



V-Spark 3.5.0 Frequently Asked Questions

V-Spark 3.5.0 Frequently Asked Questions

© Copyright 2019 Voci All Rights Reserved.

The information contained in this document is the proprietary and confidential information of Voci Technologies, Inc. You may not disclose, provide or make available this document, or any information contained in this document, to any third party, without the prior written consent of Voci.

The information in this document is provided for use with V-Spark Voice Analytics. No license, express or implied, to any intellectual property associated with this document or such products is granted by this document.

All Voci Technologies, Inc. products described in this document, including V-Spark Analytics and others prefaced by Voci are owned by Voci (or those companies that have licensed technology to Voci) and are protected by patents, trade secrets, copyrights or other industrial property rights. The Voci products described in this document may still be in development. The final form of each product and release date thereof is at the sole and absolute discretion of Voci. Your purchase, license and/or use of Voci products shall be subject to Voci's then current sales terms and conditions.

Trademarks

The following terms used in this document are trademarks of Voci Technologies, Inc. in the United States and other countries:

- Voci
- V-Blaze
- V-Cloud
- V-Discovery
- V-Ferno
- V-Purify
- V-Spark

Other third party disclaimers or notices may be set forth in Voci's online or printed documentation. All other product and service names, and trademarks not owned by Voci are the property of their respective owners.

Table of Contents

1. Sign-up and Support	1
1.1. How do I invite new users to V-Spark?	1
1.2. What are the deployment options for Voci products?	1
1.3. How do I contact support?	1
1.4. What are the default url:ports for Voci products?	2
2. User Documentation	3
2.1. Where can I find V-Spark's User Documentation?	3
2.2. What documentation is available for V-Spark?	3
3. Languages and Language Models	5
3.1. What types of language models are available?	5
3.2. What languages do you support?	5
3.3. What language model should I use for my audio?	5
3.4. What is the difference between the acoustic and language models?	5
3.5. Can I add words to the language model vocabulary?	6
3.6. Why is a language model that I added not showing up in V-Spark?	6
3.7. How do I identify the different language models?	6
3.8. How do I display the active language models in my environment?	7
4. Starting / Stopping / Checking status of V-Spark	9
4.1. How do I check the status of the server?	9
4.2. How do I start the V-Spark service?	10
4.3. How do I restart the V-Spark service?	10
4.4. How do I stop the V-Spark service?	10
5. Audio/Uploading Audio/Metadata	11
5.1. How do I upload audio?	11
5.2. How quickly will I see my audio after I upload?	11
5.3. What kind of audio can V-Spark support?	11
5.4. What audio quality works best in V-Spark?	11
5.5. How do I upload metadata with my audio?	11
5.6. How do I know what my audio source is?	12
5.7. How do I see my own usage?	12
5.8. Can Voci products process 1,000,000,000 audio hours per year?	12
5.9. Can Voci Technologies, Inc. support 10,000 call center agents?	12
5.10. Why does the API require zip files for transcription?	13
6. Folder Creation/Maintenance	15
6.1. Which language model should I pick?	15
6.2. How do I start a local callback server for testing?	15
6.3. What are the consequences of moving results after a period of time?	16
6.4. What does "Automatically send transcript (JSON) files to a URL" mean?	16
7. The Audio Evaluator	17
7.1. How do I use the Audio Evaluator?	17
7.2. Who can see the audio that I test using the Audio Evaluator?	17
7.3. How do I know how many hours of audio I will have?	17
8. Viewing/Searching Data in V-Spark	19
8.1. What kinds of search queries are supported?	19
8.2. I don't see the audio that I just uploaded. Where is it?	19
8.3. I uploaded audio a while ago, but I don't see it anymore. What happened?	19
8.4. I just created a new application - why don't I see my data?	19
8.5. I just reprocessed audio for an application - why don't I see my data?	19
8.6. Why are parts of the transcript red/green?	20
8.7. The display in V-Spark is not showing the correct information	20
8.8. Why is the search score always 1?	20

9. Audio Transcripts	21
9.1. How do I download transcripts for the audio I uploaded?	21
9.2. What's the difference between JSON and Text transcripts?	21
9.3. How do I interpret the JSON output of my transcript?	21
10. Applications	23
10.1. Why can't I edit this application?	23
10.2. What pre-made application templates does Voci provide?	23
A. Voci Glossary	25

Chapter 1. Sign-up and Support

1.1. How do I invite new users to V-Spark?

Navigate to Settings/Accounts and locate your company in the list. Right click on "Register..." and select "Copy link address". You can now copy the registration link into an email message and send it to your colleagues.

1.2. What are the deployment options for Voci products?

The following are the options that are available to you:

- In Voci's AWS cloud
- In an instance in your own AWS cloud (on Amazon instances verified for capabilities by Voci). You will need to provide Voci with an Amazon ID so that we can place an AMI in a location from which you can retrieve it.
- On premise. We supply hardware for your production environment or can verify the capabilities of hardware that you supply.

Voci can provide the first 2 environments at little or no cost for 100 hours of audio for 30 days.

Voci delivers 3 different product levels in any of these environments: V-Blaze, V-Blaze + Options, and V-Spark. These provide the following components and capabilities:

- **V-Blaze:** pure automatic speech recognition and transcription, often referred to as speech to text. (API only)
- **V-Blaze + Options:** automatic speech recognition and transcription (speech to text), plus the redaction of privacy info (PCI), diarization (speaker splitting) of mono audio, emotion/sentiment analysis, and gender analysis. (API only)
- **V-Spark:** all of V-Blaze and V-Blaze + Options, plus a speech analytics platform with a graphical user interface

If you are interested in one of these platforms, contact Voci at <support@vocitec.com>.

1.3. How do I contact support?

If you purchased V-Spark through a Voci partner, you may need to contact that company's customer support group first. That group will contact Voci support if necessary. Voci's product support can be contacted at <support@vocitec.com>.

1.4. What are the default url:ports for Voci products?

If you are using V-Blaze, Voci's automatic speech recognition engine, a web-based monitoring application is available at the following URL:

```
https://hostname:8088
```

The V-Blaze REST API is accessible at the following URL, but requires that you specify the methods that you want to use:

```
https://hostname:17171
```

The V-Blaze REST API is described in the document "*Using the V-Blaze API*".

If you are using V-Spark, Voci's analytics engine, a web-based user interface is available at the following URL:

```
https://hostname:3000
```

The V-Spark REST API can be contacted at the same URL, but requires that you specify the methods that you want to use. The V-Spark API is described in the document "*Using the V-Spark 3.5.0 API*".

Chapter 2. User Documentation

2.1. Where can I find V-Spark's User Documentation?

All of the V-Spark User Documentation can be found under the V-Spark Help menu.

2.2. What documentation is available for V-Spark?

- **V-Spark 3.5.0 Quickstart Guide** - Provides basic information to get a new V-Spark user registered for an account, logged in, processing audio, and reviewing analysis.
- **V-Spark 3.5.0 Review and Analysis Guide** - Provides more detailed information about interpreting call statistics, application scores, and setting up emailed reports.
- **V-Spark 3.5.0 Management Guide Management Guide** - Provides information about customizing analysis, uploading audio, and managing users, organizations, and folders.
- **V-Spark 3.5.0 Application Development Guide** - Provides information about creating V-Spark applications, and viewing application scores.
- **Using the V-Spark 3.5.0 API** - Provides information on using the REST API to accomplish tasks in V-Spark without using the V-Spark user interface.
- **V-Spark 3.5.0 Regular Expression Guide** - Provides information about the type and format of regular expression queries that are supported by V-Spark.
- **V-Spark 3.5.0 Frequently Asked Questions** - Answers common questions about V-Spark.
- **Voci JSON Output Format Guide** - Provides information about the JSON transcript format that is used by all Voci products and the fields that it contains.

Chapter 3. Languages and Language Models

3.1. What types of language models are available?

We have the following models to choose from:

1. Call Center
2. Financial
3. Survey
4. Telecommunications
5. Voicemail

Work with your Voci sales representative to decide which models and languages should be available in your V-Spark installation out of the box. Send a note to <support@vocitec.com> to get started.

3.2. What languages do you support?

At this time, Voci supports North American, Australian, European, and UK English as well as North American and European Spanish and Canadian French. In addition, we have early access versions of Dutch, Portuguese, German, European French and Italian. These all require additional tuning based on audio supplied by the customer.

3.3. What language model should I use for my audio?

V-Spark Typically uses the standard North American English Call Center and Survey models. If for some reason, transcription results are not good, please discuss this with Voci. There may be additional models that are a better fit for a particular use case or customer. See [Section 6.1, “Which language model should I pick?”](#) for more information about the language models that are supported by Voci products.

3.4. What is the difference between the acoustic and language models?

In Voci products, the acoustic model is how things sound. The language model is what exactly was said.

3.5. Can I add words to the language model vocabulary?

Voci uses Large-Vocabulary Continuous Speech Recognition (LVCSR, also known as speech-to-text, full transcription, ASR, or automatic speech recognition), which uses a set of words as its basic unit of understanding. Since Voci's language models are LVCSR rather than phonetics-based, they are very context-dependent. Therefore, all new additions to language model vocabulary must be tested and the language models tuned accordingly. Voci Application Engineers can provide interested customers with more information about custom tuning.

3.6. Why is a language model that I added not showing up in V-Spark?

If the models licensed on a machine change, you must restart `vociserver` and `vociwebapi`. When logged in as the `vociadmin`, the commands are:

```
service vociserver restart
service vociwebapi restart
```

When logged in as any other user, the commands are:

```
sudo service vociserver restart
sudo service vociwebapi restart
```

3.7. How do I identify the different language models?

The following is a list of common language/model codes. Your machine may not be currently licensed to use all of these models. Each model uses a significant amount of memory, so the combinations you will be able to use will vary depending on the size in memory and processing power of the machine you choose.

Languages (LANG):

- `eng1` = US/NorthAmerican English
- `eng2` = Australian English
- `eng3` = UK English
- `fre1` = French
- `spa1` = NA Spanish
- `spa2` = Castilian/Euro Spanish

- `ger` = German

Common Model Variants (MODELVAR):

- **callcenter** -tuned for typical agent/client interactions and high throughput
- **voicemail** - tuned for voicemail messages, and provides a larger vocabulary. Often effective on non-specific miscellaneous audio.
- **telecom** - large vocabulary within the overall domain of telephone communications

You will need access to <https://vocitec.com/> in order to install additional models. The rpm packages for these models can be in excess of 1GB in size and can have multiple dependencies. You'll also need corresponding licensing to enable/run each model.

3.8. How do I display the active language models in my environment?

The API call to use is the `/models` method, such as the following, where `HOSTNAME` is the name of your ASR or V-Blaze host:

```
curl HOSTNAME:17171/models
```

This call returns the list of options that are usable in a transcribe call as the value of the `model` parameter. They are of the form `LANG:MODELVAR`.

The following is an example of this call and the values that it returns:

```
curl asr.example.com:17171/models
{"models":["eng1:callcenter","eng3:callcenter","spa1:callcenter"]}
```

You can modify the list of active models by editing the `/opt/voci/models.list` file and commenting out or uncommenting model entries. You can then cause the ASR system to restart with the new model settings by executing the `sudo service vociserver restart` command.

Chapter 4. Starting / Stopping / Checking status of V-Spark

4.1. How do I check the status of the server?

To check the status of the running processes associated with V-Spark, issue the following command:

```
sudo service vspark status
```

To get detailed information about the status of the ASR Engine used by V-Spark, you can also issue a command like the following from the machine on which V-Spark is installed:

```
curl localhost:17171/status; echo
```

```
{
  "status": {
    "bytesdone": 29641600,
    "idlefor": 309,
    "totalstreams": 198,
    "framesrecv": 185260,
    "uttsdone": 78,
    "bytesproc": 29641600,
    "speed": 0.001674517984328025,
    "dectime": "12 days, 19:19:08.225184",
    "framesproc": 185260,
    "bytesrecv": 29641600,
    "framesdec": 95815,
    "speedrecv_Bps": 26.792287749248402,
    "#done": 3,
    "uttsproc": 78,
    "started": "2016-12-20 15:17:12.535789",
    "lag": 0.0,
    "speeddec": 0.0008660473964611342,
    "speedrecv": 0.001674517984328025,
    "framesdone": 185260,
    "#queued": 0,
    "#active": 0,
    "totalidle": 1901472,
    "lastactive": "2017-01-11 15:39:33.281058"
  }
}
```

Here are some other examples of other commands that will give you information from a running V-Spark server:

```
curl localhost:17171/server; echo
```

```
/server
```

```
curl localhost:17171/models; echo
```

```
{"models":["enl:callcenter","enl:voicemail","enl:survey]}
```

```
curl -F output=text -F file=@/opt/voci/server/examples/sample1.wav localhost:17171/transcribe; echo
```

```
[{"source":"sample1.wav","utterances":["And that it was resolved in a very professional manner. Your employees a very good."]}
```

4.2. How do I start the V-Spark service?

If the status of the service is stopped or you know the service is not running, issue the following command to start the V-Spark service:

```
sudo service vspark start
```

4.3. How do I restart the V-Spark service?

If you have made a change or you know you need to restart the V-Spark service, issue the following command to do so:

```
sudo service vspark restart
```

4.4. How do I stop the V-Spark service?

If you need to stop the V-Spark service, issue the following command:

```
sudo service vspark stop
```

Chapter 5. Audio/Uploading Audio/Metadata

5.1. How do I upload audio?

To upload audio, you must first create a folder by navigating to Settings/Folders and clicking the blue '+' on the right-hand side. Once the folder is created, you can upload a single audio file or a zip folder of multiple audio files and their accompanying metadata by clicking the blue upload button next to the folder.

5.2. How quickly will I see my audio after I upload?

Transcript results will begin to appear within a few minutes, but it may take a while for every file in larger batches to appear. For example, if you upload 100 hours of audio at once, you can expect all 100 hours to have appeared on the Dashboard within 1 hour. If you don't see your files on V-Spark within 24 hours, check the log next to the folder the files were uploaded to for any processing errors, then contact support.

5.3. What kind of audio can V-Spark support?

V-Spark can support most common audio formats. If you're not sure whether your audio format is supported, you can test it using the Audio Evaluator tool found under the help menu as "Will my audio work?", or at the URL <https://vspark-demo.vocitec.com/audioEvaluator>.

Voci cannot guarantee that proprietary or encrypted formats will be supported.

5.4. What audio quality works best in V-Spark?

V-Spark's results are most accurate when there is one speaker per audio channel and/or little speaker overlap, little or no background noise, and little or no audio compression. Choosing the right language model for your audio is also important for accuracy.

5.5. How do I upload metadata with my audio?

You can upload metadata with your audio by including a metadata file for each audio file in the zip folder that you want to upload. For more information about uploading metadata and using custom metadata fields, refer to the "*V-Spark Management Guide*". See [Chapter 2, User Documentation](#) for information about locating Voci documentation and the information that each document provides.

5.6. How do I know what my audio source is?

In V-Spark, audio source refers to the number of channels in the original audio file. Mono indicates that the audio file has only one channel, while Stereo indicates that it has two channels. If you are unsure what the audio source is for your files, you can use the Audio Evaluator tool to find out.

5.7. How do I see my own usage?

Login to V-Spark. Navigate to **Settings > Accounts > Companies**. Choose your company and click on the usage bar for that company. This is available to any user with **Create/Edit** permissions in the specified company.

5.8. Can Voci products process 1,000,000,000 audio hours per year?

Yes.

Per-box max throughput for V-Spark is typically 800K-1M hr/yr, depending on hardware, models, and specific audio characteristics. That equates to a 24/7/365 continuous throughput of 90-115X. You give the machine 90 to 115 hours of audio, you get your results in 1 hour.

If you anticipate higher instantaneous/avg throughput requirements or more overall volume per year, you can start stacking ASR Engines behind an HTTP/S load balancer. That is the typical configuration that is used by Voci customers when they start running high volumes. If you provision more capacity beyond the minimum number that are required, this load balancer approach also provides some levels of redundancy/failover.

Licensing is variable and can support whatever volumes you require. For information about rates for various volume amounts or about increasing your overall usage, contact your Voci sales person or send mail to <sales@vocitec.com>. Our Internet licensing solution provides even more flexibility for dynamic growth on your side and more continuous metering/measurement of usage.

Voci would be happy to see your usage to 1B hours of volume, and would be happy to support that growth. 1 billion hours is a large solution that requires significantly more than 1 server (roughly 1000+ servers). One server can handle a maximum of 800,000 to 1,000,000 hours per year, but our ASR solution stacks/scales very smoothly and easily. To get a rough estimate of the number of servers that you may need, divide your total needed volume by 800,000. Remember that this level of per-server usage means that you will be running all machines flat out. Most traffic loads are more spikey and the actual load curve and SLA requirements should be considered when sizing.

5.9. Can Voci Technologies, Inc. support 10,000 call center agents?

Voci Technologies, Inc. can provide an architecture to satisfy this using our out-of-the-box components, but it wouldn't be a simple 1-2 system installation.

For a rough frame of reference, 10,000 concurrent agents or interactions/calls would require 70-100 ASR Engine servers as well as a number of servers to support the V-Spark Analytics Engine (in order to avoid

any backlog/queuing). Voci currently benchmarks a typical single V-Spark (analytics-only) server as able to handle around 1000X of incoming audio. To optimally support the requested level of performance would probably mean both segregating sites/regions and/or offloading central components of the Analytics Engine (such as Elasticsearch and MySQL) to separate servers or clusters.

It would be highly unlikely for any company to begin a 50-100 server project without a reasonable and phased ramp-up plan. Voci will be glad to discuss possible architectures, but specifics will depend on the actual requirements.

For more information, send email that describes the scenario you are trying to address to <support@vocitec.com>.

5.10. Why does the API require zip files for transcription?

Most V-Spark customers tend to submit companion metadata files with the audio files that they are transcribing, and it often doesn't make sense to upload audio without metadata to provide additional information about that audio. Combining both of these into a zip file helps keep them together under a single trackable `request.id`. Some workarounds are available for cases in which you must upload them separately, but the standard customer usage pattern (often referred to as a *flow*) of using a zip file ensures that both arrive together.

You can also bundle multiple calls together into one zip archive. This can provide efficiency gains in some customer environments, such as those where uploading files takes a long time, but is certainly not necessary.

Using zip compression saves bandwidth and helps ensure faster uploads/transactions. The standard PCM/WAVE audio format is a pretty "heavy" format because its audio files are very large, as are those which are encoded using other audio containers/codecs. Any compression that you can do without losing audio quality will help in most cases.

Chapter 6. Folder Creation/ Maintenance

6.1. Which language model should I pick?

Voci provides three out-of-the-box language models: Call Center, Survey, and Voicemail:

- The **Call Center** model is designed specifically for call center conversations between an agent and client discussing products, services, and customer accounts.
- The **Survey** model is designed for after-call surveys where customers answer questions about their customer service experience.
- The **Voicemail** model is designed for voicemail messages. The Voicemail model therefore has a larger vocabulary and is a good choice for audio that doesn't fall into either of the previous categories.

See [Section 3.1, “What types of language models are available?”](#) for information about other language models that are available by request.

6.2. How do I start a local callback server for testing?

To start a simple version of a local callback server for testing purposes, do the following:

1. Use SSH to connect to the host on which V-Spark is installed
2. Use the `cd` command to change to a directory that you are able to write to:

```
cd /tmp
```

3. Check to see if you are using Voci Python by issuing the following command:

```
which python
```

This command should return `/opt/voci/bin/python`. If it does not and you want to make sure that you are using Voci's version of Python, add the directory where V-Spark is installed to your `PATH` by using a command like the following:

```
export PATH=DIRECTORY:${PATH}
```

4. Execute the following command to start up a callback server to listen for files:

```
python /opt/voci/vspark/examples/receiveFiles.py
```

If your session is using a different version of Python by default, issue the command:

```
/opt/voci/bin/python /opt/voci/vspark/examples/receiveFiles.py
```

When the program starts, it displays the message `Serving at: http://localhost:8000`. When you specify `localhost:8000` as the target callback URL, you will see results of files processed in the current session window.

5. For testing purposes, when you define a folder within V-Spark, you can define a callback method pointing to the server you just started at `localhost:8000/tmp` or any other path that you have write access to.

6.3. What are the consequences of moving results after a period of time?

If you remove a folder at any time, the only way to recreate it is to create it again and reprocess the audio that it originally contained.

If you are focused on saving storage space, the audio files are big, but the JSON transcripts are small. If you simply want to be able to continue to search transcripts, you can remove the audio and leave the JSON files. This will enable you to still do all your searches, but you will not be able to play back the file.



Note

If you remove files from folders, the references to the files stay in the V-Spark UI. Voci has a script that can be used to remove results on a schedule you determine. If you want to use such a script, contact Voci at support@vocitec.com.

Depending on the Voci solution that you are using, you may or may not have access to the locations containing the data. The default location is a symbolic link at `/var/lib/vspark/audio` or `/var/lib/vspark/json` which typically points to `/home/voci/vspark/database/audio` or `/home/voci/vspark/database/json`, respectively. If you have access to these links and have greater amounts of storage on some partition other than `/home`, you may want to change these links to point to directories that you create on that partition.

6.4. What does "Automatically send transcript (JSON) files to a URL" mean?

This option enables you to define a callback, which receives the entire JSON file output (including the transcript elements). The callback flow will not re-send new/updated JSON when application reprocessing takes place. If you download audio, transcripts, or text manually from the **File Details** page, you will get updated and correct `app_data/scores`, but if you do a bulk download from the folder settings, it will be the same as the original callback.

Chapter 7. The Audio Evaluator

7.1. How do I use the Audio Evaluator?

To use the Audio Evaluator, select Will My Audio Work? from the Help dropdown menu. On the screen that appears, select "Choose File" and select a single audio file to upload at a time.

7.2. Who can see the audio that I test using the Audio Evaluator?

The Audio Evaluator is completely secure - Voci employees do not have any way to view the files that you test using the Audio Evaluator and the files are not saved anywhere.

7.3. How do I know how many hours of audio I will have?

In many cases, knowing the type of audio that you will be processing doesn't really help you calculate space requirements. For example, knowing the audio is "wav" isn't helpful. There are on the order of 100 different encoding types that use the RIFF header standard. If the audio has a RIFF header, it's "wav".

The simplest method uses the per-channel bit rate and number of channels, or the overall bit rate (which already takes number of channels into account). The bit rate metric is a universal way of measuring audio size requirements - all audio formats have a bit rate (or an average bit rate for variable bit rate encodings).

For example, let's say all of your wav files contain audio encoded as u-law (common in telephony). u-law uses a fixed rate of 64 Kbps/channel. This means that 1 second of audio on a single channel is represented by 64,000 bits.

Furthermore, let's say your audio is stereo u-law. This doubles the overall bit rate to 128 Kbps.

The size of your audio set in bits is represented by:

$$2.2 \text{ GB} * 1024^3 \text{ bytes/GB} * 8 \text{ bits / byte} = 18,897,856,102 \text{ bits}$$

The duration of your audio set in seconds is represented by:

$$18,897,856,102 \text{ bits} / (128,000 \text{ bits/second}) = 147,639.5 \text{ seconds}$$

You can now convert duration to hours:

$$147,639.5 \text{ seconds} / (3600 \text{ seconds / hour}) = 41 \text{ hours}$$

If the audio was mono rather than stereo, you would have 82 hours.

If the encoding is G.729A rather than u-law, then the bit rate per channel is 8 Kbps. Sticking with the mono case, now you have:

$$82 * (64/8) = 656 \text{ hours}$$

The [mediainfo](https://mediaarea.net/en/MediaInfo/Download) [https://mediaarea.net/en/MediaInfo/Download] tool is very useful for examining audio properties. It runs on Windows, Linux, and Mac OS, and will tell you the bit rate and number of channels. On Windows, it features right-click integration, so you can just right-click on an audio file and "mediainfo" will be an available option.

Chapter 8. Viewing/Searching Data in V-Spark

8.1. What kinds of search queries are supported?

V-Spark supports both Plain Text and Regular Expression search of speaker utterances. V-Spark users can also search according to file name, agent ID, and custom metadata fields if desired. For more information about Plain Text search and custom metadata fields, please see the "*V-Spark Management Guide*". For more information about Regular Expression search, please see the "*V-Spark Regular Expression Guide*".

8.2. I don't see the audio that I just uploaded. Where is it?

While V-Spark can process audio in up to 100x real time, keep in mind that the V-Spark Demo server may be busy processing files that other clients have uploaded. As long as you see the notification that your file(s) were successfully uploaded, you should be able to view your files within the next couple of hours on the Dashboard. If you do not see your audio files within 24 hours, contact your Voci sales representative.

8.3. I uploaded audio a while ago, but I don't see it anymore. What happened?

If there is data in the last 31 days, V-Spark will display the current month on the graph. If the most recent data is older than 31 days, the graph will display that calendar month by default. You can view older transcripts by changing the date range in the Files View section of the Dashboard.

8.4. I just created a new application - why don't I see my data?

Unless you selected the option to reprocess application scores, your application will only display scores for files that you have uploaded since the time you linked the application to a folder. You can choose to reprocess application scores by navigating to Settings/Applications and selecting the green reprocessing button to the right of your application.

8.5. I just reprocessed audio for an application - why don't I see my data?

Application scores will be temporarily unavailable while the V-Spark Demo server reprocesses your data. Once reprocessing has been completed, the new application scores will be available to view. You can

monitor the progress of reprocessing by selecting the Queue icon to the right of the "Applications" heading on the Settings/Applications page.

8.6. Why are parts of the transcript red/green?

These sections indicate the emotion associated with that speaker turn/phrase. Emotion scores are calculated based on the tone of the speaker's voice, as well as word choice. Red and green highlighting of speaker turns or bolding of speaker phrases refer to Negative and Positive emotion, respectively. When speaker turns are highlighted gray, this means that the Positive and Negative portions of the speaker turn cancel each other out, and can be considered Neutral.

8.7. The display in V-Spark is not showing the correct information

First, try clearing the cache in your browser and retry your query. If this does not solve the problem, create a ticket to Voci Technologies, Inc. Support group that describes the problem. You can create a ticket by emailing <support@vocitec.com>.

8.8. Why is the search score always 1?

The score shown in the right-hand column on the V-Spark Dashboard is the search relevance ranking score. If you have not specified any search tags, the score will always be 1.000. With search tags specified, the score is calculated based on those tags and the number of files that you search matches, which helps you focus on the most relevant hits.

Chapter 9. Audio Transcripts

9.1. How do I download transcripts for the audio I uploaded?

You can download transcripts from the File Details page or from the Settings/Folders page. To download transcripts one at a time, navigate to the File Details page for the file in question. From there, you can choose to download the JSON or Text version of the transcript underneath the audio player.

9.2. What's the difference between JSON and Text transcripts?

JSON transcripts contain metadata at the file and utterance level in addition to the words spoken in the audio file. Text transcripts contain a readable, plain text transcript of the call, complete with speaker labels.

9.3. How do I interpret the JSON output of my transcript?

More information about JSON transcript output and the fields contained therein can be found in the JSON Output Guide located under the Help menu of V-Spark.

Chapter 10. Applications

10.1. Why can't I edit this application?

If the application you are trying to edit is within an Organization outside of your home Company, you do not have permissions to edit it. If you are unable to make changes to an application within your home Company, contact your Voci sales representative.

10.2. What pre-made application templates does Voci provide?

Voci provides various pre-made application templates as a starting point to customize your analysis. These templates are the following:

- The Agent Scorecard application scores calls according to agent performance for categories like Communication, Politeness, and Professionalism.
- The Call Categorization application categorizes calls based on call type and can be used to develop industry-specific call analysis.
- The Call Drivers application categorizes calls based on call type and can be used to develop industry-specific call driver analysis.
- The Customer Experience application covers all aspects of both positive and negative call center experiences, including customer sentiment, customer effort, and churn detection.

Appendix A. Voci Glossary

ACT	(See Average Call Time .)
Agent/Client Voice Clarity	A measure of the speech recognizer's confidence in its transcription accuracy. Higher clarity scores indicate that the call was transcribed with higher accuracy.
AHT	(See Average Handle Time .)
Application/Scorecard	A feature of V-Spark that allows users to create custom analysis of calls by entering sets of search queries and filters to group calls into upper-level categories, lower-level categories, and leaf-level categories. Leaf-level categories have no lower-level categories of their own.
ASR	(See Automatic Speech Recognition .)
Audio Source	Refers to the number of channels in the original audio file. Mono indicates that the audio file has only one channel, while Stereo indicates that it has two channels.
Automatic Speech Recognition	Technologies that leverage the independent, computer-driven transcription of spoken language into readable text (often known as <i>Speech to Text</i> or <i>STT</i>), enabling users to use speech rather than a keyboard as the source of input data for a computer system
Average Call Time	The total amount of time that a customer is engaged with an agent, including call transfers, hold time, talk time, and so on.
Average Handle Time	A call center metric that represents the total amount of work related to customer/agent transactions, including the Average Talk Time (ATT), hold time, and any administrative tasks that must be done for per-call record-keeping (often referred to as After Call Work, ACW), divided by the number of calls that were handled. This metric is used to summarize and analyze agent performance.
Call Silence Time	The percentage of the call that was silent, i.e. when no speakers were talking.
Diarization	The process of splitting mono audio into 2 channels, according to speaker. Higher diarization scores indicate that the speech recognizer separated speaker phrases more accurately.
JSON	Short for JavaScript Object Notation. Allows for information to be stored in a hierarchical fashion within a single document.
JSON Transcript	A transcript and associated metadata produced from an audio file, in JSON format. The Voci JSON format is documented in the <i>Voci JSON Output Format Guide</i> .
Overtalk	Refers to portions of conversation where speakers talked over one another/at the same time. Overtalk scores measure the percentage of the call that was overtalk.

Purify	Refers to V-Spark's ability to automatically redact sensitive PCI (Payment Card Industry) numbers from your calls in order to protect cardholder data. This information can be redacted from just the call audio, just the call transcript, or from both audio and transcript.
Regex	Short for <i>regular expression</i> . Allows users to perform more complex searches using variables to stand in for characters, words, and character patterns.
RIFF	An acronym for <i>Resource Interchange File Format</i> . A generic file container format for storing data in tagged chunks. It is primarily used to store multimedia such as sound and video.
Stream	A sequence of computer data that is made available over time. This term is commonly used in the context of <i>streaming audio</i> or <i>streaming video</i> which you can listen to as the data is continually being delivered to the device that you are using.
Tags	Name/value pairs that are specified when an audio stream is sent to transcription, and which are used to modify the way in which audio is transcribed.
TAT	(See Turn-Around Time .)
Turn-Around Time	The amount of time between submitting an audio file for transcription and receiving the transcribed text.
Utterance	An uninterrupted chain of spoken language by a single speaker.
VAD	(See Voice Activity Detection .)
Voci JSON	JSON output produced by one of Voci's products that typically contains all or some of a transcript produced from an audio file, plus metadata. The Voci JSON format is documented in the <i>Voci JSON Output Format Guide</i> .
Voice Activity Detection	An early portion of the Automatic Speech Recognition process, voice activity detection identifies the portions of an audio file that correspond to speech. This improves performance by not spending processing time attempting to decode audio that is not relevant to speech, such as music, background noise, and so on.