



Kick off your AI journey with Video Indexer

Automatically extract insights from media content

Jun Heider | Chief Technology + Operating Officer | RealEyes Media
in coordination with Microsoft Azure

Published September 2019





Introduction

Is your media library growing too fast for your staff to manage alone? Unfortunately, as content continually increases, the hours in a day remain the same (surprise!). This means we rarely have enough time or resources to manage content discoverability for our media libraries. But using an indexing solution that leverages artificial intelligence (AI), you can automate how you manage and index your media.

Imagine the time your staff invests in transcribing and adding metadata tags to videos. As new content volume increases, your staff needs more time to tag it. One person needs at least one hour to process a one-hour video, right? And if you add more people to a project, potential human mistakes like typos, inconsistent taxonomy, varied labeling, and so on increases.

But what if you could use AI to automatically process, index, and tag your videos? By automating tagging work, your team could recover valuable time for

contributing to other areas of the business.

Using an AI-based solution, you can automate and increase your ability to tag your content library with rich metadata, which helps you to discover your content faster and more efficiently. It also helps expose how your content relates to other items in your growing library. You can also use AI to automatically generate transcriptions and translations for your media, improving the reach of your content to a wider global audience and enabling you to meet accessibility regulations more easily.

AI provides value by augmenting your human team, not replacing it. For instance, your team can focus on determining compelling content for social media clips while using AI to surface source content for your colleagues to review, or even to automate creating clips and sharing them to social media.

By using AI to transform how you manage media content, you can ensure your staff can focus on more meaningful work, while AI manages your menial tasks.





Enter Azure Video Indexer

Microsoft Azure is a set of cloud computing services you can use to build, manage, and deploy applications on a massive, global network using your preferred tools and frameworks. Azure offers multiple services in the media and AI spaces.


This paper focuses on Azure Media Services, and specifically on how the Video Indexer service helps with your media, AI, and machine learning (ML) needs.

[Azure Media Services](#) is a robust platform for live and VOD workflows. Media Services provides all the building blocks you need to build your own video-based solution service, including services to encode, package, index, distribute, and protect your content.

Video Indexer is one of the many modular services in the Azure Media Services family. Video Indexer uses a set of ML algorithms already integrated together to extract the metadata automatically and eliminate the need for many (many) hours of manual tagging. You can use Video Indexer in combination with other services in Media Services, or as a standalone through its own API and portal experience.

Video Indexer includes a variety of features, like facial recognition, general sentiment analysis, transcripts, transcript translations, and content moderation. For instance, Video Indexer can call out adult scenes and profanity in your media content. This only scratches the surface of the value that Video Indexer provides by bringing together all these insights for you within one service, while providing predictable pricing.

Let's take a closer look at how to get started with Video Indexer.





Getting started with Video Indexer

Getting started with Video Indexer is very easy and free. The application UI is robust, so both developers and non-developers can use the VI service right away. The first thing to do is set up an account at [videoindexer.ai](https://www.videoindexer.ai), and then you can start indexing your videos and audio files.

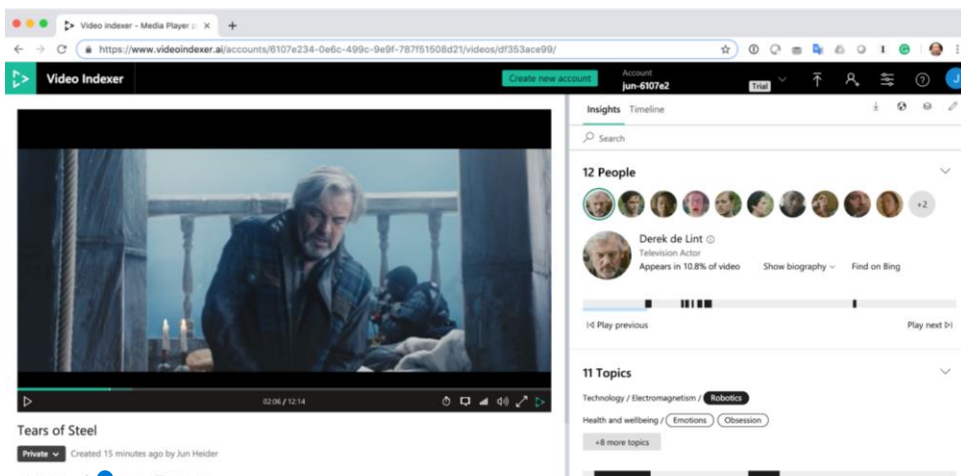
When you've created your free account, you will have a generous amount of trial hours to use for indexing your video and audio content. If you use the portal, you can index up to 10 hours of video for free. If you're a developer and use the API, you can index up to 40 hours of video. You can track how much you've used in your account settings.

The Video Indexer portal experience is very interactive. On the Insights tab, you can click through specific faces, keywords, labels, brands, sentiment, even scenes and keyframes. On the Timeline tab, you can track which entities exist at each second within the timeline of the video in addition to displaying the audio transcript.

Video Indexer also has a list view to search and navigate between videos you've uploaded, along with a video editor for making clips of the media within your library. All in all, Video Indexer is a very powerful tool for you to use right out of the box.

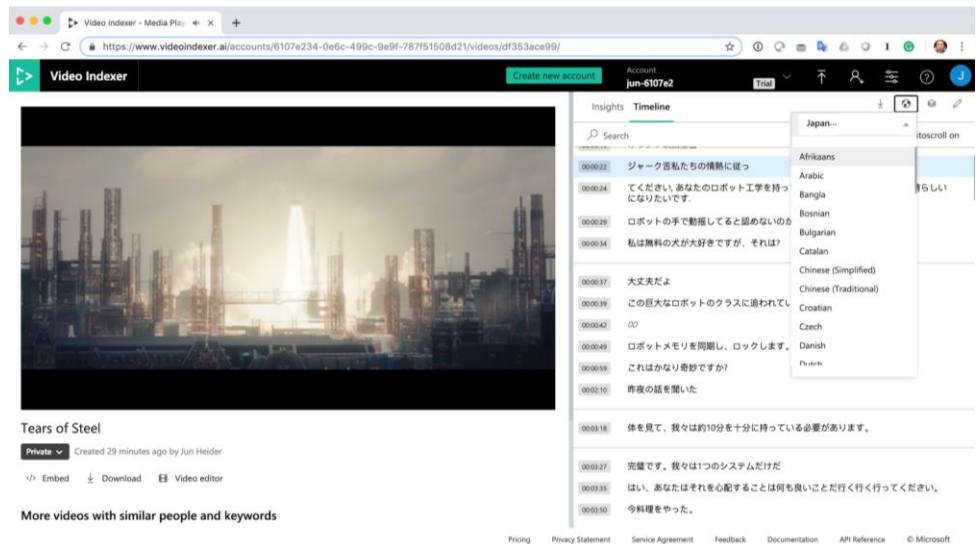
To find out more, view the [documentation](#).

Video Indexer Insights



More videos with

Video Indexer Timeline





Integrating Video Indexer into your workflows

The **Video Indexer API** makes integration into your own custom workflows a breeze. The main asset you will work with is a JSON file that contains the output of the indexing process. The API itself is easy to work with, and you can find the API definitions available as downloadable Open API (Swagger) descriptor files. You can also integrate the Video Indexer API into your own tooling or client SDK generation.

Language-specific code snippets are at the bottom of each endpoint reference. You can even use the "try it" feature to test API calls from directly within the documentation.

Although you can do generally anything in the UI that you can do via the API, in its simplest form, indexing your media via the API is a three-step process:

1. [Get your account access token](#) (token expires within an hour).
2. [Upload your video](#) for indexing.

3. In the previous POST request you can provide a callback URL for VI to call when the processing is complete. VI will invoke the callback and then you can [get the Video Indexer Insights](#) JSON.

Below is a view into the Insights JSON, which captures a wealth of information. For a deeper dive, take a closer look at [Video Indexer API output](#).

Additionally, with the Video Indexer API you can [re-use and re-embed](#) controls from the web portal directly into your own hosted web applications and CMS tooling. That means you can access the same video player and insights viewer HTML controls you see in the Video Indexer user experience, reducing your development costs and maintenance of custom controls built on the API. This saves you a lot of development effort when you are building video-based apps.

Insights
JSON

```

3835 "video": {
3836   "id": "6107e234-0e6c-499c-9e9f-787f51508d21",
3837   "accountID": "6107e234-0e6c-499c-9e9f-787f51508d21",
3838   "id": "df353ace99",
3839   "status": "Processed",
3840   "metadata": {
3841     "reviewState": "None",
3842     "privacyMode": "Private",
3843     "processingProgress": "100%",
3844     "failureCode": "None",
3845     "failureMessage": "",
3846     "externalID": null,
3847     "externalURL": null,
3848     "metadata": null,
3849     "insights": [
3850       {
3851         "version": "1.0.0.0",
3852         "duration": "0:12:14.289",
3853         "sourceLanguage": "en-US",
3854         "language": "en-US",
3855         "transcript": [
3856           ],
3857         "ocr": [
3858           ],
3859         "keywords": [
3860           ],
3861         "topics": [
3862           ],
3863         "faces": [
3864           {
3865             "id": 2420,
3866             "name": "Derek de List",
3867             "confidence": 0.9991,
3868             "description": "Derek de List",
3869             "thumbnailID": "1c120ba0-30d0-4806-9e4b-d8e255",
3870             "referenceID": "562fab91-8db-854f-16af-7378c1",
3871             "referenceType": "Bing",
3872             "title": "Television Actor",
3873             "imageURL": "https://www.bing.com/th?id=AM95...",
3874             "thumbnail": [
3875             ],
3876             "fileID": "655d19f-2f10-42e5-8718-84195f...",
3877             "fileName": "FaceThumbnail_655d19f-2f10-42e5-8718-84195f...",
3878             "instances": [
3879             ]
3880           }
3881         ]
3882       }
3883     ]
3884   }
3885 }

```

Re-use &
re-embed

Share & Embed

Copy share URL

Select widget

Copy embed code

```
<iframe width="580" height="780"
src="https://www.videoindexer.ai/embed/insights/6107e234-0e6c-499c-9e9f-787f51508d21/df353ace99/?
accessToken=eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzI1NiJ9.eyJ1IjoiYXJlY2N0dW50SWQI"

```



Model customization in Video Indexer

Although **machine learning** has come a long way in recent years, in some cases you might want to customize or train the system further to be fine-tuned for your specific content.

For instance, you might have specific employees in your organization who are not celebrities but are noteworthy nonetheless. Another case might be that your organization uses a large amount of technical or industry-specific (i.e. legal, healthcare) jargon, or you need to know when specific brands are discussed, so you might want to create your

your own organizational or brand dictionary.

Video Indexer exposes the ability for you to customize your content model. It currently supports three areas: [People](#), [Language](#), and [Brands](#). You can customize the content model from both the web portal and also via the API. Customizations are private to each account, preventing one organization's customizations from affecting another's.

For more in-depth details, please visit the [updated documentation on customization capabilities](#).

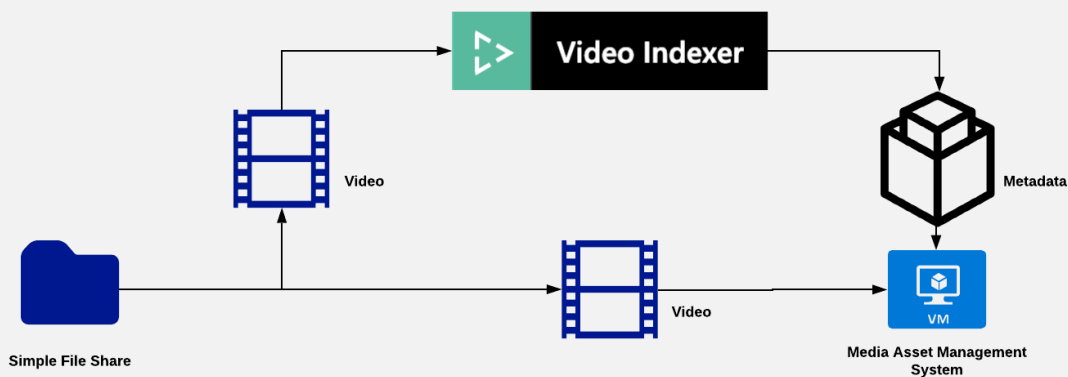
Reviewing some use cases

Video Indexer itself is a modular service in the Azure Media Services family, useful on its own. The API opens up opportunities for custom applications and integration. Let's take a look at some particularly compelling use cases



Enhancing discoverability in your media asset management (MAM) system

Imagine you have thousands of hours of videos in your media library and are moving from simple file share to MAM system. Without metadata, searching through the video becomes more difficult as your library grows. How about letting Video Indexer do the heavy lifting?

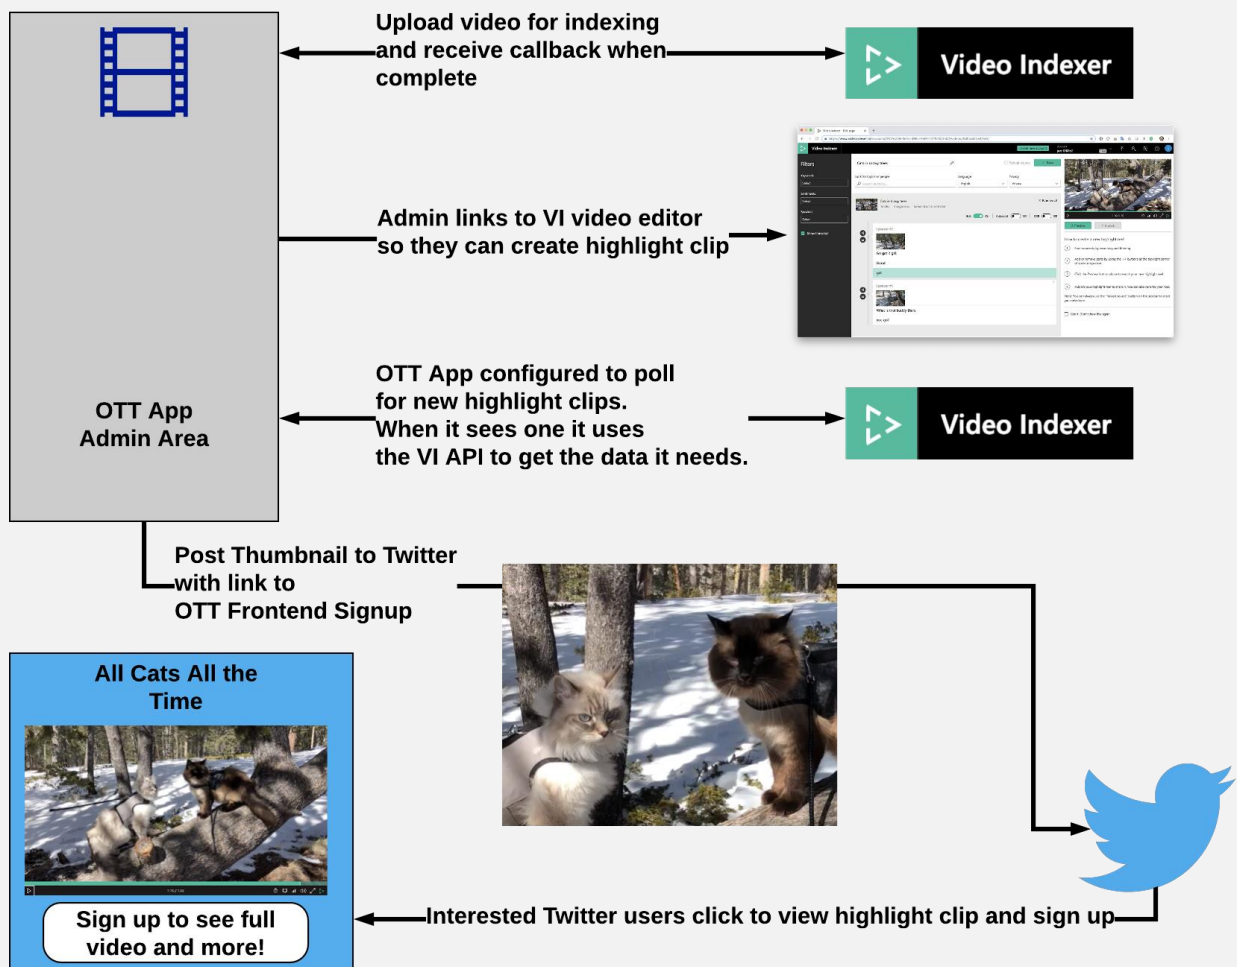


In the diagram you will notice Video Indexer as a step in the video ingestion process. Its responsibility is to extract the metadata from the uploaded audio and video files and push that over to the MAM. Once the ingestion is complete, you can immediately search and discover your videos based on the metadata automatically discovered by Video Indexer.



Sharing highlights on social media

A great way to monetize and bring traffic to your videos is to share highlight clips on social media or your own custom websites. Say you had a subscription-based over-the-top (OTT) service for cats. You can create a highlight clip from your latest cat video using Video Indexer's video editor, and then build in hooks to post a thumbnail of the highlight clip to Twitter with a link back to your OTT service sign-up page. On your sign-up page, you can easily embed the clip via the Video Indexer HTML player widget, so your audience can play it back. If they like what they see, they'll join your OTT service as a subscriber. Once they join, you can link them to the full video content, once again using the Video Indexer player widget.




The metadata generated by Video Indexer helps you find the appropriate highlights. Then, you use the video editor to create the highlight clip using search and video preview. Using the VI API from your application, you poll for new clips, get the data you need, and then post to Twitter. For more information, take a closer look at the [List Videos](#), [Get Thumbnail](#), and [Get Widget Player](#) APIs highlighted in this use case.



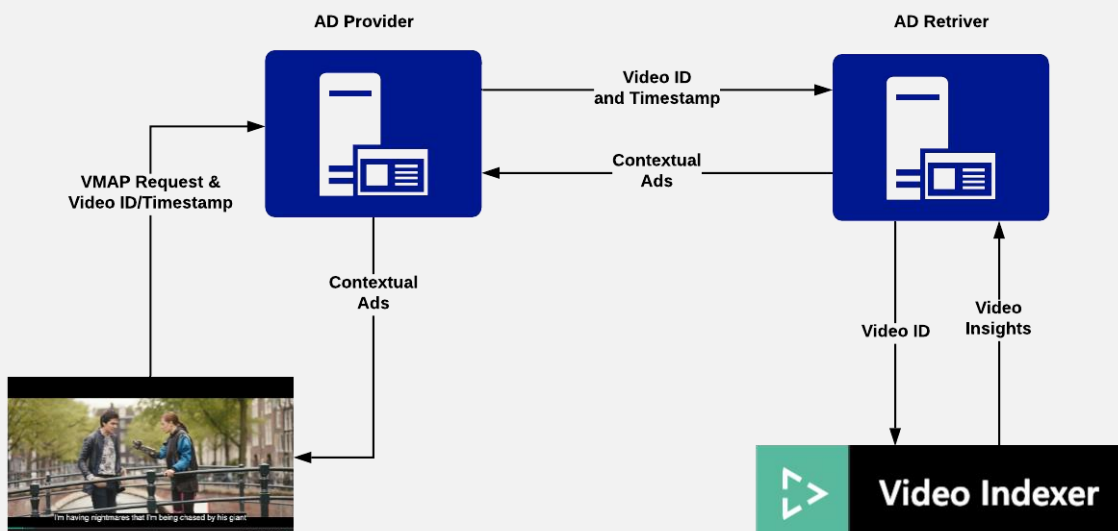
Improving online safety via automated content filtering

When your video platform is open to user-generated content, a high risk exists that inappropriate or offensive content will make its way into your system and onto your websites. Manually reviewing all content before making it available to viewers is ineffective, costly, and unscalable. Having an automated check in place before content is published will allow your platform's UGC volume to grow while helping to prevent questionable content from making its way onto your websites or reaching your audience. Video Indexer has a built-in feature to flag videos it deems racy or adult and will keep them private and safe from your viewers. You can learn more about content moderation [here](#).



Using Video Indexer Insights for contextual advertising

You must maximize the impact of commercials when your video monetization strategy involves advertising. Contextual or "targeted" advertising is a hot area that factors in your video content when triggering specific commercials. For instance, a good time to run a cat food commercial might be during a video documentary about adventure cats. You can use the insights generated by Video Indexer to drive your contextual advertising strategy.

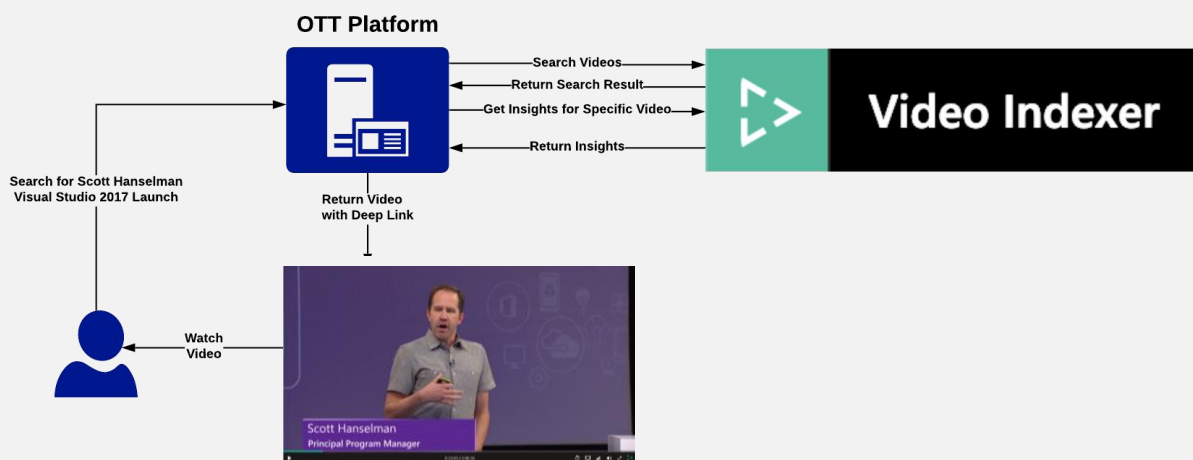




Enabling deep linking

Many viewers have a short attention span or very limited time to watch your content. If you want them to keep coming back to your OTT platform, you must make each and every visit memorable. Wouldn't it be great if a user could search for a specific object, actor, or spoken phrase in one of your videos and start playback right when the subject appears in the video?

You can use Video Indexer to search for videos that contain multiple subjects (people, topics, objects, and more) your viewers are looking for and then link those viewers directly to the point in time of the video where they can see what they're looking for. For instance, if a viewer is interested in watching the Visual Studio 2017 Launch event, but only when Scott Hanselman starts talking, they could search his name and start watching the video when he appears, thereby saving over 13 precious minutes of viewing time.

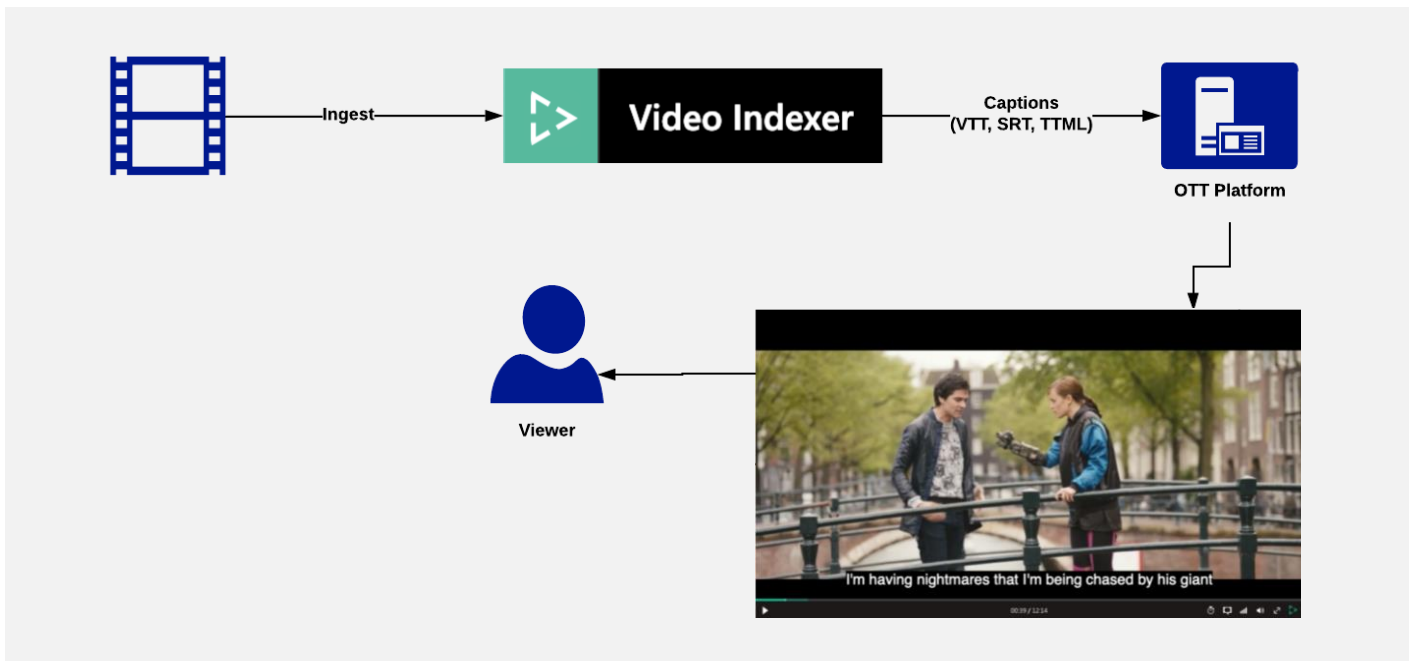


You can learn more about the Search Videos and Get Video Index [API calls](#). As a content owner, you may want to bring this power to your workflows and external applications as well. You can use Video Indexer many ways to build compelling, integrated applications. No matter the scenario, Video Indexer will help kickstart your AI transformation.



Enhancing content accessibility with automated transcription and subtitle generation

The Accessibility Act of 2010 mandated providing closed captions if technically feasible and achievable. Video Indexer transcribes content and gives you the option to generate captions in many languages. Once Video Indexer generates caption files, you can then take them and package them together with your streaming assets to enable closed captioning for your viewers. In addition, because Video Indexer can also translate the text, you can provide your customers a selectable subtitle track so they can read the translated audio in the language they choose. Read the [documentation](#) to learn more about captions in Video Indexer.



Taking the next step with Video Indexer

Now that you understand the value of Video Indexer, you can trial the solution with your own content in minutes.

If haven't yet done so, sign up for a free Video Indexer trial account at videoindexer.ai. You can then immediately start uploading and gaining valuable cognitive insights on your videos. When you are ready for production, or if want to go beyond the resources of the trial, you can quickly scale up for production-level service and full SLA by creating a Video Indexer account on your Azure subscription. This will provide you with a fully isolated set of resources in your own Azure subscription in the region of your choice, Media Services and storage accounts, and the compliance of the [Azure trusted cloud](#). Once you create an account, you have access to unlimited indexing capabilities.

Next, [familiarize yourself with VI](#). If you want to build applications, the [API reference](#) will come in handy.

If you need some inspiration, keep an eye on the [Video Indexer blog](#). For instance, [read how to use VI for live stream analysis](#).

If you have ideas for improving Video Indexer, you can provide feedback via this [forum](#). If you have any questions, email the VI team at VISupport@microsoft.com.

The Video Indexer trial does not have an official SLA or support channel outside of the community discussion site, but you can post issues to this [page](#), and the team and VI community will jump in to help. Once you've upgraded to a paid account, you will have a full SLA with the ability to submit support requests directly to [Azure support](#).

Also, you may want to start customizing content models relevant to your organization. As of this document release, Video Indexer supports customizing content models for brands, language, and people. All three of these content model types support [customization](#) through the API.

Video Indexer is a powerful ally that gives you valuable insights into your media library right away. When you're ready, you can extend Video Indexer through integrations. While on your AI journey, pull over to let Video Indexer join you!

RealEyes Media (realeyesmedia.com) delivers exceptional viewer experiences for broadcasters, media, and e-gaming companies, as well as for brands whose needs cannot be met by off-the-shelf software. From bespoke application development to complex media aggregation and intelligence, RealEyes' versatile team of technologists are adept at complementing existing technology teams to tackle some of the industry's most complex problems. With skill sets ranging from user experience design to full stack development, RealEyes can help from ideation through implementation—building highly configurable, personalization-driven applications for web-mobile and set-top—and the backend management and insights to empower scale and enable effective business decisions. RealEyes are the go-to experts in tackling hard and complex problems such as ultra-low latency at scale, dynamic ad insertion, multi-CDN delivery, and complex failover strategies for mission-critical video streaming

© 2019 Microsoft Corporation and RealEyes Media. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred. This document does not provide you with any legal rights to any intellectual property in any Microsoft or RealEyes product. You may copy and use this document for your internal reference purposes.

