": [
    {
      "Id": 1,
      "Text": "warfarin",
      "Category": "MEDICATION",
      "Type": "GENERIC_NAME",
      "Score": Float,
      "BeginOffset": 18,
      "EndOffset": 26,
      "Attributes": [],
      "Traits": [
        {
          "Name": "NEGATION",
          "Score": Float
        }
      ],
      "RxNormConcepts": [
        {
          "Description": "warfarin",
          "Code": "11289",
          "Score": Float
        },
        {
          "Description": "warfarin sodium 2 MG Oral Tablet",
          "Code": "855302",
          "Score": Float
        },
        {
          "Description": "warfarin sodium 10 MG Oral Tablet",
          "Code": "855296",

```

```
        "Score": Float
    },
    {
        "Description": "warfarin sodium 2 MG Oral Tablet [Coumadin]",
        "Code": "855304",
        "Score": Float
    },
    {
        "Description": "warfarin sodium 10 MG Oral Tablet [Jantoven]",
        "Code": "855300",
        "Score": Float
    }
]
}
],
"ModelVersion": "3.3.0.20221107"
}
```

SNOMED CT linking

Use **InferSNOMEDCT** to detect medical entities and link them to concepts from the 2022-03 version of the Systematized Nomenclature of Medicine, Clinical Terms (SNOMED CT). SNOMED CT provides you with a comprehensive vocabulary of medical concepts, including medical conditions and anatomy, medical tests, treatments, and procedures. To learn more about SNOMED CT, visit [SNOMED CT](#).

For each detected medical entity, Amazon Comprehend Medical lists the top five SNOMED CT concept IDs and descriptions associated with the medical concept, along with a confidence score to indicate the confidence of the model in its prediction. The SNOMED CT concept IDs are listed in descending order of confidence along with the confidence scores. The SNOMED CT concept IDs can then be used to structure patient clinical data for medical coding, reporting, or clinical analytics when you use them with the SNOMED CT poly-hierarchy.

InferSNOMEDCT is available for customers in the US. For information on SNOMED CT in other countries, and information on SNOMED CT licensing, see [SNOMED CT](#).

InferSNOMEDCT is well suited for the following scenarios:

- Assistance for professional medical coding in patient records
- Clinical studies and trials

- Population health management

InferSNOMEDCT detects entities in the following categories. Additional contextual information is also detected and linked as attributes or traits.

- **MEDICAL_CONDITION**: The signs, symptoms, and diagnoses of medical conditions.
- **ANATOMY**: The parts of the body or body systems and the locations of those parts or systems.
- **TEST_TREATMENT_PROCEDURE**: The procedures that are used to determine a medical condition.

Anatomy category

The **ANATOMY** category detects references to the parts of the body or body systems and the locations of those parts or systems.

Attributes

The following attributes are detected for the **ANATOMY** category:

- **DIRECTION**: Directional terms. For example, left, right, medial, lateral, upper, lower, posterior, anterior, distal, proximal, contralateral, bilateral, ipsilateral, dorsal, or ventral.
- **SYSTEM_ORGAN_SITE**: Body systems, anatomic locations or regions, and body sites.

Medical condition category

The **MEDICAL_CONDITION** category detects the signs, symptoms, and diagnoses of medical conditions.

Type

For the **MEDICAL_CONDITION** category, the following type is detected:

- **DX_NAME**: An identification of a medical condition that is determined by evaluation of the symptoms.

Attributes

The following attributes are detected for the **MEDICAL_CONDITION** category:

- **ACUITY**: Determination of disease instance, such as chronic, acute, sudden, persistent, or gradual.
- **QUALITY**: Any descriptive term of the medical condition, such as stage or grade.
- **DIRECTION**: Directional terms. For example, left, right medial, lateral, upper, lower, posterior, anterior, distal, proximal, contralateral, bilateral, ipsilateral, dorsal, or ventral.
- **SYSTEM_ORGAN_SITE**: Body systems, anatomic locations or regions, and body sites.

Traits

The following traits are detected for the **MEDICAL_CONDITION** category:

- **DIAGNOSIS**: A medical condition that is determined as the cause or result of the symptoms. Symptoms can be found through physical findings, laboratory or radiological reports, or other means.
- **HYPOTHETICAL**: An indication that a medical condition is expressed as a hypothesis.
- **LOW_CONFIDENCE**: An indication that a medical condition is expressed as having high uncertainty. This is not directly related to the confidence scores provided.
- **NEGATION**: An indication that a medical condition is not present.
- **PERTAINS_TO_FAMILY**: An indication that a medical condition is relevant to the patient's family, not the patient.
- **SIGN**: A medical condition that is reported by the physician.
- **SYMPTON**: A medical condition that is reported by the patient.

Test, treatment, and procedure category

The **TEST_TREATMENT_PROCEDURE** category detects the procedures that are used to determine a medical condition.

Type

For the **TEST_TREATMENT_PROCEDURE** category, the following types are detected:

- **PROCEDURE_NAME**: Interventions performed on the patient to treat a medical condition or to provide patient care.

- **TEST_NAME**: Procedures performed on a patient for diagnosis, measurement, screening, or a rating that might have a resulting value. This includes any procedure, process, evaluation, or rating to determine a diagnosis, to rule out or find a condition, or to scale or score a patient.
- **TREATMENT_NAME**: Interventions performed to combat a disease or disorder. This includes medications, such as antivirals and vaccinations.

Attributes

For the **TEST_TREATMENT_PROCEDURE** category, the following attributes are detected:

- **TEST_NAME**: The diagnostic test performed.
- **TEST_VALUE**: The numeric results from a diagnostic test.
- **TEST_UNIT**: The units associated with a **TEST_VALUE**: result.
- **PROCEDURE_NAME**: The name of a surgery or medical procedure performed.
- **TREATMENT_NAME**: The name of a treatment administered to a patient.

Traits

- **FUTURE**: An indication that a test, treatment, or procedure refers to an action or event that will occur after the subject of the notes.
- **HYPOTHETICAL**: An indication that a test, treatment, or procedure is expressed as a hypothesis
- **NEGATION**: An indication that a result or action is negative or not being performed.
- **PAST_HISTORY**: An indication that a test, treatment, or procedure is from the patient's past (prior to the current encounter).

SNOMED CT details

Included in the JSON response are the SNOMED CT details, which includes the following information:

- **EDITION**: Only the US edition is supported.
- **VERSIONDATE**: The date stamp of the SNOMED CT version used.
- **LANGUAGE**: Analysis on English (US-EN) language is supported.

Input and response examples

Note

For specific API input and response syntax, see [InferSNOMEDCT](#) in the *Amazon Comprehend Medical API Reference*.

The following example input text shows how the InferSNOMEDCT operation works. To view all input text, scroll over the **Copy** button.

```
"HEENT : Boggy inferior turbinates, No oropharyngeal lesion"
```

The InferSNOMEDCT operation returns the following output in JSON format.

```
{
  "Entities": [
    {
      "Category": "ANATOMY",
      "BeginOffset": 0,
      "EndOffset": 5,
      "Text": "HEENT",
      "Traits": [],
      "SNOMEDCTConcepts": [
        {
          "Code": "69536005",
          "Score": Float,
          "Description": "Head structure (body structure)"
        },
        {
          "Code": "429031000124106",
          "Score": Float,
          "Description": "Review of systems, head, ear, eyes, nose and throat
(procedure)"
        },
        {
          "Code": "385383008",
          "Score": Float,
          "Description": "Ear, nose and throat structure (body structure)"
        },
        {
          "Code": "64237003",
```

```

        "Score": Float,
        "Description": "Structure of left half of head (body structure)"
    },
    {
        "Code": "113028003",
        "Score": Float,
        "Description": "Ear, nose and throat examination (procedure)"
    }
],
"Score": Float,
"Attributes": [],
"Type": "SYSTEM_ORGAN_SITE",
"Id": 0
},
{
    "Category": "MEDICAL_CONDITION",
    "BeginOffset": 8,
    "EndOffset": 33,
    "Text": "Boggy inferior turbinates",
    "Traits": [
        {
            "Score": Float,
            "Name": "SIGN"
        }
    ],
    "SNOMEDCTConcepts": [
        {
            "Code": "254477009",
            "Score": Float,
            "Description": "Tumor of inferior turbinate (disorder)"
        },
        {
            "Code": "260762006",
            "Score": Float,
            "Description": "Choroidal invasion status (attribute)"
        },
        {
            "Code": "2455009",
            "Score": Float,
            "Description": "Revision of lumbosubarachnoid shunt (procedure)"
        },
        {
            "Code": "19883003",
            "Score": Float,

```

```

        "Description": "Atrophy of nasal turbinates (disorder)"
    },
    {
        "Code": "256723009",
        "Score": Float,
        "Description": "Inferior turbinate flap (substance)"
    }
],
"Score": Float,
"Attributes": [
    {
        "Category": "ANATOMY",
        "RelationshipScore": Float,
        "EndOffset": 5,
        "Text": "HEENT",
        "Traits": [],
        "SNOMEDCTConcepts": [
            {
                "Code": "69536005",
                "Score": Float,
                "Description": "Head structure (body structure)"
            },
            {
                "Code": "429031000124106",
                "Score": Float,
                "Description": "Review of systems, head, ear, eyes, nose
and throat (procedure)"
            },
            {
                "Code": "385383008",
                "Score": Float,
                "Description": "Ear, nose and throat structure (body
structure)"
            },
            {
                "Code": "64237003",
                "Score": Float,
                "Description": "Structure of left half of head (body
structure)"
            },
            {
                "Code": "113028003",
                "Score": Float,

```

```

        "Description": "Ear, nose and throat examination
(procedure)"
    }
  ],
  "Score": Float,
  "RelationshipType": "SYSTEM_ORGAN_SITE",
  "Type": "SYSTEM_ORGAN_SITE",
  "Id": 0,
  "BeginOffset": 0
}
],
"Type": "DX_NAME",
"Id": 1
},
{
  "Category": "ANATOMY",
  "BeginOffset": 23,
  "EndOffset": 33,
  "Text": "turbinates",
  "Traits": [],
  "SNOMEDCTConcepts": [
    {
      "Code": "310607007",
      "Score": Float,
      "Description": "Sarcoidosis of inferior turbinates (disorder)"
    },
    {
      "Code": "80153006",
      "Score": Float,
      "Description": "Segmented neutrophil (cell)"
    },
    {
      "Code": "46607005",
      "Score": Float,
      "Description": "Nasal turbinate structure (body structure)"
    },
    {
      "Code": "6553002",
      "Score": Float,
      "Description": "Inferior nasal turbinate structure (body
structure)"
    },
    {
      "Code": "254477009",

```

```

        "Score": Float,
        "Description": "Tumor of inferior turbinate (disorder)"
    }
],
"Score": Float,
"Attributes": [],
"Type": "SYSTEM_ORGAN_SITE",
"Id": 3
},
{
    "Category": "ANATOMY",
    "BeginOffset": 39,
    "EndOffset": 52,
    "Text": "oropharyngeal",
    "Traits": [],
    "SNOMEDCTConcepts": [
        {
            "Code": "31389004",
            "Score": Float,
            "Description": "Oropharyngeal structure (body structure)"
        },
        {
            "Code": "33431000119109",
            "Score": Float,
            "Description": "Lesion of oropharynx (disorder)"
        },
        {
            "Code": "263376008",
            "Score": Float,
            "Description": "Entire oropharynx (body structure)"
        },
        {
            "Code": "716151000",
            "Score": Float,
            "Description": "Structure of oropharynx and/or hypopharynx and/or
larynx (body structure)"
        },
        {
            "Code": "764786007",
            "Score": Float,
            "Description": "Oropharyngeal (intended site)"
        }
    ],
    "Score": Float,

```

```
    "Attributes": [],
    "Type": "SYSTEM_ORGAN_SITE",
    "Id": 5
  },
  {
    "Category": "MEDICAL_CONDITION",
    "BeginOffset": 39,
    "EndOffset": 59,
    "Text": "oropharyngeal lesion",
    "Traits": [
      {
        "Score": Float,
        "Name": "SIGN"
      }
    ],
    "SNOMEDCTConcepts": [
      {
        "Code": "31389004",
        "Score": Float,
        "Description": "Oropharyngeal structure (body structure)"
      },
      {
        "Code": "33431000119109",
        "Score": Float,
        "Description": "Lesion of oropharynx (disorder)"
      },
      {
        "Code": "764786007",
        "Score": Float,
        "Description": "Oropharyngeal (intended site)"
      },
      {
        "Code": "418664002",
        "Score": Float,
        "Description": "Oropharyngeal route (qualifier value)"
      },
      {
        "Code": "110162001",
        "Score": Float,
        "Description": "Abrasion of oropharynx (disorder)"
      }
    ],
    "Score": Float,
    "Attributes": [
```

```

    {
      "Category": "ANATOMY",
      "RelationshipScore": Float,
      "EndOffset": 5,
      "Text": "HEENT",
      "Traits": [],
      "SNOMEDCTConcepts": [
        {
          "Code": "69536005",
          "Score": Float,
          "Description": "Head structure (body structure)"
        },
        {
          "Code": "429031000124106",
          "Score": Float,
          "Description": "Review of systems, head, ear, eyes, nose
and throat (procedure)"
        },
        {
          "Code": "385383008",
          "Score": Float,
          "Description": "Ear, nose and throat structure (body
structure)"
        },
        {
          "Code": "64237003",
          "Score": Float,
          "Description": "Structure of left half of head (body
structure)"
        },
        {
          "Code": "113028003",
          "Score": Float,
          "Description": "Ear, nose and throat examination
(procedure)"
        }
      ],
      "Score": Float,
      "RelationshipType": "SYSTEM_ORGAN_SITE",
      "Type": "SYSTEM_ORGAN_SITE",
      "Id": 0,
      "BeginOffset": 0
    }
  ],

```

```
        "Type": "DX_NAME",
        "Id": 4
    }
],
"SNOMEDCTDetails": {
    "Edition": "US",
    "VersionDate": "20200901",
    "Language": "en"
},
"Characters": {
    "OriginalTextCharacters": 59
},
"ModelVersion": "3.3.0.20220301"
}
```

Ontology linking batch analysis

Use Amazon Comprehend Medical to detect entities in clinical text stored in an Amazon Simple Storage Service (Amazon S3) bucket and to link those entities to standardized ontologies. You can use ontology linking batch analysis to analyze either a collection of documents or a single document with up to 20,000 characters. By using either the console or the ontology linking batch API operations, you can perform operations to start, stop, list, and describe ongoing batch analysis jobs.

For pricing information for batch analysis and other Amazon Comprehend Medical operations, see [Amazon Comprehend Medical Pricing](#).

Performing batch analysis

You can run a batch analysis job using either the Amazon Comprehend Medical console or the Amazon Comprehend Medical batch API operations.

Performing batch analysis using the API operations

Prerequisites

When you are using the Amazon Comprehend Medical API, create an AWS Identity Access and Management (IAM) policy and attach it to an IAM role. To learn more about IAM roles and trust policies, see [IAM Policies and Permissions](#).

1. Upload your data into an S3 bucket.

2. To start a new analysis job, use the **StartICD10CMInferenceJob**, **StartSNOMEDCTInferenceJob**, or the **StartRxNormInferenceJob** operations. Provide the name of the Amazon S3 bucket that contains the input files and the name of the Amazon S3 bucket where you want to send the output files.
3. Monitor the progress of the job by using **DescribeICD10CMInferenceJob**, **DescribeSNOMEDCTInferenceJob**, or **DescribeRxNormInferenceJob** operations. Additionally, you can use **ListICD10CMInferenceJobs**, **ListSNOMEDCTInferenceJobs**, and **ListRxNormInferenceJobs** to see the status of all ontology linking batch analysis jobs.
4. If you need to stop a job in progress, use **StopICD10CMInferenceJob**, **StopSNOMEDCTInferenceJob**, or **StopRxNormInferenceJob** to stop analysis.
5. To view the results of your analysis job, see the output S3 bucket that you configured when you started the job.

Performing batch analysis using the console

1. Upload your data into an S3 bucket.
2. To start a new analysis job, select the type of analysis you will be performing. Then, provide the name of the S3 bucket that contains the input files and the name of the S3 bucket where you want to send the output files.
3. Monitor the status of your job while it is ongoing. From the console, you can view all batch analysis operations and their status, including when analysis was started and ended.
4. To see the results of your analysis job, see the output S3 bucket that you configured when you started the job.

IAM policies for batch operations

The IAM role that calls the Amazon Comprehend Medical batch API operations must have a policy that grants access to the S3 buckets that contain the input and output files. The IAM role must also be assigned a trust relationship so that the Amazon Comprehend Medical service can assume the role. To learn more about IAM roles and trust policies, see [IAM Roles](#).

The role must have the following policy:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::input-bucket/*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::input-bucket",
        "arn:aws:s3:::output-bucket"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "s3:PutObject"
      ],
      "Resource": [
        "arn:aws:s3:::output-bucket/*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

The role must have the following trust relationship. It is recommended that you use the `aws:SourceAccount` and `aws:SourceArn` condition keys to prevent the confused deputy security issue. To learn more about the confused deputy problem and how to protect your AWS account, see [The confused deputy problem](#) in the IAM documentation.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "comprehendmedical.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "account_id"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:comprehendmedical:us-east-1:account_id:*"
        }
      }
    }
  ]
}
```

Batch analysis output files

Amazon Comprehend Medical creates one output file for each input file in the batch. The file has the extension `.out`. Amazon Comprehend Medical first creates a directory in the output S3 bucket using the *`AwsAccountId-JobType-JobId`* as the name, and then it writes all of the output files for the batch to this directory. Amazon Comprehend Medical creates this new directory so that output from one job doesn't overwrite the output of another job.

A batch operation produces the same output as a synchronous operation.

Each batch operation produces the following three manifest files that contain information about the job:

- **Manifest** – Summarizes the job. Provides information about the parameters used for the job, the total size of the job, and the number of files processed.
- **Success** – Provides information about the files that were successfully processed. Includes the input and output file name and the size of the input file.
- **Unprocessed** – Lists files that the batch job did not process with error codes and error messages per file.

Amazon Comprehend Medical writes the files to the output directory that you specified for the batch job. The summary manifest file will be written to the output folder, along with a folder titled `Manifest_AccountId-Operation-JobId`. Within the manifest folder is the success folder, which contains the success manifest, and the failed folder, which contains the unprocessed file manifest. The following sections show the structure of the manifest files.

Batch manifest file

The following is the JSON structure of the batch manifest file.

```
{
  "Summary" :
  {
    "Status" : "COMPLETED | FAILED | PARTIAL_SUCCESS | STOPPED",
    "JobType" : "ICD10CMInference | RxNormInference | SNOMEDCTInference",
    "InputDataConfiguration" : {
      "Bucket" : "input bucket",
      "Path" : "path to files/account ID-job type-job ID"
    },
    "OutputDataConfiguration" : {
      "Bucket" : "output bucket",
      "Path" : "path to files"
    },
    "InputFileCount" : number of files in input bucket,
    "TotalMeteredCharacters" : total characters processed from all files,
    "UnprocessedFilesCount" : number of files not processed,
    "SuccessFilesCount" : total number of files processed,
    "TotalDurationSeconds" : time required for processing,
    "SuccessfulFilesListLocation" : "path to file",
    "UnprocessedFilesListLocation" : "path to file",
    "FailedJobErrorMessage": "error message or if not applicable,
      The status of the job is completed"
  }
}
```

Success manifest file

The following is the JSON structure of the file that contains information about successfully processed files.

```
{
  "Files": [{
    "Input": "input path/input file name",
    "Output": "output path/output file name",
    "InputSize": size in bytes of input file
  },
  {
    "Input": "input path/input file name",
    "Output": "output path/output file name",
    "InputSize": size in bytes of input file
  }
]
```

Unprocessed manifest file

Following is the JSON structure of the manifest file that contains information about unprocessed files.

```
{
  "Files" : [ {
    "Input": "file_name_that_failed",
    "ErrorCode": "error code for exception",
    "ErrorMessage": "explanation of the error code and suggestions"
  },
  { ...}
]
```

Security in Amazon Comprehend Medical

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Comprehend Medical, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Comprehend Medical. The following topics show you how to configure Amazon Comprehend Medical to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon Comprehend Medical resources.

Topics

- [Data protection in Amazon Comprehend Medical](#)
- [Identity and access management in Amazon Comprehend Medical](#)
- [Logging Amazon Comprehend Medical API calls using AWS CloudTrail](#)
- [Compliance validation for Amazon Comprehend Medical](#)
- [Resilience in Amazon Comprehend Medical](#)
- [Infrastructure security in Amazon Comprehend Medical](#)

Data protection in Amazon Comprehend Medical

The AWS [shared responsibility model](#) applies to data protection in Amazon Comprehend Medical. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Comprehend Medical or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Identity and access management in Amazon Comprehend Medical

Access to Comprehend Medical requires credentials that AWS can use to authenticate your requests. Those credentials must have permissions to access Comprehend Medical actions. [AWS Identity and Access Management \(IAM\)](#) can help secure your resources by controlling who can access them. The following sections provide details on how you can use IAM with Comprehend Medical.

- [Authentication](#)
- [Access Control](#)

Authentication

You must give users permissions to interact with Amazon Comprehend Medical. For users who need full access use `ComprehendMedicalFullAccess`.

To provide access, add permissions to your users, groups, or roles:

- Users and groups in AWS IAM Identity Center:

Create a permission set. Follow the instructions in [Create a permission set](#) in the *AWS IAM Identity Center User Guide*.

- Users managed in IAM through an identity provider:

Create a role for identity federation. Follow the instructions in [Create a role for a third-party identity provider \(federation\)](#) in the *IAM User Guide*.

- IAM users:

- Create a role that your user can assume. Follow the instructions in [Create a role for an IAM user](#) in the *IAM User Guide*.
- (Not recommended) Attach a policy directly to a user or add a user to a user group. Follow the instructions in [Adding permissions to a user \(console\)](#) in the *IAM User Guide*.

To use Amazon Comprehend Medical's asynchronous operations you also need a service role.

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

To learn more about specifying Amazon Comprehend Medical as the service in principal, see [Role-based Permissions required for batch operations](#).

Access Control

You must have valid credentials to authenticate your requests. The credentials must have permissions to call an Amazon Comprehend Medical action.

The following sections describe how to manage permissions for Amazon Comprehend Medical. We recommend that you read the overview first.

- [Overview of managing access permissions to Amazon Comprehend Medical resources](#)
- [Using Identity-Based policies \(IAM policies\) for Amazon Comprehend Medical](#)

Topics

- [Overview of managing access permissions to Amazon Comprehend Medical resources](#)
- [Using Identity-Based policies \(IAM policies\) for Amazon Comprehend Medical](#)
- [Amazon Comprehend Medical API Permissions: actions, resources, and conditions reference](#)
- [AWS managed policies for Amazon Comprehend Medical](#)

Overview of managing access permissions to Amazon Comprehend Medical resources

Permissions policies govern the access to an action. An account administrator attaches permissions policies to IAM identities to manage access to actions. IAM identities include users, groups, and roles.

Note

An *account administrator* (or administrator user) is a user with administrator privileges. For more information, see [IAM Best Practices](#) in the *IAM User Guide*.

When you grant permissions, you decide both who and what actions get the permissions.

Topics

- [Managing access to actions](#)
- [Specifying policy elements: actions, effects, and principals](#)
- [Specifying conditions in a policy](#)

Managing access to actions

A *permissions policy* describes who has access to what. The following section explains the options for permissions policies.

Note

This section explains IAM in the context of Amazon Comprehend Medical. It doesn't provide detailed information about the IAM service. For more about IAM, see [What Is IAM?](#) in the *IAM User Guide*. For information about IAM policy syntax and descriptions, see [AWS IAM Policy Reference](#) in the *IAM User Guide*.

Policies attached to an IAM identity are *identity-based* policies. Policies attached to a resource are *resource-based* policies. Amazon Comprehend Medical supports only identity-based policies.

Identity-based policies (IAM policies)

You can attach policies to IAM identities. Here are two examples.

- **Attach a permissions policy to a user or a group in your account.** To allow a user or a group of users to call an Amazon Comprehend Medical action, attach a permissions policy to a user. Attach a policy to a group that contains the user.
- **Attach a permissions policy to a role to grant cross-account permissions.** To grant cross-account permissions, attach an identity-based policy to an IAM role. For example, the administrator in Account A can create a role to grant cross-account permissions to another account. In this example, call it Account B, which could also be an AWS service.
 1. Account A administrator creates an IAM role and attaches a policy to the role that grants permissions to resources in Account A.

2. Account A administrator attaches a trust policy to the role. The policy identifies Account B as the principal who can assume the role.
3. Account B administrator can then delegate permissions to assume the role to any users in Account B. This allows users in Account B to create or access resources in Account A. If you want to grant an AWS service the permissions to assume the role, the principal in the trust policy can also be an AWS service principal.

For more information about using IAM to delegate permissions, see [Access Management](#) in the *IAM User Guide*.

For more information about using identity-based policies with Amazon Comprehend Medical, see [Using Identity-Based policies \(IAM policies\) for Amazon Comprehend Medical](#). For more information about users, groups, roles, and permissions, see [Identities \(Users, Groups, and Roles\)](#) in the *IAM User Guide*.

Resource-based policies

Other services, such as AWS Lambda, support resource-based permissions policies. For example, you can attach a policy to an S3 bucket to manage access permissions to that bucket. Amazon Comprehend Medical doesn't support resource-based policies.

Specifying policy elements: actions, effects, and principals

Amazon Comprehend Medical defines a set of API operations. To grant permissions for these API operations, Amazon Comprehend Medical defines a set of actions that you can specify in a policy.

The four items here are the most basic policy elements.

- **Resource** – In a policy, use an Amazon Resource Name (ARN) to identify the resource to which the policy applies. For Amazon Comprehend Medical, the resource is always "*".
- **Action** – Use action keywords to identify operations that you want to allow or deny. For example, depending on the specified effect, `comprehendmedical:DetectEntities` either allows or denies the user permission to perform the Amazon Comprehend Medical `DetectEntities` operation.
- **Effect** – Specify the effect of the action that occurs when the user requests the specific action—either allow or deny. If you don't explicitly grant access to (allow) a resource, access is implicitly denied. You can also explicitly deny access to a resource. You might do this to make sure that a user cannot access the resource, even if a different policy grants access.

- **Principal** – In identity-based policies, the user that the policy is attached to is the implicit principal.

To learn more about IAM policy syntax and descriptions, see [AWS IAM Policy Reference](#) in the *IAM User Guide*.

For a table showing all of the Amazon Comprehend Medical API actions, see [Amazon Comprehend Medical API Permissions: actions, resources, and conditions reference](#).

Specifying conditions in a policy

When you grant permissions, you use the IAM policy language to specify the conditions under which a policy should take effect. For example, you might want a policy to be applied only after a specific date. For more information about specifying conditions in a policy language, see [Condition](#) in the *IAM User Guide*.

AWS provides a set of predefined condition keys for all AWS services that support IAM for access control. For example, you can use the `aws:user:uid` condition key to require a specific AWS ID when requesting an action. For more information and a complete list of AWS keys, see [Available Keys for Conditions](#) in the *IAM User Guide*.

Amazon Comprehend Medical does not provide any additional condition keys.

Using Identity-Based policies (IAM policies) for Amazon Comprehend Medical

This topic shows example identity-based policies. The examples show how an account administrator can attach permissions policies to IAM identities. This enables users, groups, and roles to perform Amazon Comprehend Medical actions.

Important

To understand permissions, we recommend [Overview of managing access permissions to Amazon Comprehend Medical resources](#).

This example policy is required to use the Amazon Comprehend Medical document analysis actions.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowDetectActions",
      "Effect": "Allow",
      "Action": [
        "comprehendmedical:DetectEntitiesV2",
        "comprehendmedical:DetectPHI",
        "comprehendmedical:StartEntitiesDetectionV2Job",
        "comprehendmedical:ListEntitiesDetectionV2Jobs",
        "comprehendmedical:DescribeEntitiesDetectionV2Job",
        "comprehendmedical:StopEntitiesDetectionV2Job",
        "comprehendmedical:StartPHIDetectionJob",
        "comprehendmedical:ListPHIDetectionJobs",
        "comprehendmedical:DescribePHIDetectionJob",
        "comprehendmedical:StopPHIDetectionJob",
        "comprehendmedical:StartRxNormInferenceJob",
        "comprehendmedical:ListRxNormInferenceJobs",
        "comprehendmedical:DescribeRxNormInferenceJob",
        "comprehendmedical:StopRxNormInferenceJob",
        "comprehendmedical:StartICD10CMInferenceJob",
        "comprehendmedical:ListICD10CMInferenceJobs",
        "comprehendmedical:DescribeICD10CMInferenceJob",
        "comprehendmedical:StopICD10CMInferenceJob",
        "comprehendmedical:StartSNOMEDCTInferenceJob",
        "comprehendmedical:ListSNOMEDCTInferenceJobs",
        "comprehendmedical:DescribeSNOMEDCTInferenceJob",
        "comprehendmedical:StopSNOMEDCTInferenceJob",
        "comprehendmedical:InferRxNorm",
        "comprehendmedical:InferICD10CM",
        "comprehendmedical:InferSNOMEDCT"
      ],
      "Resource": "*"
    }
  ]
}

```

The policy has one statement that grants permission to use the `DetectEntities` and `DetectPHI` actions.

The policy doesn't specify the `Principal` element because you don't specify the principal who gets the permission in an identity-based policy. When you attach a policy to a user, the user is the implicit principal. When you attach a policy to an IAM role, the principal identified in the role's trust policy gets the permission.

To see all the Amazon Comprehend Medical API actions and the resources that they apply to, see [Amazon Comprehend Medical API Permissions: actions, resources, and conditions reference](#).

Permissions required to use the Amazon Comprehend Medical console

The permissions reference table lists the Amazon Comprehend Medical API operations and shows the required permissions for each operation. For more information, about Amazon Comprehend Medical API permissions, see [Amazon Comprehend Medical API Permissions: actions, resources, and conditions reference](#).

To use the Amazon Comprehend Medical console, grant permissions for the actions shown in the following policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:CreateRole",
        "iam:CreatePolicy",
        "iam:AttachRolePolicy"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "*",
      "Condition": {
```

```
        "StringEquals": {
            "iam:PassedToService": "comprehendmedical.amazonaws.com"
        }
    }
}
]
```

The Amazon Comprehend Medical console needs these permissions for the following reasons:

- iam permissions to list the available IAM roles for your account.
- s3 permissions to access the Amazon S3 buckets and objects that contain the data.

When you create an asynchronous batch job using the console, you can also create an IAM role for your job. To create an IAM role using the console, users must be granted the additional permissions shown here to create IAM roles and policies, and to attach policies to roles.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "iam:CreateRole",
        "iam:CreatePolicy",
        "iam:AttachRolePolicy"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

The Amazon Comprehend Medical console needs these permissions to create roles and policies and to attach roles and policies. The `iam:PassRole` action enables the console to pass the role to Amazon Comprehend Medical.

AWS managed (predefined) policies for Amazon Comprehend Medical

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. These AWS managed policies grant necessary permissions for common use cases so that you can avoid having to investigate what permissions are needed. For more information, see [AWS Managed Policies](#) in the *IAM User Guide*.

The following AWS managed policy, which you can attach to users in your account, is specific to Amazon Comprehend Medical.

- **ComprehendMedicalFullAccess** – Grants full access to Amazon Comprehend Medical resources. Includes permission to list and get IAM roles.

You must apply the following additional policy to any user using Amazon Comprehend Medical:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "iam:PassedToService": "comprehendmedical.amazonaws.com"
        }
      }
    }
  ]
}
```

You can review the managed permissions policies by signing in to the IAM console and searching for specific policies there.

These policies work when you are using AWS SDKs or the AWS CLI.

You can also create your own IAM policies to allow permissions for Amazon Comprehend Medical actions and resources. You can attach these custom policies to the IAM users or groups that require them.

Role-based Permissions required for batch operations

To use the Amazon Comprehend Medical asynchronous operations, grant Amazon Comprehend Medical access to the Amazon S3 bucket that contains your document collection. Do this by creating a data access role in your account to trust the Amazon Comprehend Medical service principal. For more information about creating a role, see [Creating a Role to Delegate Permissions to an AWS Service](#) in the *AWS Identity and Access Management User Guide*.

The following is the role's trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "comprehendmedical.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

After you have created the role, create an access policy for it. The policy should grant the Amazon S3 `GetObject` and `ListBucket` permissions to the Amazon S3 bucket that contains your input data. It also grants permissions for the Amazon S3 `PutObject` to your Amazon S3 output data bucket.

The following example access policy contains those permissions.


JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::input bucket/*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::input bucket"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "s3:PutObject"
      ],
      "Resource": [
        "arn:aws:s3:::output bucket/*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

Customer managed policy examples

In this section, you can find example user policies that grant permissions for various Amazon Comprehend Medical actions. These policies work when you are using AWS SDKs or the AWS CLI. When you are using the console, you must grant permissions to all the Amazon Comprehend

Medical APIs. This is discussed in [Permissions required to use the Amazon Comprehend Medical console](#).

 **Note**

All examples use the us-east-2 Region and contain fictitious account IDs.

Examples

Example 1: Allow all Amazon Comprehend Medical actions

After you sign up for AWS, you create an administrator to manage your account, including creating users and managing their permissions.

You can choose to create a user who has permissions for all Amazon Comprehend actions. Think of this user as a service-specific administrator for working with Amazon Comprehend. You can attach the following permissions policy to this user.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowAllComprehendMedicalActions",
    "Effect": "Allow",
    "Action": [
      "comprehendmedical:*"
    ],
    "Resource": "*"
  }
]
```

Example 2: Allow only DetectEntities actions

The following permissions policy grants user permissions to detect entities in Amazon Comprehend Medical, but not to detect PHI operations.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowDetectEntityActions",
      "Effect": "Allow",
      "Action": [
        "comprehendmedical:DetectEntitiesV2"
      ],
      "Resource": "*"
    }
  ]
}
```

Amazon Comprehend Medical API Permissions: actions, resources, and conditions reference

Use the following table as a reference when setting up [Access Control](#) and writing a permissions' policy that you can attach to a user. The list includes each Amazon Comprehend Medical API operation, the corresponding action for which you can grant permissions to perform the action, and the AWS resource for which you can grant the permissions. You specify the actions in the policy's `Action` field, and you specify the resource value in the policy's `Resource` field.

To express conditions, you can use AWS condition keys in your Amazon Comprehend Medical policies. For a complete list of keys, see [Available Keys](#) in the *IAM User Guide*.

Note

To specify an action, use the `comprehendmedical:` prefix followed by the API operation name, for example, `comprehendmedical:DetectEntities`.

AWS managed policies for Amazon Comprehend Medical

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

Topics

- [AWS managed policy: ComprehendMedicalFullAccess](#)
- [Amazon Comprehend Medical updates to AWS managed policies](#)

AWS managed policy: ComprehendMedicalFullAccess

You can attach the `ComprehendMedicalFullAccess` policy to your IAM identities.

This policy grants administrative permission to all Amazon Comprehend Medical actions.

JSON

```
{
  "Version": "2012-10-17",
  "Statement" : [
    {
```

```
"Action" : [
  "comprehendmedical:*"
],
"Effect" : "Allow",
"Resource" : "*"
}
]
}
```

Amazon Comprehend Medical updates to AWS managed policies

View details about updates to AWS managed policies for Amazon Comprehend Medical since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Document history](#) page.

Change	Description	Date
Amazon Comprehend Medical started tracking changes	Amazon Comprehend Medical started tracking changes for its AWS managed policies.	November 27, 2018

Logging Amazon Comprehend Medical API calls using AWS CloudTrail

Amazon Comprehend Medical is integrated with AWS CloudTrail. CloudTrail is a service that provides a record of actions taken by a user, role, or an AWS service from within Amazon Comprehend Medical. CloudTrail captures all API calls for Amazon Comprehend Medical as events. The calls captured include calls from the Amazon Comprehend Medical console and code calls to the Amazon Comprehend Medical API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Comprehend Medical. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine several things such as:

- The request that was made to Amazon Comprehend Medical
- The IP address from which the request was made
- Who made the request
- When the request was made
- Other details

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Amazon Comprehend Medical information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Amazon Comprehend Medical, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for Amazon Comprehend Medical, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

All Amazon Comprehend Medical actions are logged by CloudTrail and are documented in the [Amazon Comprehend Medical API Reference](#). For example, calls to the `DetectEntitiesV2`, `DetectPHI` and `ListEntitiesDetectionV2Jobs` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity Element](#).

Understanding Amazon Comprehend Medical log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source. The event includes information about the requested action, such as the date and time or request parameters. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the DetectEntitiesV2 action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/Mateo_Jackson",
    "accountId": "123456789012",
    "accessKeyId": "ASIAXHKUFODNN8EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::123456789012:user/Mateo_Jackson",
        "accountId": "123456789012",
        "userName": "Mateo_Jackson"
      },
      "webIdFederationData": {},
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2019-09-27T20:07:27Z"
      }
    }
  }
}
```

```
    },
    "eventTime": "2019-09-27T20:10:26Z",
    "eventSource": "comprehendmedical.amazonaws.com",
    "eventName": "DetectEntitiesV2",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "702.21.198.166",
    "userAgent": "aws-internal/3 aws-sdk-java/1.11.590
Linux/4.9.184-0.1.ac.235.83.329.metal1.x86_64 OpenJDK_64-Bit_Server_VM/25.212-b03
java/1.8.0_212 vendor/Oracle_Corporation",
    "requestParameters": null,
    "responseElements": null,
    "requestID": "8d85f2ec-EXAMPLE",
    "eventID": "ae9be9b1-EXAMPLE",
    "eventType": "AwsApiCall",
    "recipientAccountId": "123456789012"
}
```

Compliance validation for Amazon Comprehend Medical

Third-party auditors assess the security and compliance of Amazon Comprehend Medical as part of multiple AWS compliance programs. These include PCI, FedRAMP, HIPAA, and others. You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using Amazon Comprehend Medical is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [Architecting for HIPAA Security and Compliance Whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [AWS Config](#) – This AWS service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.

- [AWS Security Hub CSPM](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

For a list of AWS services in scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

Resilience in Amazon Comprehend Medical

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

Infrastructure security in Amazon Comprehend Medical

As a managed service, Amazon Comprehend Medical is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

To access Amazon Comprehend Medical through the network, you use AWS published API calls. Clients must support Transport Layer Security (TLS) 1.0 or later. We recommend TLS 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS), such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems, such as Java 7 and later, support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an AWS Identity and Access Management (IAM) principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Guidelines and quotas

Keep the following information in mind when using Amazon Comprehend Medical.

Note

- Amazon Comprehend Medical supports character encoding in UTF-8 English (EN).
- Amazon Comprehend Medical does not allow sequential forward slash characters (//) in file paths for async jobs.

Important notice

Amazon Comprehend Medical is not a substitute for professional medical advice, diagnosis, or treatment. Amazon Comprehend Medical provides confidence scores that indicate the level of confidence in the accuracy of the detected entities. Identify the right confidence threshold for your use case, and use high confidence thresholds in situations that require high accuracy. For certain use cases, results should be reviewed and verified by appropriately trained human reviewers. Use Amazon Comprehend Medical in patient care scenarios only after trained medical professionals have reviewed results for accuracy and sound medical judgment.

Supported regions

For a list of AWS Regions where Amazon Comprehend Medical is available, see [AWS Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

Throttling

For information about throttling and quotas for Amazon Comprehend Medical, and to request a quota increase, see [AWS Service Quotas](#).

Overall quotas

Amazon Comprehend Medical real time (sync) analysis operations have the following quotas:

Resource	Default
Transcriptions per second (TPS) for the DetectIntent v2, DetectIntent, and DetectEntities	40 TPS
Transcriptions per second (TPS) for the InferText EDCT operations	2 TPS
Characters	40,000 CPS

Resource	Default
per second (CPS) for the Detective Intelligence v2, Detective Intelligence Detective Intelligence, Inference, and Inference OCM operations	
Characters per second (CPS) for the Inference EDCT operations	5,000 CPS

Resource	Default
Maximum document size (UTF-characters) for the Detentite Detentite v2, and Detentite operations	20 KB

Resource	Default
Maximum document size (UTF-characters) for the Inference-0-CM and Inference operations	10 KB
Maximum document size (UTF-characters) for the Inference-EDCT operations	5 KB

Amazon Comprehend Medical batch analysis (async) operations have the following quotas:

Description	Quota
Transactions per second (TPS) for the <code>StartEntitiesDetectionV2Job</code> , <code>StartPHIDetectionJob</code> , <code>StopEntitiesDetectionV2Job</code> , <code>StopPHIDetectionJob</code> , <code>StartICD10CMInferenceJob</code> , <code>StartRxNormInferenceJob</code> , <code>StopICD10CMInferenceJob</code> , <code>StopRxNormInferenceJob</code> , <code>StartSNOMEDCTInferenceJob</code> and <code>StopSNOMEDCTInferenceJob</code> operations	5 TPS
Transactions per second (TPS) for the <code>ListEntitiesDetectionV2Jobs</code> , <code>ListPHIDetectionJobs</code> , <code>DescribeEntitiesDetectionV2Job</code> , <code>DescribePHIDetectionJob</code> , <code>ListICD10CMInferenceJobs</code> , <code>ListRxNormInferenceJobs</code> , <code>DescribeICD10CMInferenceJob</code> , <code>DescribeRxNormInferenceJob</code> , <code>ListSNOMEDCTInferenceJobs</code> and <code>DescribeSNOMEDCTInferenceJob</code> operations	10 TPS
Maximum individual file size for batch jobs for all operations	70 KB
Maximum number of individual files in a batch job	5000000
Maximum size of batch jobs (sum total of all files submitted in a batch job)	1 GB
Maximum number of active running batch jobs for each operation	10 jobs

If your text is larger than the character quotas, use [segment.py](#) to create smaller segments that can be analyzed.

Document history for Amazon Comprehend Medical

The following table describes the documentation for this release of Amazon Comprehend Medical.

Change	Description	Date
Updated API version for InferSNOMEDCT	The InferSNOMEDCT API operation now uses version 3.3.0.20220301 .	March 19, 2026
Updated API version for InferRxNorm	The InferRxNorm API operation now uses version 3.3.0.20221107 .	March 19, 2026
Updated API version for InferICD10CM	The InferICD10CM API operation now uses version 3.3.0.20251001 .	March 19, 2026
Updated API version for DetectEntitiesV2	The DetectEntitiesV2 API operation now uses version 3.3.0.	March 19, 2026
Updated API version for DetectEntities	The DetectEntities API operation now uses version 3.3.0.	March 19, 2026
Updated API version for InferSNOMEDCT	The InferSNOMEDCT API operation now uses version 3.2.0.20220301 .	November 3, 2025
Updated API version for InferRxNorm	The InferRxNorm API operation now uses version 3.2.0.20221107 .	November 3, 2025
Updated API version for InferICD10CM	The InferICD10CM API operation now uses version 3.2.0.20251001 .	November 3, 2025

Updated API version for InferSNOMEDCT	The InferSNOMEDCT API operation now uses version 3.1.0.20220301 .	August 5, 2025
Updated API version for InferRxNorm	The InferRxNorm API operation now uses version 3.1.0.20221107 .	August 5, 2025
Updated API version for InferICD10CM	The InferICD10CM API operation now uses version 3.1.0.20250401 .	August 5, 2025
Updated API version for InferSNOMEDCT	The InferSNOMEDCT API operation now uses version 3.0.0.20220301 .	June 26, 2024
Updated API version for InferRxNorm	The InferRxNorm API operation now uses version 3.0.0.20221107 .	June 26, 2024
Updated API version for InferICD10CM	The InferICD10CM API operation now uses version 3.0.0.20231001 .	June 26, 2024
Updated API version for DetectEntitiesV2	The DetectEntitiesV2 API operation now uses version 3.0.0.	June 26, 2024
Updated API version for DetectEntities	The DetectEntities API operation now uses version 3.0.0.	June 26, 2024
Updated API version for InferICD10CM	The API operation InferICD10CM now uses version 2.7.0.20231001 .	February 7, 2024
Updated API version for InferICD10CM	The API operation InferICD10CM now uses version 2.6.0.20220401 .	September 27, 2023

Updated API version for InferSNOMEDCT	The API operation <code>InferSNOMEDCT</code> now uses version <code>3.3.0.20220301</code> .	June 12, 2023
New trait and attribute added for DetectEntitiesV2	The API operation <code>DetectEntitiesV2</code> now uses version <code>3.3.0</code> . This update adds new traits and attributes.	June 12, 2023
New trait and attribute added for DetectEntities	The API operation <code>DetectEntities</code> now uses version <code>3.3.0</code> . This update adds new traits and attributes.	June 12, 2023
New trait added to InferRxNorm	The API operation <code>InferRxNorm</code> now uses version <code>3.3.0.20221107</code> . This update adds the <code>PAST_HISTORY</code> trait to medications detected.	June 12, 2023
New attribute added to InferICD10CM	The API operation <code>InferICD10CM</code> now uses version <code>2.5.0.20220401</code> . This update adds the <code>QUALITY</code> attribute to the medical condition category.	June 12, 2023
Updated API version for InferRxNorm	The API operation <code>InferRxNorm</code> now uses version <code>2.1.0.20221003</code> . This update improves the performance of the <code>InferRxNorm</code> model.	April 28, 2023

[Updated API version for InferSNOMEDCT](#)

The API operation `InferSNOMEDCT` now uses `2.5.0.20220301` . This update improves the model's performance.

April 21, 2023

[Updated API version for InferICD10CM](#)

The API operation `InferICD10CM` now uses version `2.4.0.20220401` . This update improves the detection of the `DX_NAME` entity and `DIRECTION` attributes.

March 17, 2023

[Updated API version for InferSNOMEDCT](#)

The API operation `InferSNOMEDCT` now uses version `2.4.0.20220301` . This update improves the detection of cardiology concepts in the `MEDICAL_CONDITION` and `TEST_TREATMENT_PROCEDURE` categories.

March 15, 2023

[Updated API version for InferICD10CM](#)

The API operation `InferICD10CM` now uses version `2.3.0.20220401` . This update improves the detection of cardiology concepts in the `MEDICAL_CONDITION` category.

March 15, 2023

[Updated API version for DetectEntitiesV2](#)

The API operation `DetectEntitiesV2` now uses version `2.3.0`. This update improves the detection of cardiology concepts in the `MEDICAL_CONDITION` and `TEST_TREATMENT_PROCEDURE` categories. .

March 15, 2023

[Updated API version for DetectEntities](#)

The API operation `DetectEntities` now uses version `0.6.0`. This update improves the detection of cardiology concepts in the `MEDICAL_CONDITION` and `TEST_TREATMENT_PROCEDURE` categories.

March 15, 2023

[Update API version for InferSNOMEDCT](#)

The API operation `InferSNOMEDCT` now uses version `2.3.0.20220301` . This update improves the detection of negations when using the `InferSNOMEDCT` API operation.

February 8, 2023

[Update API version for InferICD10CM](#)

The API operation `InferICD10CM` now uses version `2.2.0.20220401` . This update improves the detection of negations when using the `InferICD10CM` API operation.

February 8, 2023

[Update API version for DetectEntitiesV2](#)

The API operation `DetectEntitiesV2` now uses version `2.2.0`. This update improves the detection of negations when using the `DetectEntitiesV2` API operation.

February 8, 2023

[Update API version for DetectEntities](#)

The API operation `DetectEntities` now uses version `0.5.0`. This update improves the detection of negations when using the `DetectEntities` API operation.

February 8, 2023

[Update API version for InferSNOMEDCT](#)

The API operation `InferSNOMEDCT` now uses version `2.2.0.20220301`. This update fixes a bug that prevented the total set of attributes identified from being returned in the API output.

December 9, 2022

[Updated API version for InferSNOMEDCT](#)

The API operation `InferSNOMEDCT` now uses version `2.1.0.20220301`. This update includes accuracy improvements, and new traits are detected.

November 16, 2022

Updated API version for InferRxNorm	The API operation <code>InferRxNorm</code> now uses version 2.1.0.20221003. This update fixes a bug that prevented the total set of attributes identified from being returned.	November 16, 2022
Updated API version for InferICD10CM	The API operation <code>InferICD10CM</code> now uses version 2.1.0.20220401 . This update includes accuracy improvements, and new traits are detected.	November 16, 2022
Updated API version for DetectEntitiesV2	The API operation <code>DetectEntitiesV2</code> now uses version 2.1.0. This update includes accuracy improvements. It also adds new entities, attributes, relationships, traits, and a new category.	November 16, 2022
Updated API version for DetectEntities	The API operation <code>DetectEntities</code> now uses version 0.4.0. This update includes accuracy improvements, and new traits are detected.	November 16, 2022
Updated limits for StartSNOMEDCTInferenceJob and InferSNOMEDCT	The maximum individual file sizes <code>StartSNOMEDCTInferenceJob</code> and <code>InferSNOMEDCT</code> API operations are now 5 KB.	September 22, 2022

Text processing update	Amazon Comprehend Medical now provides improved support for correctly parsing non-relevant white space and carriage returns in all text analysis API operations.	September 12, 2022
Text processing update	Amazon Comprehend Medical now provides improved support for correctly parsing non-relevant white space and carriage returns in all ontology API operations.	September 12, 2022
Updated API version for InferSNOMEDCT	The Amazon Comprehend Medical API operation <code>InferSNOMEDCT</code> now uses the version 1.1.0.20220301.	July 8, 2022
Updated API version for InferICD10CM	The Amazon Comprehend Medical API operation <code>InferICD10CM</code> now uses the version 1.1.0.20220401.	July 8, 2022
Updated API version for InferRxNorm operation in Amazon Comprehend Medical	Amazon Comprehend Medical now uses the 2022-03-07 RxNorm and RxTerms Release for each RxCUI.	June 29, 2022
Updated models for Amazon Comprehend Medical	Amazon Comprehend Medical now provides improved support for NEGATION <i>trait</i> detection using Detect Entities V2.	April 25, 2022

[New feature for Amazon Comprehend Medical](#)

You can now use Amazon Comprehend Medical to establish a private connection with your virtual private cloud (VPC) by creating an interface VPC endpoint. For more information, see [VPC endpoints\(PrivateLink\)](#).

June 13, 2021

[New feature for Amazon Comprehend Medical](#)

Amazon Comprehend Medical now provides batch operations for ontology linking. This enables the service to detect entities in medical text stored in an S3 bucket and link those entities to standardized ontologies. For more information, see [Ontology Linking Batch Analysis](#).

May 4, 2020

[New feature for Amazon Comprehend Medical](#)

You can now use Amazon Comprehend Medical to extract and relate the date or time expression to any of the entities currently detected by Amazon Comprehend Medical. For more information, see [Detect Entities Version 2](#).

March 2, 2020

[New features](#)

Amazon Comprehend Medical now detects both medications and medical conditions and links the entities to established national ontologies RxNorm and ICD-10-CMs. For more information, see [Ontology Linking APIs](#).

[New feature](#)

Amazon Comprehend Medical now provides batch operations so that you can process medical text stored in an S3 bucket. It also provides a new model that you can use for examining your medical text. For more information, see [Detect Entities Version 2](#).

[Major new feature](#)

Amazon Comprehend Medical is a new service that detects useful information in unstructured clinical text: physicians' notes, discharge summaries, test results, case notes, and so on. Amazon Comprehend Medical uses natural language processing (NLP) models to leverage the latest advances in machine learning to sort through this enormous quantity of data and retrieve valuable information that is otherwise difficult to retrieve and use without significant manual effort. For more information, see [Amazon Comprehend Medical](#).

November 27, 2018

Amazon Comprehend Medical change log

The following sections detail updates to the Amazon Comprehend Medical service.

Updated API version for DetectEntitiesV2

Release date: 2026-03-19

Type: Automatic

The DetectEntitiesV2 API operation now uses model version 3.3.0.

Updated API version for DetectEntities

Release date: 2026-03-19

Type: Automatic

The DetectEntities API operation now uses model version 3.3.0.

Updated API version for InferICD10CM

Release date: 2026-03-19

Type: Automatic

The InferICD10CM API operation now uses model version 3.3.0.20251001.

Updated API version for InferRxNorm

Release date: 2026-03-19

Type: Automatic

The InferRxNorm API operation now uses model version 3.3.0.20221107.

Updated API version for InferSNOMEDCT

Release date: 2026-03-19

Type: Automatic

The InferSNOMEDCT API operation now uses model version 3.3.0.20220301.

Updated API version for InferICD10CM

Release date: 2025-11-03

Type: Automatic

The InferICD10CM API operation now uses model version 3.2.0.20251001. This update introduces support for the newly released 2026 ICD10CM code set that went into effect on October 1, 2025.

Updated API version for InferRxNorm

Release date: 2025-11-03

Type: Automatic

The InferRxNorm API operation now uses model version 3.2.0.20221107.

Updated API version for InferSNOMEDCT

Release date: 2025-11-03

Type: Automatic

The InferSNOMEDCT API operation now uses model version 3.2.0.20220301.

Updated API version for InferICD10CM

Release date: 2025-08-05

Type: Automatic

The InferICD10CM API operation now uses model version 3.1.0.20250401. This update introduces support for the newly released 2025 ICD10CM code set that went into effect on April 1, 2025.

Updated API version for InferRxNorm

Release date: 2025-08-05

Type: Automatic

The InferRxNorm API operation now uses model version 3.1.0.20221107.

Updated API version for InferSNOMEDCT

Release date: 2025-08-05

Type: Automatic

The InferSNOMEDCT API operation now uses model version 3.1.0.20220301.

Updated API version for DetectEntities

Release date: 2024-06-26

Type: Automatic

The DetectEntities API operation now uses model version 3.0.0.

Updated API version for DetectEntitiesV2

Release date: 2024-06-26

Type: Automatic

The DetectEntitiesV2 API operation now uses model version 3.0.0.

Updated API version for InferRxNorm

Release date: 2024-06-26

Type: Automatic

The InferRxNorm API operation now uses model version 3.0.0.20221107.

Updated API version for InferSNOMEDCT

Release date: 2024-06-26

Type: Automatic

The InferSNOMEDCT API operation now uses model version 3.0.0.20220301.

Updated API version for InferICD10CM

Release date: 2024-06-26

Type: Automatic

The InferICD10CM API operation now uses model version 3.0.0.20231001.

Changes to the Amazon Comprehend Medical API operations

Release date: 2024-02-07

Type: Automatic

In this update the InferICD10CM API operation was updated. The InferICD10CM API operation now uses model version 2.7.0.20231001. This update introduces support for the newly released 2024 ICD10CM code set that went into effect on October 1, 2023. This update also improves overall model accuracy.

Changes to the Amazon Comprehend Medical API operations

Release date: 2023-09-27

Type: Automatic

In this update, the InferICD10CM API operation was updated to use model version 2.6.0.20220401. This update improves low confidence scores and model accuracy.

Changes to the Amazon Comprehend Medical API operation

Release date: 2023-06-12

Type: Automatic

In this update the following API operations now use a new model version.

- DetectEntities now uses model version 3.3.0.
- DetectEntitiesV2 now uses model version 3.3.0.
- InferRxNorm now uses model version 3.3.0.20221107.
- InferICD10CM now uses model version 2.5.0.20220401.
- InferSNOMEDCT now uses model version 3.3.0.20220301.

This update adds the trait PAST_HISTORY to the medication category in the DetectEntities, InferRxNorm and DetectEntitiesV2 API operations.

This update also adds the attribute QUALITY to the medical condition category in the DetectEntitiesV2, DetectEntities, and InferICD10CM.

To learn more about Amazon Comprehend Medical's text analysis API operations, see [Text analysis API operations](#).

To learn more about Amazon Comprehend Medical's ontology linking API operations, see [Ontology linking](#).

Changes to the Amazon Comprehend Medical API operations

Release date: 2023-04-28

Type: Automatic

In this update the InferRxNorm API operations were updated. The InferRxNorm API operation now uses model version 2.1.0.20221003. This update improves the model's performance.

Changes to the Amazon Comprehend Medical API operations

Release date: 2023-04-21

Type: Automatic

In this update the InferSNOMEDCT API operation was updated. The InferSNOMEDCT API operation now uses model version 2.5.0.20220301. This update improves the model's performance.

Changes to the Amazon Comprehend Medical API operations

Release date: 2023-03-17

Type: Automatic

In this update the InferICD10CM API operation was updated. This update improves the detection of the DX_NAME entity and DIRECTION attributes. It also makes a minor system processing adjustment that impacts the detection of traits and attributes and the frequency/confidence scores in the API output.

Changes to the Amazon Comprehend Medical API operations

Release date: 2023-03-15

Type: Automatic

In this update, the following Amazon Comprehend Medical API operations were updated: InferSNOMEDCT, InferICD10CM, DetectEntities, and DetectEntitiesV2. These updates are automatic, and improve the detection of medical conditions, tests, treatments, and procedures, along with related attributes and traits, within the cardiology specialty.

Changes to Amazon Comprehend Medical API operations

Release date: 2023-02-08

Type: Automatic

In this update, the following Amazon Comprehend Medical API operations were updated: `InferSNOMEDCT`, `InferICD10CM`, `DetectEntities`, and `DetectEntitiesV2`. These updates are automatic, and improve the detection of negations in medical conditions when scanning clinical text.

Changes to the Amazon Comprehend Medical `InferSNOMEDCT` API operations

Release date: 2022-12-09

Type: Automatic

The Amazon Comprehend Medical API operation, `InferSNOMEDCT` was updated to use version `2.2.0.20220301`. This update fixes a bug that prevented the total set of attributes identified from being returned.

Batch analysis carried out using the `StartSNOMEDCTInferenceJob` API operation uses the same API version, `2.2.0.20220301`. To learn more about using these API operations, see [SNOMED CT linking](#).

Changes to Amazon Comprehend Medical API operations

Release date: 2022-11-16

Type: Automatic

In this update, the following Amazon Comprehend Medical API operations were updated: `InferSNOMEDCT`, `InferICD10CM`, `DetectEntities`, `DetectEntitiesV2`

DetectEntities

The Amazon Comprehend Medical API operation `DetectEntities` now uses API version **0.4.0**. This update includes accuracy improvements, and new traits are detected.

For **Category:** `TEST_TREATMENT_PROCEDURE`, **Type:** `PROCEDURE_NAME`, the following traits are added.

- `NEGATION`
- `PAST_HISTORY`
- `HYPOTHETICAL`

- FUTURE

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** TEST_NAME, the following traits are added.

- PAST_HISTORY
- HYPOTHETICAL
- FUTURE

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** TREATMENT_NAME, the following traits are added.

- NEGATION
- PAST_HISTORY
- HYPOTHETICAL
- FUTURE

For **Category:** MEDICAL_CONDITION, **Type:** DX_NAME, the following traits are added.

- HYPOTHETICAL
- LOW_CONFIDENCE
- PERTAINS_TO_FAMILY

DetectEntitiesV2

The Amazon Comprehend Medical API operation DetectEntitiesV2 now uses API version **2.1.0**. The update includes a new category BEHAVIORAL_ENVIRONMENTAL_SOCIAL along with updates to existing categories and traits.

For the new **Category:** BEHAVIORAL_ENVIRONMENTAL_SOCIAL, the following types have been added: GENDER, RACE_ETHNICITY, ALLERGIES, TOBACCO_USE, ALCOHOL_CONSUMPTION, REC_DRUG_USE, UnmappedAttributes.

For **Type:** GENDER

- No attributes or traits are available.

For **Type:** RACE_ETHNICITY

- No attributes or traits are available.

For **Type:** ALLERGIES, the following traits are added.

- NEGATION

- PAST_HISTORY

For **Type**: TOBACCO_USE, the following attributes and traits are added.

- Attributes
 - AMOUNT
 - DURATION
 - FREQUENCY
- Trait
 - NEGATION
 - PAST_HISTORY

For **Type**: ALCOHOL_CONSUMPTION, the following attributes and traits are added.

- Attributes
 - AMOUNT
 - DURATION
 - FREQUENCY
- Trait
 - NEGATION
 - PAST_HISTORY

For **Type**: REC_DRUG_USE, the following attributes and traits are added.

- Attributes
 - AMOUNT
 - DURATION
 - FREQUENCY
- Trait
 - NEGATION
 - PAST_HISTORY

For **Type**: UnmappedAttributes, the following attribute is added.

- AMOUNT

For **Category**: MEDICAL_CONDITION, **Type**: DX_NAME, the following traits are added.

- PERTAINS_TO_FAMILY

- HYPOTHETICAL
- LOW_CONFIDENCE

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** PROCEDURE_NAME, the following traits are added.

- NEGATION
- PAST_HISTORY
- HYPOTHETICAL
- FUTURE

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** TEST_NAME, the following traits are added.

- PAST_HISTORY
- HYPOTHETICAL
- FUTURE

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** TREATMENT_NAME, the following traits are added.

- NEGATION
- PAST_HISTORY
- HYPOTHETICAL
- FUTURE

Batch analysis carried out using the `StartEntitiesDetectionV2Job` API operation uses the same API version, **2.1.0**. To learn more about using these API operations, see [Detect entities \(Version 2\)](#).

InferSNOMEDCT

The Amazon Comprehend Medical API operation `InferSNOMEDCT` now uses API version **2.1.0.20220301**. This version update adds new traits to the following categories and types.

For **Category:** MEDICAL_CONDITION, **Type:** DX_NAME, the following traits are added.

- PERTAINS_TO_FAMILY
- HYPOTHETICAL
- LOW_CONFIDENCE

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** TEST_NAME, the following traits are added.

- PAST_HISTORY
- FUTURE
- HYPOTHETICAL

For **Category:** TEST_TREATMENT_PROCEDURE, **Type:** TREATMENT_NAME, the following traits are added.

- NEGATION
- PAST_HISTORY
- FUTURE
- HYPOTHETICAL

Bug fixes

- This update fixes a bug that prevented the total set of attributes identified from being returned.

Batch analysis carried out using the StartSNOMEDCTInferenceJob API operation uses the same version. To learn more about using these API operations, see [SNOMED CT linking](#).

InferICD10CM

The Amazon Comprehend Medical API operation InferICD10CM now uses API version 2.1.0.20220401. This version update adds new traits to the following categories and types.

For **Category:** MEDICAL_CONDITION, **Type:** DX_NAME, the following traits are added.

- PERTAINS_TO_FAMILY
- HYPOTHETICAL
- LOW_CONFIDENCE

Bug fixes

- This update fixes a bug that prevented the total set of attributes identified from being returned.

Batch analysis carried out using the StartRxNormInferenceJob API operation uses the same version. To learn more about using the InferRxNorm and StartRxNormInferenceJob API operations, see [RxNorm linking](#).

InferRxNorm

The Amazon Comprehend Medical API operation InferRxNorm now uses API version 1.3.1.20221003.

Bug fixes

- This update fixes a bug that prevented the total set of attributes identified from being returned.