

Release notes

V-Spark 4.4.1-1

Medallia Voci

Contents

Release notes	7
V-Spark 4.4	8
4.4.0.....	9
V-Spark 4.3	11
4.3.3.....	12
4.3.2.....	13
4.3.1.....	14
4.3.0.....	17
V-Spark 4.2	20
Fixes in V-Spark 4.2	24
V-Spark 4.2.0 Fixes	24
Known Issues in V-Spark 4.2	25
V-Spark 4.1	26
Fixes in V-Spark 4.1	28

V-Spark 4.1.0 Fixes	28
Known Issues in V-Spark 4.1	29
Maintenance Updates Since the Initial 4.1 Release.....	29
V-Spark 4.1.0-2 Updates	29
V-Spark 4.0	30
Announcements with V-Spark 4.0.....	36
Fixes in V-Spark 4.0	40
V-Spark 4.0.2 Fixes	40
V-Spark 4.0.1 Fixes	41
Maintenance Updates Since the Initial 4.0 Release.....	42
V-Spark 4.0.2 Updates	43
V-Spark 4.0.1-10 Updates	45
V-Spark 4.0.1-9 Updates	46
V-Spark 4.0.1-8 Updates	47
V-Spark 4.0.1-7 Updates	47

V-Spark 4.0.1-6 Updates	48
V-Spark 4.0.1-5 Updates	48
Known Issues in V-Spark 4.0	49
Previous versions	50
V-Spark 3.5.0.....	51
V-Spark 3.5.0 Release Notes	52
Fixes in V-Spark 3.5.0	55
Maintenance Updates Since the Initial 3.5.0 Release.....	56
Announcements with 3.5.0	57
Known Issues in 3.5.0	58
V-Spark 3.4.3.....	60
V-Spark 3.4.3 Release Notes	61
Fixes in V-Spark 3.4.3	64
Maintenance Updates Since the Initial 3.4.3 Release.....	66
Announcements with 3.4.3	67

Known Issues in 3.4.3	68
V-Spark 3.4.2.....	70
V-Spark 3.4.2 Release Notes	71
Fixes in 3.4.2	73
Maintenance Fixes Since the Initial 3.4.2 Release.....	75
Announcements with 3.4.2	76
Known Issues in 3.4.2	77
V-Spark 3.4.1	79
V-Spark 3.4.1 Release Notes	79
Maintenance Fixes Since the Initial 3.4.1 Release.....	82
Announcements with 3.4.1	83
Known Issues in 3.4.1	84
V-Spark 3.4.0.....	85
V-Spark 3.3.1	86
V-Spark 3.3.1 Release Notes	87

Fixes in 3.3.1	88
Known Issues in 3.3.1	89
V-Spark 3.3.0.....	90
V-Spark 3.2.1.....	91
V-Spark3.2.1 Release Notes	92
V-Spark 3.2.....	94
V-Spark 3.2 Release Notes	95
Known Issues in 3.2	97
V-Spark 3.1.....	98
V-Spark 3.0.0.....	99
V-Spark 3.0.0 Release Notes	100
Known Issues in 3.0	106
V-Spark 2.2.5.....	107
V-Spark 2.2.4.....	108
V-Spark 2.2.3 Release Notes	109

V-Spark 2.2.0 Release Notes	110
Third-party service QA tests	112

Release notes

V-Spark version 4.4 / October 2022

- Better CSV report generation, metadata handling, and user experience

V-Spark version 4.3.3 / July 2022

- Better sorting, logging, and error handling

V-Spark version 4.3.2 / May 2022

- Better sorting in exported CSV reports

V-Spark version 4.3.1 / January 2022

- Example and reference system configuration files
- Better character handling in CSV reports

V-Spark version 4.3.0 / December 2021

- Added custom metadata to CSV reports
- Better tagging, logging, indexing, and queuing
- Better database handling and package organization

V-Spark 4.4

V-Spark 4.4 is a major release with enhancements to CSV report generation and custom metadata value handling, user experience improvements, and security and bug fixes.

IMPORTANT:

System administrators must run the `vspark-admin` script with the `core-update` parameter on any V-Spark system that has been upgraded from a version earlier than 4.3.1. Systems upgraded from V-Spark 4.3.1 or later do not need `core-update`.

Refer to the `core-update` parameter in the V-Spark 4.0 release notes for more information about this script.

4.4.0

1. CSV reports may now include an unlimited number of entries, and the configurable maximum limit has been removed.

With this change, users initialize report generation using the UI, but if the report is not immediately available, V-Spark continues to generate the report in the background. V-Spark emails the requesting user when generation is complete and the report is ready to download. Reports generated with a relatively large number of entries may take a considerable amount of time to complete.

Previously, the number of entries in CSV reports was limited to a maximum defined with the `export_filelist_limit` system configuration option. This option has been removed.

2. Renamed and clarified expected values for the **Servers** field used when connecting a company entity to an automatic speech recognition (ASR) engine. The **Servers** field is now named **ASR API URL**. Examples in the field now show the format of the required URL, and UI text explains to use a semicolon (;) to separate multiple URLs.
3. Added the new `restrict_application_editor` system configuration file option (`00base.cfg`) to prohibit users with only **View** permissions from seeing application configurations. If `restrict_application_editor = on` , only users with edit permissions for the associated organization are able to open the application editor from the dashboard. This new configuration does not change any existing behavior related to accessing the application editor from the **Settings** menu.
4. Addressed cross-site scripting and other security vulnerabilities.

Fixes in 4.4.0

1. Resolved an issue with the transcript view's highlight navigation buttons (⬆ and ⬇) on the **File Details** page. The buttons now correctly cycle through multiple highlighted application matches. This resolves a regression first introduced with V-Spark version 4.1.
2. Resolved a page routing issue that could sometimes cause certain valid URLs to be inaccessible and return a 404 error. Previously, attempting to access certain URLs resulted in a 404 error. These URLs now properly load.

3. Resolved an issue with **NaN**, **+Infinity**, and **-Infinity** values for custom metadata fields. These values are now converted and stored as strings.

Previously, when these values were submitted in JSON metadata, that JSON was accepted by the system but could result in records that could not be viewed on the **File Details** page. Additionally, those metadata values could not be changed with the API.

On the **File Details** page, hyphens are displayed as hash symbols (#), so **-Infinity** values appear as **#Infinity**.

Fixes in 4.4.1

1. Resolved an issue with exported CSV reports. Previously, generating a CSV report from the dashboard with **All** folders selected returned a report with incorrect values for `folder` and `url` column data. CSV report `folder` and `url` columns now contain the correct values when exporting all folders.

V-Spark 4.3

V-Spark 4.3 includes security updates, bug fixes, and user experience improvements. The following release notes contain fixes from multiple iterations of 4.3 because these iterations were released in close proximity. The most recent updates are listed first.

IMPORTANT:

System administrators must run the `vspark-admin` script with the `core-update` parameter on any system that has been upgraded to V-Spark 4.3.1 from an earlier version. Systems upgraded to 4.3.2 and 4.3.3 from 4.3.1 do not require a `core-update`.

Refer to the `core-update` row in the V-Spark 4.0 release notes for more information about this script.

4.3.3

1. **Updated the retry logic for handling database errors.** V-Spark now retries any database error with an error code between 2000 and 2099 (inclusive), along with the codes 9001, 9002, 9003, and 9006. Previously, V-Spark retried only a more limited set of 2000-level errors.
2. **Improved logging functionality.** Log messages generated when detecting language model availability now include a URL for the request that generated the message.

Fixes in 4.3.3

1. **Resolved an issue with default sort for search results in the Files View.** Search results once again display in descending order of **Time** by default. This fix resolves a regression first introduced with version 4.3.0. Manual sorting was unaffected by this regression.
2. **Resolved an issue for systems hosted on AWS that caused some folder data to be deleted erroneously.** Previously, deleting a folder could cause the JSON and mp3 data in long-term storage for other folders starting with the deleted folder's name to also be deleted.

For example, for systems hosted on AWS, long-term storage data associated with the folder `foobar` was erroneously deleted along with data in the deleted folder `foo`. Folder configuration and Elasticsearch and database records were not affected by this issue.

4.3.2

1. Updated the sorting logic for columns in CSV exports generated from the **Files View** list. Columns in reports created with the **Files View** are now sorted consistently and independently of the order shown in the UI.

Columns are sorted alphabetically, from left to right, in two groups. The first group contains general columns like `agent_id`, `duration`, and other audio record attributes. The second group contains application-related columns, the names of which begin with the `app.` prefix.

This change ensures that column order will be consistent for every report generated, including paginated reports.

4.3.1

1. **Updated the Elasticsearch package included with V-Spark to account for the security vulnerability associated with the Log4j security flaw.** The `voci-config-elasticsearch` package included with V-Spark 4.3.1, and initially released with version 4.3.0-2, allows users to upgrade to any minor version of Elasticsearch 7 with the caveat that versions newer than 7.16.2-1 have not been tested for use with V-Spark.

The full list of service dependencies tested with various V-Spark versions is available at the following link: [Third-party service QA tests](#). The versions tested with this release are included with these release notes.

For more information about V-Spark and the Log4j vulnerability, please refer to the following statement from the Voci support team:

- a. Voci V-Spark **does not** use Log4j directly, so we are not planning any updates to V-Spark software for this issue.
- b. V-Spark **does** have an operational dependency on Elasticsearch. The Elasticsearch team has determined that there is no remote execution risk in version 7.x releases, and that a minor issue with potential DNS-based information exposure is easily mitigated with a basic configuration update. Out of an abundance of caution, Voci advises that the mitigation be performed.
- c. The best solution to this issue is an upgrade to Elasticsearch version 7.16.2-1 or greater, which users can perform at any time convenient to them and their teams. Voci does not anticipate any issues with using Elasticsearch versions greater than 7.16.2-1 with V-Spark, but those versions have not been tested by the V-Spark Engineering team.

More information about the risks and remediation associated with Log4j and Elasticsearch can be found on [Elasticsearch's website](#).

2. **Increased the timeout threshold for jobs in the report worker queue.** Jobs now wait in the report worker queue for a maximum of 1 hour before timing out and being requeued.
3. **Added functionality to escape certain characters in the field names for exported CSV reports.** Custom metadata fields in exported CSV reports whose names begin with `=`, `+`, `-`, or `@` are prefixed with an apostrophe `'` to prevent spreadsheet applications from interpreting these labels as formulas.

4. **Improved error handling when Elasticsearch is unavailable.** Transcript worker processes now pick up new jobs only when these processes detect that Elasticsearch is up and running. Previously, jobs accepted while Elasticsearch was unavailable were picked up and immediately placed in the errors folder.
5. **Added example and reference files to V-Spark's system configuration settings directory.** The four new example files show system configurations for various use cases and are located in the `config/vspark.config.d.examples` directory. The reference copies of default configuration files are located in the `config/vspark.config.d.examples/base.defaults` directory.
6. **Improved V-Spark security to help protect against remote code execution and other vulnerabilities.**

Fixes in 4.3.1

1. **Resolved an issue with residual memory usage associated with report worker processes that caused them to use an abnormally large amount of memory.** These processes now terminate and restart when they have exceeded a configurable memory threshold after finishing a report job. The new system configuration option `report_worker_idle_memory_threshold` sets the amount of memory at which report worker processes terminate and restart, takes an integer value expressed in MB, and has a default value of 200.

Report worker process memory usage grows linearly as a function of the number of calls included in daily and monthly reports at an approximate rate of 10Kb per call. Although it is thus still possible for report worker processes to consume a large amount of memory, with this update, memory used by report worker processes is released back to the operating system as soon as reports are generated.

Report worker jobs that terminate and restart generate a `WARNING` -level message in `backendWorker.log`.

Known issues in 4.3.1

1. **Application changes may not display in real time when made by another user from a different host.** Although application editing works, users editing an application simultaneously from different hosts must refresh the **Application Editor** to see changes made by another user. This issue does not typically occur when both users are being served by the same host.


4.3.1 Tests

The following table shows the versions of third-party dependencies used to test V-Spark version 4.3.1.

V-Spark 4.3.1 Dependency Versions Tested

Dependency	Version(s) Tested
Elasticsearch	7.16.2-1
Database	MariaDB 5.5 (EPEL repository)
Redis	3.2 (EPEL repository)

4.3.0

1. Added functionality to display and export columns for custom metadata fields and their values using the **Show/Hide Columns** dropdown in the  **Dashboard Files View**. The new **Custom metadata** section of the dropdown list includes all metadata fields associated with the folder currently selected in the dashboard. If the dashboard is set to display all folders, the dropdown includes all custom metadata fields for every folder the user has permission to view.

Enabling custom metadata fields using the dropdown adds a column for each enabled field to the table of search results. If a record in the list of search results has a value for the enabled field, that value will display on the search result line in the cell associated with the custom metadata field column; otherwise, the cell is blank. Enabled fields and their values are also included in search result CSV exports.

As a part of this change, the **✕ Delete** icon for search results displayed in the files view has been moved to the left side of each line in the results list to improve user experience. Previously, the icon appeared on the right side of each result line.

2. Improved user experience when tagging search results in bulk. The dashboard now displays a **Saving...** message while adding tags to multiple search results. Previously, when tagging a large number of files, the page would sometimes display a **Successfully Added Tags** message before the new tags were visible in search results.
3. Increased the timeout threshold for jobs in the dashboard worker queue. Jobs now wait in the dashboard worker queue for a maximum of 4 hours before timing out and being requeued.
4. Added log messaging for all requeued jobs. Requeued jobs generate a **WARNING** -level message with the date and time, type of worker process, and other details in `backendWorker.log` with the format used in the following example:

```
2021-12-01 18:01:00.108 TranscriptWorker 28387 WARNING Job Requeued: Found inflight job
(transcripts_f145c2a1-2a91-4cef-a245-ac4521e8c17b) with an expired TTL --
<requestid=06babe23-2ad5-44b2-8153-4ceed1f76cd8>
```

5. Changed the format for temporary Elasticsearch index names. Previously, V-Spark assigned the prefix `.vspark-temp-` to temporary

Elasticsearch indices created for application scoring. The default value for this prefix has been changed to `@vstemp` to account for changes to index name requirements in the next major version of Elasticsearch.

In addition to the prefix change, temporary index names now include date, time, and detailed request components. The new complete format for temporary index names is `prefix_YYYYMMDD_HHMMSSz_requestid_orgshort_tid` and includes the following components:

- `prefix` – The prefix configured in V-Spark, which is `@vstemp` by default.
- `YYYYMMDD` – The year, month, and day on which the index was created.
- `HHMMSSz` – The hour, minute, second, and time zone for the index, expressed in UTC.
- `requestid` – The unique identifier for the request that generated the transcript.
- `orgshort` – The short name of the organization associated with the transcript.
- `tid` – The unique `transcriptID` identifier for the transcript.

IMPORTANT: Although a `core-update` is not required to implement the temporary index change, a full system restart of all backend nodes is required. Otherwise, the system will be in an inconsistent state with some temporary indices using the deprecated format and some using the new format.

6. Addressed an issue with startup verification of database names that contain both lowercase and uppercase characters. Previously, verification would sometimes fail when checking the names of database servers with `lower_case_table_names` set to 1, such as is common on Azure-hosted databases. For more information, refer to the external [Azure database documentation](#).
7. The third-party Python and NodeJS packages, and the optional third-party Elasticsearch and database packages, used with V-Spark are now offered as independent packages in the Voci repository.
 - Python – `voci-spark-pylib`

- NodeJS – `voci-spark-nodelib`
- Elasticsearch – `voci-config-elasticsearch`
- Databases – `voci-config-database`

8. Improved V-Spark security to help protect against configuration request forgery and remote command execution.

Fixes in 4.3.0

1. Updated the **vspark-admin check-health** maintenance script to account for changes to Elasticsearch and Redis security configurations made with V-Spark version 4.2.0. This resolves the known issue described in the 4.2.0 release notes.
2. Addressed an issue that prevented application score data from exporting correctly in certain conditions. Previously, exporting CSV reports from the dashboard with **All** folders selected in the dashboard dropdown would generate a file with only application names and no score data. Those reports now contain the correct data.
3. Addressed issues with both the **/sysinfo** endpoint and the system status page in the V-Spark UI that caused Elasticsearch to leave too many TCP connections open when the system is under heavy load. This fix resolves a regression first introduced with V-Spark 4.2.0 initially released on 2021-09-23.
4. Addressed a rare issue with the Job Manager component's drive failure recovery and retry logic where a Job Manager process could fail to recover from operating-system-level drive failures when moving certain files.

V-Spark 4.2

V-Spark 4.2 is a major release that includes new API endpoints, service dependency upgrades and security enhancements, and bug fixes.

IMPORTANT:

System administrators must run the `vspark-admin` script with the `core-update` parameter on any system that has been upgraded to V-Spark 4.2 from an earlier version.

NOTE: AWS recently rebranded their managed AWS Elasticsearch service to AWS OpenSearch, which is derived from an earlier version of Elasticsearch that is not supported by V-Spark. As of version 4.1, V-Spark does not support AWS OpenSearch for its Elasticsearch requirements.

1. Added the `/appmatches` endpoint, used to retrieve phrase matches from applications. The `/appmatches` endpoint returns the scores and phrase matches for all or specified applications and application categories for the provided `transcriptID`. To limit results, submit the names of individual applications and categories as JSON-formatted data with the request.
2. Added the `/metadata` endpoint, used to add, modify, or delete metadata values for previously processed transcription results. The `/metadata` endpoint uses the `PATCH` method to add or modify metadata values, uses the `DELETE` method to delete them, and requires JSON-formatted data specifying metadata changes to be provided with each request.
3. Added the `lastmodifiedrange` parameter to the `/search` endpoint. The `lastmodifiedrange` parameter filters search results by a specified date range. Date ranges may be expressed using any combination of years, months, days, hours, minutes, and seconds.
4. Added the `last_modified` field to the Elasticsearch record associated with each newly-created transcript. The `last_modified` field stores the date and time at which the Elasticsearch record was last changed. Data for the `last_modified` field is stored using Coordinated Universal Time (UTC). The `last_modified` field is not included in JSON records in long-term storage, but it is always included in JSON output that is dynamically generated. Dynamically generated output includes, for example, JSON transcripts downloaded

using the **Files View**.

5. Added a new folder setting that rejects duplicate file uploads based on their filenames. This setting is off by default and must be enabled for each folder that uses it. As of version 4.2.0, this setting can only be configured for folders using the UI.

When a folder has the deduplication setting enabled, that folder will reject file uploads in the following scenarios:

- A file is uploaded with the same name as a previously uploaded file.
- A zip file contains a file with the same name as a previously uploaded file.
- A zip file contains two or more files with the same name.

Note that the entire zip is rejected when a duplicate file is detected, and duplicate file rejection for zip files nested inside other zip files is not supported. When folder-level deduplication causes a file to be rejected, V-Spark generates a `WARNING` -level message in `server.log` and the **Activity Log**.

When duplicate file rejection is enabled for a folder, audio may still be reloaded for reprocessing, but reloaded MP3 audio data will be rejected when the original audio file was also an MP3. A previously processed file that was originally uploaded in any other format is accepted. Reloading JSON data is not affected by the folder-level duplicate rejection setting.

6. Added support for SSL/TLS and password authentication when V-Spark connects to Elasticsearch. These features are disabled by default. The following system configuration options have been added:

- `es_user` – Specify the username defined in Elasticsearch. The default value is `elastic`.
- `es_pwd` – Specify the password defined in Elasticsearch.
- `es_ssl` – Set to `on` to enable SSL/TLS connections.
- `es_ssl_ca` – Specify the path to the CA certificate bundle to be used.
- `es_ssl_cert` – Specify the path to the client certificate to be used.

- `es_ssl_key` – Specify the path to the client certificate key to be used.

Note that setting these configuration options does not secure the connection; they enable communication with a secured Elasticsearch implementation. For information about configuring security for Elasticsearch, refer to the following topics in the external Elasticsearch documentation:

- [🔗 TLS/SSL configuration](#)
- [🔗 Set up minimal security for Elasticsearch](#)
- [🔗 Encrypt internode communications with TLS](#)

When securing Elasticsearch with SSL/TLS in a clustered implementation, if nodes in the Elasticsearch cluster access V-Spark instances locally, two additional configuration file options, `elasticsearch_hosts` and `app_score_es_hosts`, should be defined to match the domain used to generate the certificate.

The `app_score_es_hosts` option is not typically included in V-Spark configuration files, in which case it should be added manually. `app_score_es_hosts` is set to `http://localhost:9200` by default. If its value has not been changed from the default, `app_score_es_hosts` must be set to the same value as `elasticsearch_hosts` in order for V-Spark to communicate with a secured Elasticsearch service.

If the `app_score_es_hosts` and `elasticsearch_hosts` configuration file options are not set to match the certificate domain, starting V-Spark produces a warning banner in the UI with the following message: "The Elasticsearch host(s) configured on {hostname} (from parameter 'app_score_es_hosts') are not running; using the cluster for app scoring instead." V-Spark will eventually connect to the cluster, but the banner displays until the configuration is corrected and V-Spark successfully uploads audio.

If encrypted internode communication is the only required security configuration, and all nodes can be accessed locally from their respective V-Spark instances, there is no need to set `es_ssl` to `on`, and the `app_score_es_hosts` and `elasticsearch_hosts` configuration options do not need to be configured from the default `http://localhost:9200` value.

7. Added support for SSL/TLS and password authentication when V-Spark connects to Redis. These features are disabled by default. The

following system configuration options have been added:

- `redis_ssl` – Set to `on` to enable SSL/TLS connections.
- `redis_auth_token` – Specify the Redis password to be used.
- `redis_ssl_ca` – Specify the path to the CA certificate bundle to be used.
- `redis_ssl_cert` – Specify the path to the client certificate to be used.
- `redis_ssl_key` – Specify the path to the client certificate key to be used.

Note that setting these configurations does not secure the connection; they enable communication with a secured Redis implementation. For information about configuring security for Redis, refer to the external topics [Redis Security](#) and [TLS Support](#) in the Redis documentation.

8. Made improvements to the failure and recovery logic used when the connection to Redis Sentinel fails. Clear error messages are now properly generated, and connection retry timing now uses a geometric backoff.
9. Upgraded the Lodash third-party library dependency to version 4.17.21 to address possible security concerns.
10. Improved the data validation behavior for data fields in the company creation and modification dialogs. Previously, unrelated user events would trigger validation on the **Servers** field before that field had been used.
11. Reclassified log entries generated as a result of SSO configuration errors as **ERROR** in `server.log` and the **Activity Log**. Previously, these entries were classified as `WARN` in system logs and `WARNING` in the **Activity Log**.
12. Changed the labels of the **Purify Text** and **Purify Audio** options in the **Create New Folder** and **Modify Folder** dialogs. Purify Text is now labeled **PCI Redact Text**, and Purify Audio is now labeled **PCI Redact Audio**. The field names used for these options in JSON input and output for the `/config` API endpoint were not changed, but those field names will be updated with the next major V-Spark release.
13. Added startup checks to the hostname configuration setting. If the hostname is misconfigured, V-Spark will not start, and will log `ERROR`

-level messages in `backendWorker.log` and `server.log`.

Fixes in V-Spark 4.2

The following issues have been resolved in the V-Spark 4.2 release:

V-Spark 4.2.0 Fixes

1. Addressed issues that sometimes caused matched phrases to fail to correctly highlight in the **File Details** page when viewing application phrase matches.
2. Addressed an issue that caused MySQL uptime and usage metrics to fail to properly display on the **System Status** page. These values now display correctly.
3. Addressed an issue that caused category names with Unicode characters to fail to properly display in dropdown menus on the Application Dashboard and the Agents View.
4. Addressed an issue that caused an Elasticsearch security configuration error to be reported in `server.log` as a long string of hexadecimal characters. The message no longer appears in `server.log`, but will appear legibly in `elasticsearch.log`. Users must address Elasticsearch security warnings to eliminate the message altogether by configuring Elasticsearch to require username and password credentials.
5. Addressed an issue with input guards for custom metadata field names that erroneously allowed some reserved field names to be configured for folders.

Known Issues in V-Spark 4.2

1. The `vspark-admin check-health` script does not connect to Elasticsearch and Redis on V-Spark systems configured to use secure connections to those services.

V-Spark 4.1

V-Spark 4.1 is a major release that includes support for single sign-on (SSO) authentication, updated service dependencies, and numerous security enhancements. Many changes included with 4.1 build on architecture upgrades completed as part of the 4.0 release.

IMPORTANT:

System administrators must run the `vspark-admin` script with the `core-update` parameter on any system that has been upgraded to V-Spark 4.1 from an earlier version.

Refer to the `core-update` row in the V-Spark 4.0 release notes for more information about this script.

NOTE: AWS recently rebranded their managed AWS Elasticsearch service to AWS OpenSearch, which is derived from an earlier version of Elasticsearch that is not supported by V-Spark. As of version 4.1, V-Spark does not support AWS OpenSearch for its Elasticsearch requirements.

Major Changes in the 4.1 Release

1. **V-Spark now supports single sign-on (SSO) authentication** using the OpenID Connect (OIDC) protocol. SSO allows V-Spark user accounts to be managed and authenticated by a third-party identity provider (IDP). Note the following points about using SSO with V-Spark:
 - SSO capability is configured at the system and company levels, and companies that share a system may use distinct SSO configurations.
 - Once the system and company have been configured, user accounts that use SSO authentication can be created or requested. After an administrator verifies the account and assigns the user's permissions, the user may log in with the newly-created SSO account.
 - User accounts with username-password authentication cannot be converted to SSO accounts; accounts must be created with SSO enabled in order to use it.

- Some account functionality is different when the account is authenticated with SSO: an SSO user account's email and profile company can't be changed once specified during creation, and change and forgot password functionality is disabled.

Refer to the **Using Single Sign-On (SSO) with V-Spark** section of the **Management Guide** for more comprehensive information about configuring and using SSO with V-Spark.

2. **V-Spark now supports Elasticsearch version 7.13.** No breaking changes related to this upgrade were identified. V-Spark supports a range of Elasticsearch versions, from version 7.6 up to and including version 7.13. This range of supported versions is now more strictly constrained than in previous versions of V-Spark.

System administrators should note that previous versions of V-Spark may be able to run some unsupported versions of Elasticsearch, including versions older than 7.6. Additionally, systems that do not rely on the `vspark-all` package may use versions of Elasticsearch greater than 7.13, but these versions have not been tested for use with V-Spark.

3. **Added a configurable limit to the number of consecutive failed login attempts for user accounts.** Exceeding this attempt limit will disable the user account, and it cannot be used until reactivated by a system administrator. By default, the maximum number of attempts is 6. This limit may be changed using the new system configuration option `max_passwd_failed_attempts`. Specifying a value of -1 will disable this feature.

Events to indicate failed login attempts and accounts locked for exceeding the limit are logged in `server.log` and shown in the **Activity Log** as `INFO` and `WARNING` messages that include the IP address, username, and timestamp.

4. **Added log reporting for session cookie security mismatch errors.** An `ERROR` entry is now added to `server.log` whenever the V-Spark session was not served over a secure connection even though the `secure_cookie` system configuration setting was enabled.
5. **Addressed cross-site scripting vulnerabilities.**

Fixes in V-Spark 4.1

The following issues have been resolved in the V-Spark 4.1 release:

V-Spark 4.1.0 Fixes

1. Added recovery logic for deletion and report worker processes following a temporary database failure. When deletion or report workers temporarily lose contact with the database, their respective processes will now resume once the database connection is reestablished.
2. Addressed errors associated with custom database name configuration. Previously, specifying a value other than `vspark` for the database name in the system configuration settings would cause some processing flows to fail.
3. Addressed an issue with configuring multiple ASR servers for a folder. Previously, specifying multiple ASR servers in a semicolon-delineated list would cause an error when specifying the folder's language models and application templates.
4. Modified the descriptor for **WARN**-level system events. These events now appear with the label **WARNING** in the Activity Log. Previously, some **WARN** events would not be properly filtered in the UI.
5. Addressed an issue that prevented the **+ Add** button from displaying when creating a company in a new system. Previously, the button would not display on the **Organizations** and **Users** sections of the **Accounts Settings** page until the user manually refreshed the page.
6. Addressed an issue with application shutdown functionality that caused cleanup not to finish in some circumstances. Previously, rate limiting keys were not removed from Redis until they expired. All rate limiting Redis keys are now removed when V-Spark shuts down.

Known Issues in V-Spark 4.1

1. Application changes may not display in real time when made by another user from a different host. Although application editing works, users editing an application simultaneously from different hosts must refresh the **Application Editor** to see changes made by another user. This issue does not typically occur when both users are being served by the same host.

Maintenance Updates Since the Initial 4.1 Release

V-Spark 4.1.0-2 Updates

1. Addressed an issue that caused application reprocessing jobs to fail and to appear stuck in **Waiting...** status.

V-Spark 4.0

V-Spark 4.0 is a major release with significantly changed architecture, along with numerous improvements and bug fixes.

1. **V-Spark now runs on CentOS 7.** Changes to V-Spark's technical foundation are primarily related to the CentOS 7 upgrade and are the main drivers for this release.
2. **V-Spark services are now managed by `systemd`.** `SysVinit` has been phased out and replaced with `systemd` as part of the upgrade to CentOS 7.
 - a. This update removes `init.d` and intermediate scripts where possible, and introduces a hierarchy of individually managed services defined in unit files. The new service hierarchy can be represented as follows:

```
/usr/lib/systemd/system/vspark.service
```

```
- /usr/lib/systemd/system/vspark-front.service
```

```
- /usr/lib/systemd/system/vspark-back.service
```

```
- /usr/lib/systemd/system/vspark-jobmgr.service
```

```
- /usr/lib/systemd/system/vspark-scc.service
```

```
- /usr/lib/systemd/system/vspark-scc-broker.service
```

```
- /usr/lib/systemd/system/vspark-scc-state.service
```

- b. V-Spark services can be controlled as a group by using the top-level `vspark.service`, which acts as a wrapper around the others. To start this service, use the following command:

```
$ sudo systemctl start vspark
```

- c. The services started by `vspark.service` are specified in the `Wants` directive in the `/usr/lib/systemd/system/vspark.service` file. If `vspark.service` is enabled, it will be started automatically on boot and reboot, and any services in the `Wants` directive will also try to start. Any excluded services will not start.
- d. Each service can be started, stopped, and restarted separately as needed without invoking `vspark.service` using `systemctl`, as in these examples:

```
$ sudo systemctl start vspark-back
$ sudo systemctl start vspark-front
$ sudo systemctl start vspark-jobmgr
```

Other options for `systemctl vspark` include the following:

```
$ sudo systemctl restart vspark
$ sudo systemctl stop vspark
$ sudo systemctl status vspark
```

3. V-Spark now has a dedicated command line script for system administration. To use it, run `sudo vspark-admin OPTION` with any of the parameters in the following table.

Option parameters for `vspark-admin`

Parameter	Description
<code>status</code>	Show the status of all V-Spark services.
<code>status OPTION</code>	Show the status of a specific service. To specify that service, replace <code>OPTION</code> with one of these parameters: <code>front</code> , <code>back</code> , <code>jobmgr</code> , <code>sccluster</code>
<code>version</code>	Display the current version of V-Spark that is installed.
<code>core-check</code>	Shows whether or not service dependencies are in allowed version range.
<code>core-update</code>	Shows and applies available updates for V-Spark dependencies. To preview required changes to core schemas and data structures, invoke <code>core-update</code> as in this example: <pre>\$ sudo vspark-admin core-update</pre> To apply required changes to core schemas and data structures, invoke <code>core-update</code> with the commit parameter <code>-c</code> as in this example:

Parameter	Description
	<pre>\$ sudo vspark-admin core-update -c</pre>
<code>check-health</code>	Shows status information about service dependencies.
<code>show-config</code>	Display a list of every config setting (represented in value pairs) that V-Spark is using. NOTE: Any configuration changes made to a running version of V-Spark must be reloaded for changes to take effect.

For example, to check the status of V-Spark's front-end services, run the following command:

```
$ sudo vspark-admin status front
```

4. **V-Spark now attempts to reconnect to service dependencies automatically.** When certain required services become unavailable, V-Spark will attempt to reconnect to them in order to minimize service disruption. By default, there is no limit to the number of times V-Spark will attempt to reconnect.
5. **V-Spark now uses improved logic to handle connectivity issues with the license server.** When revalidation fails due to loss of network connectivity or other scenarios, V-Spark will respect the time to live (TTL) value associated with the requested license. As long as the TTL with the license server hasn't expired, the license will stay validated in order to minimize service interruptions.
6. **V-Spark now includes the `voci-spark-tools` package,** which contains two utilities—Datatool and Config Manager—that facilitate the transfer of installation data and configuration settings.
 - a. Datatool is a command-line tool for importing (`load`) and exporting (`backup`) audio and transcription data. For more details, see the `Datatool.MD` file in the `utils/datatool` directory.
 - b. Config Manager is a script for importing and exporting installation configuration settings. This script can be found in the `utils/` directory.
7. **Improved V-Spark security to help protect against remote code executions.** For more secure file handling, name validation rules apply to audio files and ZIP archives uploaded for transcription via the GUI and API, and also to files uploaded to the Audio Evaluator. Files inside a zip are not checked. This feature was implemented with release version 4.0.1-3. Default naming rules forbid these characters:
`#*<>:~/\|{}$!'`"=^`
Filename validation is enabled by default. To disable it, set the `filename_validation` system configuration setting to off. To define custom character requirements, specify a regular expression via `filename_validation_pattern`.
8. **Improved V-Spark security to help protect against SQL injections.**

9. V-Spark can now log which users have viewed the File Details page. To minimize excessive or unnecessary logging noise, this setting is off by default. To enable it, set the new system configuration option `audit_filedetails_pageviews` to on. When enabled, an `audit` entry is logged in the Activity Log and recorded in `server.log` as an `INFO` entry.
10. Links to HTML documentation in V-Spark's help menu now point to Voci's Online Help. This update deprecates the `release_note` configuration option because the HTML version of the release notes has moved from the installation to the website. Note that the `release_note_dl` configuration still exists because PDF versions of documentation are still bundled with the software.

Announcements with V-Spark 4.0

1. System architecture has significantly changed with the V-Spark 4.0 release. System administrators should take note of the new front-end libraries and other dependencies listed in the following table as they evaluate installation requirements.

Changes in V-Spark dependencies from version 3.5 to 4.0

Dependency	Version used with V-Spark 3.5	Version used with V-Spark 4.0
CentOS	6	7
Node	6	14
Elasticsearch	5.6	7.6
MySQL and MariaDB	MySQL 5.1	MariaDB 5.5 (EPEL repository)
Redis	3.2	3.2 (no changes, EPEL repository)
Moment.js	2.17.1	2.29.1
Bootstrap	3.2.0	3.4.1
jQuery	2.2.4	3.6.0

NOTE: Customers upgrading from a 3.4.3 or 3.5 V-Spark system running on CentOS 6 should contact customer support for recommended upgrade paths.

Preliminary testing shows that enhancements in the 4.0 release lead to a 10% average performance improvement during data ingestion over previous versions.

2. The location of some V-Spark log files has changed. Some logged information now goes into CentOS 7's `journald` and can be visualized using `journalctl` from the command line, as in the following example:

```
$ sudo journalctl --unit=vspark-jobmgr
```

This change keeps V-Spark log files consistent with `systemd` service handling best practices. V-Spark 4.0 log locations are listed in the following tables:

`/var/log/vspark/`

Original Location	New Location
<code>/var/log/vspark/front-err.log</code>	journal
<code>/var/log/vspark/back-err.log</code>	journal
<code>/var/log/vspark/backend_stdio.log</code>	No change
<code>/var/log/vspark/backendWorker.log</code>	No change

Original Location	New Location
<code>/var/log/vspark/license.log</code>	No change
<code>/var/log/vspark/search.log</code>	No change
<code>/var/log/vspark/server.log</code>	No change
<hr/>	
<code>/var/log/vocijobmgr/</code>	
Original Location	New Location
<code>/var/log/vocijobmgr/init.err</code>	journal
<code>/var/log/vocijobmgr/init.out</code>	journal
<hr/>	
<code>/var/lib/vspark/managers/{company-org-folder}/logs/</code>	
Original Location	New Location
*	No change

3. The **voci-spark-hdfs-lib** package is now optional, and it is no longer installed by default, as most installations don't require HDFS

support.

Fixes in V-Spark 4.0

The following issues have been resolved in the V-Spark 4.0 release.

V-Spark 4.0.2 Fixes

1. Action buttons will no longer appear enabled until the required form criteria are met. Previously, **Login** and **Upload** buttons would appear to be active even though using them would cause errors.
2. Resolved an issue with error and warning logging that caused some details to be discarded, or to be spread out across multiple lines. Log messages for most system errors and warnings now include more complete and consistent detail.

V-Spark 4.0.1 Fixes

1. Folder permissions were sometimes incorrectly set when using the API to create folders. This caused the API call to fail sporadically, and caused the folders to be hidden and unusable until the permissions were correctly set.
2. Dashboard performance when viewing individual folders has been improved. Note that only single-folder dashboard performance was affected by this issue.
3. V-Spark services now automatically restart during hardware boot and reboot.
4. Backend V-Spark services now properly stop when Redis is down. Prior to this change, stopping the backend service while Redis was down could lead to errors that required killing backend processes manually in order to restart them.
5. The status indicator of Folders' server status now consistently shows the red or green indicator independently of user role.
6. The call volume data display on the Overview Dashboard now behaves more consistently. Previously, page refreshes would show either the monthly view or the 31-day view at random. The dashboard now shows either the monthly or the 31-day view consistently across subsequent page loads, depending on the last selected.
7. Organizations created via the /config endpoint will now use the "US/Eastern" time zone if none is specified in the initial API call. Previously, no time zone would be assigned to the organization. This would eventually cause the UI to fail to render some pages, as certain pages expect time zone information to exist for organizations.

Maintenance Updates Since the Initial 4.0 Release

This section describes the bug fixes and other enhancements included in the latest version of the 4.0 release. Maintenance versions and associated changes are grouped by release version and listed most recent to earliest.

IMPORTANT: Any updates to V-Spark 4.x from versions earlier than 4.0.1-8 require running the `core-update` script before starting the upgraded installation. See the `core-update` row in release note 3 for instructions for running this script.

V-Spark 4.0.2 Updates

1. System architecture has significantly changed with the V-Spark 4.0.2 release. These front-end library updates are the primary driver for the 4.0.2 release. System administrators should take note of the new libraries listed in the following table:

Changes in V-Spark front-end libraries from version 4.0.1 to 4.0.2

Dependency	Version used with V-Spark 4.0.1	Version used with V-Spark 4.0.2
Node	10.19.0	14.15.0
Moment.js	2.17.1	2.29.1
Bootstrap	3.2.0	3.4.1
jQuery	2.2.4	3.6.0

2. Overhauled the layout of the **System Status** page the **System Status** page to emphasize the most significant points of data flow in the system. The following changes have been made:
 - System information has been divided into the **System Status** and **Processing & Performance** tabs.
 - The **System Status** tab displays host hardware and installation attributes, along with uptimes, status information, and usage for Elasticsearch, Redis, SQL, and local storage.
 - The **Processing & Performance** tab displays queue status and folder activity metrics. This information includes the following:
 - ASR and analytics queue and processing status for audio calls

- Time to ASR and transcript analysis completion (in seconds) over the last 15 minutes, the last hour, the last 4 hours, the last day, and the last week
 - Two lists of folders: one that displays folders by most calls submitted to the system, and another that displays folders by most calls in the ASR queue
 - A sortable table of folder activity that displays the number of files ingested, the duration of that folder's audio (in seconds), the number of files completed, and the last active date for all system folders that can be filtered by various time periods
3. **Improved response handling for systems using V-Cloud.** Eliminated extraneous 503 errors generated when communicating with V-Cloud to minimize unnecessary noise in Job Manager log files.
 4. **Updated the system password policy with more rigorous requirements.** The following policy changes have been made:
 - A user must provide the current password to change the account's email address.
 - A user may not change the account password to any of the previous four used.
 - Passwords must contain both letters and numbers.
 5. **V-Spark now supports the `secure` attribute for sensitive cookies.** To use cookies with the secure attribute, set the new system configuration option `secure_cookies` to on.

WARNING: Enabling `secure_cookies` on a non-HTTPS system will make the system unusable.
 6. **Applied additional security updates, including the following:**
 - Replaced feedback message on the **Forgot Password** page to obfuscate whether the provided email address exists in the system.
 - Addressed all cross-site scripting vulnerabilities in the system discovered up to this point.

V-Spark 4.0.1-10 Updates

1. Added new configuration file options to support encrypted TLS/SSL database connections. The following options have been added:
 - `db_ssl`
 - `db_ssl_ca`
 - `db_ssl_cert`
 - `db_ssl_key`
 - `db_ssl_enable_azure_connections`
2. Improved logic for handling Job Manager processes exhibiting continuous memory growth. V-Spark now waits a maximum of 10 minutes for in-progress jobs to finish before terminating and replacing a worker process whose memory usage exceeds the configured threshold.

V-Spark 4.0.1-9 Updates

1. Added folder activity metrics to the **System Status** page. Folder activity metrics may be limited to certain time ranges using the time selector. The following folder metric fields have been added:

Ingestion Volume

Displays the total number of audio and metadata files added to Elasticsearch and assigned a `transcriptID` in each listed folder. Files deleted after ingestion are not represented in this count.

Throughput

Displays the total audio time, measured in seconds, of all files in each listed folder.

Analysis Completion

Displays the total number of files analyzed as counted by `transcriptID` in each listed folder.

The **System Status** page is in active development, and additional metrics will become available in subsequent releases.

2. Applied additional security updates to enhance application resiliency.
3. Fixed an issue that caused Chrome to auto-fill the folder page search box with the current username.

V-Spark 4.0.1-8 Updates

1. Added system summary and performance metrics sections to the System Status page.

Overall Metrics shows the host system's CPU specifications, memory and storage usage, and information about front- and back-end services.

Total Jobs shows the number of jobs the system is currently processing.

Time to Analysis Completion shows the minimum, average, and maximum analysis completion times over various increments.

2. Changed a default Elasticsearch setting to allow for greater application complexity. This issue is identified by an Elasticsearch error in `search.log` containing a single numeric value without any additional context. The incomplete error message was caused by a faulty truncation of the full Elasticsearch error message, which could lead to difficulty diagnosing the root cause.

V-Spark 4.0.1-7 Updates

1. Applied additional security updates to enhance application resiliency.
2. Enabled a new System Status page. for system administrator user accounts accessible using the **System** option in the **Settings** menu.

The first iteration of the status page shows

- the names of active and queued folders,
- the number of active and queued transcript analysis requests, and
- the number of active and queued callbacks.

V-Spark 4.0.1-6 Updates

1. Applied multiple security updates to enhance application resiliency.
2. Improved error logging in `server.log` to include SMTP configuration errors when applicable.
3. Defined **unzip** as a package dependency so that it no longer needs to be installed manually on a host system.

V-Spark 4.0.1-5 Updates

1. Addressed a continuous memory growth issue affecting Job Manager processes. A new V-Spark option (`jm_memory_threshold`) has been added that specifies the memory threshold in bytes for Job Manager folder workers.

When the worker process exceeds the memory threshold, it is internally paused until the current jobs finish. It then terminates and is replaced by a new Job Manager folder worker process.

2. Fixed an issue that caused file scoring to fail without retry if Elasticsearch failed to create certain temporary indices.
3. Fixed a file permission issue with the **vspark-admin** script's **core-check** option. The script now reports the correct filesystem version when executing after a clean V-Spark install.
4. Fixed V-Spark's integration with ASR servers that use a Python API. Prior to this fix, when companies were configured to use an ASR server's Python API, files would get stuck in the processing queue. This was a rare edge case.
5. Fixed an issue that prevented updating an organization's retention policy via the UI. Prior to this fix, the retention policy could not be updated in certain scenarios. Any upgrade from a previous release to 4.0.1-5 requires running `core-update` to fully address this issue.

Known Issues in V-Spark 4.0

1. Audio files ingested with incorrectly formatted JSON metadata files may be imported without their expected metadata fields. These files are flagged as BAD METADATA , but are still imported, transcribed, and analyzed, which could lead to files being imported without expected metadata.
2. Application changes may not display in real time when made by another user from a different host. Although application editing works, users editing an application simultaneously from different hosts must refresh the **Application Editor** to see changes made by another user. This issue does not typically occur when both users are being served by the same host.

Previous versions

V-Spark 3.5.0

The next few sections provide release notes for V-Spark 3.5.0, describing updates and improvements that were introduced in this version of the product. This document also provides information about Known Issues in 3.5.0 that could affect your use and/or upgrade of the product. You may need to make changes after an upgrade to retain existing behavior or plan your upgrade process in order to prepare appropriately.

V-Spark 3.5.0 Release Notes

V-Spark Version 3.5.0 provides improvements in a number of areas.

1. V-Spark now supports distributed processing across multiple networked "nodes". This increases the availability, performance, and maintainability of your V-Spark installation while providing a single integrated interface.
 - Running multiple components on multiple hosts increases availability by duplicating vital functions. Components can be started and stopped individually, enabling you to update single components without disabling the entire system.
 - Nodes share a modular storage architecture and distributed filesystem. This enables the system to handle uploads and downloads from more clients without storage bottlenecks.
 - V-Spark now uses an improved queue-based processing flow. This reduces resource contention, increases the system's ability to utilize multiple resources, and improves your ability to monitor the progress of your data.

If your V-Spark processing needs exceed the computing capacity of one host, contact support@vocitec.com for more information about multi-node processing.

2. There have been multiple changes and enhancements made to the V-Spark REST API. If your workflow depends on the API, you will want to review this section carefully.
 - a. **Log API no longer supported** - The functionality of the `/log` API has been incorporated into the `/request` API, and the `/log` API is now deprecated and its jobmanager "processed.log" files are no longer supported. Any external tools that depend on this API call will need to be re-written.
 - b. **Status API modified** - The JSON schema of the data returned by the `/status` API has been changed. The organization and content of the output has been significantly modified, and the value of the queued count now reports the number of queued requests and not the number of queued files. Any external tools that depend on this API call may need to be modified.
 - c. **Partial /config calls may now be posted for companies** - When using the V-Spark API to create a new company, it is no longer necessary to POST a complete JSON configuration. You may now POST a partial config that only includes mandatory values for the

new company. When using the V-Spark API to modify a company's configuration, it is no longer necessary to POST a complete config that includes even unchanged values. You may now POST a partial config that only includes any changes to existing configuration values.

This change only applies to company configurations. Posting of partial configuration updates for organizations, folders, applications, or users using the REST API are not supported.

- d. **New /config/perms API** - The REST API now has a `/config/perms` call that returns information about which users have which permissions to which parts of the system.
 - e. **Error message for invalid JSON** - any V-Spark API that accepts JSON content does a check to determine the validity of any JSON content that is POSTed, and returns an `Invalid JSON` error if the JSON is not valid.
 - f. Requests to the `/transcribe` API that include the `s3key` option to submit a file stored in AWS S3 **must** specify the Amazon S3 region of the S3 bucket using the `region` option or the request fails silently.
 - g. **Failed File Details in /request** - The V-Spark `/request` endpoint displays files that were unable to be transcribed. Details explaining why the files were not transcribed are included; for example, the file type was not supported or the file contained bad metadata.
3. **Data loading directly from file system no longer supported** - V-Spark no longer supports direct ingestion of data from the local filesystem. Always upload audio via the V-Spark UI or API so that all inputs are properly queued and logged. If your workflow requires file system audio upload, contact support@vocitec.com to discuss solutions.
 4. **Improvements to Folder processing status dialog** - The **ASR processing status** section of this dialog has been improved to better show the progress of requests and files through V-Spark's processing queues. The **Queued** column now shows the number of processing requests waiting to be processed, and not the number of files in the queue.
 5. The following improvements have been made to V-Spark Application building and processing.
 - a. **Reprocess Applications by date range** - When reprocessing Applications, you may now choose a date *range* instead of simply a start

date for reprocessing. Only files within the selected date range will be reprocessed. You may still choose to have all files reprocessed, regardless of date.

- b. **Category names can now include Unicode characters** - The names of Application Categories may now include any character in the UTF-8 Unicode Standard. This includes accented characters and non-English characters of many types. The only restricted character is "." (Unicode 002E) which is used in API searches. New categories can be created with these characters in their names, and existing categories can be modified using the Application Editor so that their names better match the style of their source language.
 - c. **Metadata filter import and export** - Exporting a single category now includes metadata filters. This makes it easier to build applications and categories, as you can now export and import metadata filters, including both built-in and custom metadata filters.
6. **V-Cloud tokens are now verified when entered** - When you create a new Company that has a V-Cloud authorization token, add a V-Cloud authorization token to an existing Company, or update a Company's V-Cloud authorization token, the system now verifies connectivity to V-Cloud and the validity of the token before enabling V-Cloud for the Company. This check occurs whether you are working through the V-Spark user interface or the API.
7. **Callback Improvements** - Callback status messages were improved to include more specific information about status codes and whether the callback failed or succeeded. The improved callback messages will also make it easier to identify company, organization, and folder names in log entries.
8. **Delete Files from V-Spark** - Transcript files can be deleted using the V-Spark API or user interface. Deleted files will no longer be available in the file list `/transcript` view. Note that the summary data, stats, and dashboard views will not be updated to reflect deletion.
9. For greater system security, systems administrators can now configure a set of password policies.
 - a. **Minimum password length** - Administrators can now configure a minimum password length. Account passwords cannot be set or changed to passwords that do not have at least the required number of characters. The minimum password length is configurable, and defaults to 7 characters.
 - b. **Password aging** - Passwords can now expire after a configurable number of days. The default is for passwords to expire after 90 days. When an account password expires, the next time the user attempts to log in they will be redirected to a page where they will

update their account password. This policy can also be turned off so that passwords never expire.

Fixes in V-Spark 3.5.0

This section identifies specific issues that have been resolved in the 3.5.0 release of V-Spark:

1. **Companies and Orgs being deleted are restricted** - Both the V-Spark UI and API prevent you from creating new organizations, folders, or users if the company or organization that would own those items is in the process of being deleted.
2. **No usernames that are reserved words** - Accounts cannot be created with usernames that are on the list of words that are reserved for use by V-Spark internals. As of this release, the list of reserved words is: 'company', 'companies', 'org', 'orgs', 'app', 'apps', 'folder', 'folders', 'jobmgr', 'jobmgrs', 'user', 'users', 'all', 'system', and 'readonly'.
3. **Improved clean up of temporary files** - Temporary files that are created during searches and downloads were not being deleted promptly and could result in filesystem capacity problems. Temporary files are now deleted properly.

Maintenance Updates Since the Initial 3.5.0 Release

This section identifies fixes that were introduced in patch releases and hot fixes since the initial 3.5.0 release, and are therefore now included in the latest version of the 3.5.0 release. Maintenance versions and associated fixes are listed most recent to earliest:

1. **Improved response handling for systems using V-Cloud.** Eliminated extraneous 503 errors generated when communicating with V-Cloud to minimize unnecessary noise in Job Manager log files. (3.5.1-5)
2. **Backend V-Spark services now properly stop when Redis is down** - Prior to this change, stopping the backend service while Redis was down could lead to errors that required killing backend processes manually in order to restart them. (3.5.1-4)
3. **Update to support the latest version of the Datatool** - if your V-Spark installation has a copy of the Datatool provided by the Voci support team, this release will add value by providing some core changes that enable the Datatool to exercise its sampling feature with V-Spark 3.5.x releases. (3.5.1-4)
4. **Improved handling of hardware boot and reboot events** - The V-Spark initialization process was updated to permit V-Spark services to start automatically during hardware boot and reboot. (3.5.1-3)
5. **Dashboard performance optimization** - Loading optimization updates were made to allow all folder-level dashboard widgets to load faster. (3.5.1-3)
6. **Google Chrome 83 support update** - Chrome 83 now blocks all downloads initiated from or instantiated in a sandboxed iframe by default. This update enables V-Spark to override that default behavior and allow file downloads. (3.5.1-2)
7. **Failed file details in /request** - The V-Spark /request endpoint displays files that were unable to be transcribed. Details explaining why the files were not transcribed are included; for example, the file type was not supported or the file contained bad metadata. (3.5.0-17)
8. **Callback improvements**- Callback status messages were improved to include more specific information about status codes and whether the callback failed or succeeded. (3.5.0-11)
9. **Delete files from V-Spark**- Transcript files can be deleted using the V-Spark API or user interface. Deleted files will no longer be available in the file list /transcript view. However, summary data will not be updated to reflect deletion at this time. (3.5.0-10)

Announcements with 3.5.0

The V-Spark 3.5.0 Application Development Guide has been significantly updated, reorganized, and enhanced for this release. Improvements include:

1. The "Application Scoring" chapter now explains the different types of scoring in more detail.
2. The "Viewing Application Scores" chapter has been reorganized and now contains significantly more detail about V-Spark features that enable you to view and interpret the analysis results produced by your applications.
3. Information on reprocessing audio and searching transcripts has been expanded with more detail.
4. Most figures have been updated and their callouts clarified.
5. **Application Leaf-level Scores** - Leaf-level application scores have been renamed from "count scores" to "leaf scores" throughout documentation.

Known Issues in 3.5.0

1. Partial configuration updates for a
 - organization
 - folder
 - application
 - or user

using the REST API are not currently supported. Attempting to update a subset of one of these types using the REST API currently returns HTTP error code 400. You must supply all fields for the configuration to update any field(s) using the REST API.

2. Application re-processing of more than 100,000 files in a single operation may time out.
3. Issues with Dashboard Displays
 - Files uploaded to a folder that is linked to a disabled application or that have not been re-processed will still appear in the Files View on the application dashboard
4. Issues with Search
 - Time-based search
 - Searches are turn-based and may return results for phrases that are outside of the requested time frame if those phrases are within the same speaker turn
 - Ad-hoc Search snippets in the Files View may display results outside the specified time interval
 - Exceptions to standard File Details highlighting:
 - Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone". Searches using (~t) and (~n) will only be highlighted if the words matched

appear in the order specified in the search tag.

- Highlighting may be incomplete for matches marked with emotion
- Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds
- Matches for "@"* (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
- Excluded application search phrases may be highlighted on the File Details page. Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.

5. Other

- Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.
- App Reprocessing will encounter an error if the reprocessing occurs while an App or Category is being deleted. This error is recorded in the server logs.
- Searches that contain both wildcards and spaces do not return results.
- Most punctuation characters are not supported in search strings.
- The `/search` API allows searches for application categories that have hyphens and other restricted characters in their names, however these searches fail. Voci application templates that had hyphens and other restricted characters in their names have been modified to remove the restricted characters. All existing application category names should be modified to remove restricted characters.

V-Spark 3.4.3

The next few sections provide release notes for V-Spark 3.4.3, describing updates and improvements that were introduced in this version of the product. This document also provides information about Known Issues in 3.4.3 that could affect your use and/or upgrade of the product. You may need to make changes after an upgrade to retain existing behavior, or plan your upgrade process in order to prepare appropriately.

V-Spark 3.4.3 Release Notes


V-Spark Version 3.4.3 provides improvements in a number of areas including analysis metrics, API usage, metadata handling, search, and general system operation.

Daily statistics are now generated in scheduled batches. This replaces the old behavior of generating statistics one file at a time.

The database queries that are used during the calculation of Application summaries have been optimized to be more computationally efficient.

1. **Improved Diarization Scores** - Diarization now uses a new scoring algorithm that better represents how well mono audio was split into two speaker channels. Scoring used to be in relation to previously uploaded transcripts, but is now independent of transcript history.
2. **Better File Revision Tracking** - The system now allows more than one transcript to have the same filename and received date. Duplicated transcripts used to be given a different timestamp to differentiate them from previously uploaded transcripts. Duplicates are now differentiated by transcript ID (tID).
3. **Improved Loading of System Configuration** - System configuration is now stored in a directory of multiple files instead of a single file. Configuration files in the **vspark.config.d** directory are read in alphabetical order at server startup. Legacy configuration files are still read, but the files in **vspark.config.d** always take precedence.
4. **"Last Generated" Time Display** - Statistics are now shown with the time they were last generated. The display may lag behind actual statistics, as the analysis only runs every five minutes, (This interval is configurable.)
5. **New /sysinfo API** - The REST API now has a /sysinfo call that returns system information.
6. **Improved /config/users API** - New user accounts created using the /config/users API can now be configured to use external authorization methods. Note that existing users cannot be reconfigured in this way.
7. **Additional Fields Supported in Search API** - The Search API feature allows four additional output fields: "tags", "folder", "url", and "es_doc_id" (Elasticsearch document ID).
8. **Search Results Now Display Custom Metadata** - When doing a search for custom metadata, the values of the custom metadata field being

searched are displayed in the search results. It is no longer necessary to click through to the file details to see how results match the custom search criteria.

9. **Custom Metadata Now Appear in Filter Dropdown** - If your Application is linked to a folder with custom metadata fields, those fields now appear directly in the "Filter" drop-down menu instead of a separate "Custom Metadata" menu. This makes metadata fields easier to find and faster to utilize.
10. **Unicode Support in Metadata** - When uploading audio files with metadata, fields in the metadata may now contain any valid Unicode values, except for reserved characters Colon (":", Unicode 003A) and Equals Sign ("=", Unicode 003D).
11. **Improved Login with External Authorization Methods** - Improved linking for systems that have external authorization enabled. Contact support@vocitec.com  if you are interested in this capability.
12. **Logging of Unsuccessful Login Attempts** - For security purposes, the activity log now tracks all unsuccessful login attempts along with the reason the attempt failed, such as incorrect password or username, unverified user account, user account disabled, or unmapped user account.
13. **Improved View of User Permissions** - In the Company table of the Accounts Settings page, there is now a "Permissions" column that contains a "View users" link. Follow this link to view a list of all users who have permissions that apply to this company and its organizations. Permissions are displayed in a dialog that shows all V-Spark users with permissions, whether those permissions are explicit to the company you chose, or implicit in the permissions those users have in the system. If a user's permissions in one company or organization are based on the permissions they hold in another company, the name of that other company is displayed in parentheses next to their username. By default, permissions for the entire company are displayed. To show permissions for a specific organization within the company, choose the name of the organization from the drop-down menu above the list.
14. **Notification of Usage Limits** - The system generates a system notification when a company nears their usage limit, and another if the usage limit is reached. The warning level is configurable, and defaults to 80%. These notifications are visible to system administrators and to users with write access to the affected company. Notifications are dismissible within V-Spark, and will only display once per day. These notifications are also sent via email to system administrators and company administrators, and recorded in the activity log.

15. Folder processing can now be paused via the API - Using the `/config/folders` API, you can now pause the processing of a Folder by POSTing a JSON configuration file for the Folder that has the `mode` property of the Folder set to the value `paused` . Resume processing by POSTing JSON for the Folder that has the `mode` property set to `active` .

Fixes in V-Spark 3.4.3

This section identifies specific issues that have been resolved in the 3.4.3 release of V-Spark:

1. **Deletion Jobs Now Span System Restarts** - If the system is shut down or restarted while a deletion job was executing, the job was not being completed. Incomplete deletion jobs, whether user-initiated or automatic, are now re-queued for execution when the system starts.
2. **Improved Handling of Application Scoring Failures** - Previously, uploaded files that failed scoring would be moved to the error directory, even if other analysis succeeded. These files will no longer be moved to the error directory. Scoring failures are now recorded in the activity log and backend logs, showing which file failed scoring, and that none of the Applications that use that file will have scores for that file.
3. **Transcript Processing is Now Atomic** - If an Elasticsearch or MySQL failure is encountered during transcript loading, all related updates are rolled back cleanly to preserve data integrity.
4. **System Mode Changes are Now Logged** - System Mode changes between "normal", "safe", and "readonly" were not being logged, which made diagnosing trouble more difficult. The system now logs when mode changes occur, and (depending on warning level) sends email to system administrators.
5. **Trending Phrases No Longer Reported** - Trending phrases were being included in daily and monthly reports, even if Trending was disabled. This has been corrected.
6. **Appstats Now Returns Categories with Scores of Zero** - The `/appstats` API was not returning any categories if the top category score was zero, even if the depth level was defined to return categories with scores of zero. This defect was corrected, and categories are now returned properly.
7. The `/transcribe` API was not processing audio files if they were uploaded from Amazon's Simple Storage Service (S3). The system was returning an empty reply. This defect has been corrected, and the system now returns a transcript as expected.
8. **Files with Empty Utterances Now Update Properly** - Uploaded files that result in an empty Utterances field were ignored during application reprocessing. This issue has been resolved and these files are now updated correctly.

9. **Config Now Returns Correct Template Value** - The `/config/CO_SHORT/ORG_SHORT/apps/APPNAME` API was not returning the correct template value for Applications that were created from templates. It was returning the "template" value of "custom" even for Applications that were created from templates. This issue has been resolved, and the system now returns the correct template value.

Maintenance Updates Since the Initial 3.4.3 Release

This section identifies fixes that were introduced in patch releases and hot fixes since the initial 3.4.3 release, and are therefore now included in the latest version of the 3.4.3 release. Maintenance versions and associated fixes are listed most recent to earliest:

1. **Improved response handling for systems using V-Cloud.** Eliminated extraneous 503 errors generated when communicating with V-Cloud to minimize unnecessary noise in Job Manager log files. (3.4.3-15)
2. Fixed an issue that would cause Chrome to auto-fill the search box on a folder page with the current username. (3.4.3-14)
3. Fixed a typo in the How it Works page image. (3.4.3-14)
4. **Google Chrome 83 Support Update** - Chrome 83 now blocks all downloads initiated from or instantiated in a sandboxed iframe by default. This update enables V-Spark to override that default behavior and allow file downloads. (3.4.3-13)
5. **Improved handling of email communication** - Email notifications were updated to list email recipients in the BCC field of the message. Also, the TO field is now set to the application's global address. (3.4.3-12)
6. **Improved support to AWS S3 downloads** - Improved download retry logic to account for possible issues during network transport. (3.4.3-11)
7. **Run V-Spark on a Single Core** - Enabled V-Spark to run on a machine with a single core on its single CPU. (3.4.3-10)
8. **Eliminated Temporary File Leaks** - Improved cleanup of temporary files. (3.4.3-10)
9. **Fix Application Reprocessing Failure** - Improved mysql retry logic to reduce application processing errors. (3.4.3-10)
10. **Audio Duration Update** - Improved how a file's audio duration is determined. (3.4.3-10)
11. **S3 Update** - V-Spark now fully supports signed S3 bucket requests for all AWS regions by adding support for the AWS Signature Version 4 Signing Process. Previously only Version 2 was supported. (3.4.3-10)
12. **Partial Content Requests** - Some browsers, such as Google Chrome, specify a desired content range for cases such as fetching a piece of

an audio file to play/seek in its embedded audio/video player. V-Spark now honors partial content requests on the part of a browser by responding with 206 HTTP code and the piece of audio that was requested. (3.4.3-10)

13. **Elasticsearch connection management** - Connections to Elasticsearch clients are now managed more effectively to prevent connection errors. (3.4.3-9)
14. **Folder ASR options** - Expanded support for folder ASR options to allow adding values with multiple elements via the `/config/folders` API and GUI. Previously, both methods only supported simple ASR options with string, boolean, or numeric values. Boolean ASR option values must now be formatted properly for them to be correctly recognized. The Boolean values `true` and `false` must be entered in all lower case with no quotation marks. Values containing capital letters or contained within quotes (such as `True` or `"false"`) are interpreted as string values, and *not* as Booleans. (3.4.3-8)
15. **Request API Content-type** - The `/request` API `status` and `details` calls were returning output with the incorrect content-type value. The content-type has been restored to "application/json". (3.4.3-8)
16. **Cloud Tokens Checked Properly on Company Update** - Updating a company that uses cloud tokens was pausing all of the company's folders due to a token check that was not matching the tokens to the folder servers properly. Tokens are now checked properly, enabling companies to update without their folders being paused. (3.4.3-7)
17. **Emotion Word Highlighting** - Turning off "emotion" view on file details page will mitigate a known issue where phrase matches that overlap words marked with emotion are not highlighted. Now, opening a file details page from the application dashboard will turn off emotion by default, to improve the display of application phrase matches. (3.4.3-7)

Announcements with 3.4.3

V-Spark 3.4.3 does not add any new announcements of upcoming changes and deprecation. Please see Announcements with 3.4.2 for announcements made with the 3.4.2 release.

Known Issues in 3.4.3

1. Partial configuration updates for a company, organization, folder, application, or user using the REST API are not currently supported. Attempting to update a subset of a company or organization using the REST API currently returns HTTP error code 400. You must supply all fields for the configuration in order to update any field(s) using the REST API.
2. Application re-processing of more than 100,000 files in a single operation may time out.
3. Issues with Dashboard Displays
 - Files uploaded to a folder that is linked to a disabled application or that have not been re-processed will still appear in the Files View on the application dashboard
4. Issues with Search
 - Time-based search
 - Searches are turn-based and may return results for phrases that are outside of the requested time frame if those phrases are within the same speaker turn
 - Ad-hoc Search snippets in the Files View may display results outside the specified time interval
 - Exceptions to standard File Details highlighting:
 - Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone". Searches using (~t) and (~n) will only be highlighted if the words matched appear in the order specified in the search tag.
 - Highlighting may be incomplete for matches marked with emotion
 - Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds

- Matches for "@"* (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
- Excluded application search phrases may be highlighted on the File Details page. Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.

5. Other

- Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.
- App Reprocessing will encounter an error if the reprocessing occurs while an App or Category is being deleted. This error is recorded in the server logs.
- Searches that contain both wildcards and spaces are not returning results.
- Most punctuation characters are not supported in search strings.
- The `/search` API allows searches for application categories that have hyphens and other restricted characters in their names, however these searches fail. Voci application templates that had hyphens and other restricted characters in their names have been modified to remove the restricted characters. All existing application category names should be modified to remove restricted characters.
- If the server is restarted while files are waiting in the Callback queue, those files will not be delivered and the callback will not complete. Allow Callback delivery to complete before restarting the server.

V-Spark 3.4.2

The next few sections provide release notes for V-Spark 3.4.2, describing updates and improvements that were introduced in this version of the product. This document also provides information about Known Issues in 3.4.2 that could affect your use and/or upgrade of the product. You may need to make changes after an upgrade to retain existing behavior or plan your upgrade process in order to prepare appropriately.

V-Spark 3.4.2 Release Notes

V-Spark Version 3.4.2 focuses on internal improvements in performance and stability, and also adds new administrative capabilities for managing and monitoring disk use.

1. **Configurable data retention policies** - Companies and organizations can now specify the period of time for which call and analysis data is preserved on their systems. This helps reduce storage costs, and also reduces the chance that your V-Spark system will encounter problems due to running out of storage space.
2. **Improved handling of low storage conditions** - System administrators can now configure the minimum amount of free local storage space that must be available in order for V-Spark to continue transcribing and analyzing new data.
3. **External Authentication** - System administrators can now configure V-Spark to use certain external authentication mechanisms to enable user logins. Contact support@vocitec.com if you are interested in this capability.
4. **Search Unification** - Adhoc search and search within applications are substantially unified at customer request to make it easier to predict the behavior of search within applications.
5. **Transcribe API now accepts single files** - When using the `/transcribe` API to submit individual files for transcription, single audio or JSON files can now be submitted individually. This was always possible using the GUI and is now possible using the API. Files submitted individually will not be associated with each other. Multiple files, such as audio files and metadata files that should be associated, must still be combined into a single ZIP file and submitted together.
6. **Search Expression Improvements** - In application scoring searches, the underscore character ('_'), a wildcard matching 0 or more words, and the '->' character combination, matching phrases in consecutive speaker turns, are no longer supported. When upgrading to V-Spark 3.4-2, the underscore character will be automatically scored as '~t' (search for phrases appearing within a single speaker turn) in existing applications. The '->' character combination is no longer supported in apps, and must be replaced manually in applications because there is no exact replacement.
7. **Application Search Scoring Improvements** - Leaf-level categories without include phrases will now receive a count score of 1 to indicate

that files have passed filters in the leaf's upper-level categories. Scoring for time-based phrases is available on a limited/early-access basis. Contact support@vocitec.com for more information.

Fixes in 3.4.2

This section identifies specific issues that have been resolved in the 3.4.2 release of V-Spark:

1. **Wildcard search works for custom metadata date fields** - Previously, custom metadata in `datetime` format could not be searched using wildcards, and trying to do so would generate an error.
2. **Time-based adhoc search results for `~e` are now correct** - If any call data was uploaded for analysis on versions 3.3.0, 3.3.1, and 3.4.1, an additional update script must be run to enable the use of `~e` in adhoc searches. Contact support@vocitec.com for more information.
3. **Complex searches for transcript exclude phrases with other search terms are now supported** - Previously, adhoc searches for different types of terms (custom metadata, tags, etc.) with transcript exclude phrases would ignore the exclude phrase.
4. **App reprocessing of folders with a very large number of files was timing out and producing an error**. The timeout value for search operations included in App reprocessing is now configurable. If you need to reprocess folders that contain a very large number of files, you can avoid errors by increasing the timeout value.
5. **Clicking the Call duration graphs now updates the file list** - In versions 3.3.1 and 3.4.1, no results were displayed when the Daily Charts graphs for Call duration were clicked.
6. **User subscriptions requests to Daily Report emails were sometimes being duplicated by the UI, resulting in multiple copies of some reports being attached to the email**. Subscription requests are now checked against existing values before they are added to the description table, which avoids this duplication. Users who are receiving duplicate reports should unsubscribe from all reports, and re-subscribe manually.
7. **Spaces in custom metadata field names are no longer saved as non-breaking spaces** - Fixed an issue where spaces in custom metadata field names were being saved as non-breaking spaces which could cause an issue with app processing if custom metadata was used as a category filter.
8. **Text files sent via callbacks for non-English calls now include the transcript content** - In 3.4.1, text files sent via callbacks were empty for transcripts that contained non-ASCII characters. JSON files were not affected.

9. **New fields cannot be given reserved names** - It was possible to add custom metadata fields that had the same name as reserved built-in fields. The UI and API now restrict the creation of fields so that only non-reserved names can be used.
10. **CSV export supports all available fields** - All available fields can now be selected in the V-Spark user interface for inclusion when exporting data in CSV (Comma-Separated Value) format. UI CSV export now matches `/search` API export results.

Maintenance Fixes Since the Initial 3.4.2 Release

This section identifies fixes that were introduced in patch releases and hot fixes since the initial 3.4.2 release, and are therefore now included in the latest version of the 3.4.2 release. Maintenance versions and associated fixes are listed earliest to most recent:

1. **Folder callbacks send only specified files types** - Callback worker was sending JSON, Text, and MP3 files, regardless of which file types a folder's callback was configured to send. (3.4.2-2)
2. **Content-type for MP3 files is specified in http callbacks** - HTTP callbacks were not specifying the content-type when sending MP3 files. (3.4.2-2)
3. **Application scoring does not fail for files with no utterances** - System was reporting processing errors when a file with no utterances was uploaded to a folder linked to an application or included in a folder that was reprocessed for an application. (3.4.2-2)
4. The `/config/C0_SHORT/users/USERNAME` API call was returning user information within an extraneous username layer. The response has been modified to remove that layer. (3.4.2-5)
5. The `/config` API call was returning Application information within an extraneous Apps layer, even if only one Application was requested. The response has been modified to remove that layer when configuration information for only one App has been requested. (3.4.2-5)
6. The `/search` API was returning an error when searching for App scores. The command has been fixed to prevent the error and return App scores as expected. (3.4.2-5)
7. Daily and monthly reports would fail to generate if any `agentid` had more than 24 hours of talk time in one day. This can occur if talk time for multiple agents is reported with no associated `agentid` or with the same `agentid`. The reports have been corrected to report daily agent talk times in HH+:MM:SS format if the number of talk time hours for that agent is greater than 24. (3.4.2-6)
8. The `/search` API was not allowing searches for categories that had spaces in their names. These searches would fail. This defect has been corrected, and searches for categories that have spaces in their names return proper results. (3.4.2-6)
9. Application scoring was not producing correct scores for categories with "Speakers:" phrases. This defect has been corrected, and the scores now reflect both speakers. (3.4.2-6)

10. Application scoring was not producing correct scores for categories with phrases containing boolean operators OR (|) and AND (&). (3.4.2-7)
11. Application scoring was not producing correct scores for categories with custom metadata exclude filters if the file being scored did not contain the metadata *field* being filtered. (3.4.2-7)
12. Additional security fixes were applied. (3.4.2-7)
13. Improve handling of concurrent `/transcribe` API requests. (3.4.2-8)

Announcements with 3.4.2

V-Spark 3.4.2 does not add any new announcements of upcoming changes and deprecation. Please see Announcements with 3.4.1 for announcements made with the 3.4.1 release.

Known Issues in 3.4.2

1. Partial configuration updates for a company (/config/ CO_SHORT) or organization (/config/ CO_SHORT/ORG_SHORT) using the REST API are not currently supported. Attempting to update a subset of a company or organization using the REST API currently returns HTTP error code 400. You must supply all fields for the configuration in order to update any field(s) using the REST API.
2. Application re-processing of more than 100,000 files in a single operation may time out.
3. Custom metadata field names may only contain standard ASCII characters, excluding = and : , which are reserved characters.
4. Issues with Dashboard Displays
 - Files uploaded to a folder that is linked to a disabled application or that have not been reprocessed will still appear in the Files View on the application dashboard
5. Issues with Search
 - Time-based search
 - Searches are turn-based and may return results for phrases that are outside of the requested timeframe if those phrases are within the same speaker turn
 - Ad-hoc Search snippets in the Files View may display results outside the specified time interval
 - Exceptions to standard File Details highlighting:
 - Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone". Searches using (~t) and (~n) will only be highlighted if the words matched appear in the order specified in the search tag.
 - Highlighting may be incomplete for matches marked with emotion

- Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds
- Matches for "@"* (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
- Excluded application search phrases may be highlighted on the File Details page. Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.

6. Other

- Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.
 - The `/appstats` API is not returning any categories if the top category score is zero, even if the depth level is defined to return categories with scores of zero.
 - The `/transcribe` API is not processing audio files if they were uploaded from Amazon's Simple Storage Service (S3). The system is returning an empty reply.
 - Uploaded JSON files that contain an empty Utterances field are not being reprocessed properly.
7. The `/search` API allows searches for application categories that have hyphens and other restricted characters in their names, however these searches fail. Voci application templates that had hyphens and other restricted characters in their names have been modified to remove the restricted characters. All existing application category names should be modified to remove restricted characters.
 8. The `/config/CO_SHORT/ORG_SHORT/apps/APPNAME` API is not returning the correct template value for Applications that were created from templates. It is returning the "template" value of "custom" even for Applications that were created from templates.

V-Spark 3.4.1

The next few sections provide release notes for V-Spark 3.4.1, describing new features and capabilities that were introduced in the 3.4 version of the product. This document also provides information about Known Issues in 3.4.1 with V-Spark 3.4 that could affect your use of the product.

V-Spark 3.4.1 Release Notes

V-Spark Version 3.4.1 provides performance, usability, and stability improvements, and also provides significant API enhancements.

1. **Call Tagging** - makes it easy to identify specific transcriptions or the results of previous searches by enabling you to apply a **tag** (a logical label) to them and then subsequently search for transcriptions that are associated with that tag
2. **User management enhancements** - more granular permissions and a substantially enhanced GUI simplify setting new and existing permissions and associated roles. These user management improvements are also reflected in the updated REST API, as described later in this section.

If you are upgrading from a previous release of V-Spark, existing user permissions will be converted into those that work with the new permissions model, as shown in the following table:

Mapping Between pre-3.4 and 3.4 Permissions

Pre-3.4 Permission/Role	3.4 Permission/Role
Admin	System admin
Manager	Create/Edit rights on all organizations that they had been given access to within the company under which their account was created (their home company). View rights on any other organization that they had access to
User	View rights on any organization that they had access to

3. **RESTful API - Phase 2** - significant API enhancements make it easier to deploy and manage V-Spark programmatically:
 - a. Enhanced user permissions provide easier delegation of administrative responsibilities within a V-Spark deployment and increase data privacy by separating administrative roles from the ability to view data
 - b. API access to daily overview and agent application statistics makes it easier to extract reports and use extracted data in third-party reporting applications
 - c. Applications can now be created and modified programmatically, making it even easier to use, customize, and deploy applications from templates or simply from outside the V-Spark environment
4. **Expanded email configuration for partners** - Email templates and improved configuration files make it easy for partners to customize system email sender and content
5. **Trending support is disabled** - terminology trend tracking and analysis is now disabled by default, based on its limited use, capabilities, and impact on system performance. We have decided that it needs additional development in order to perform as originally envisioned. At this point in time, this additional development is not a line item in our development plans. The current implementation and of tracking and analyzing terminology trends can be re-enabled by an administrator who has access to the V-Spark installation, by modifying the V-Spark configuration file. Re-enabling the current implementation of trending support is not recommended.
6. **End of day prediction for call volume** (also referred to as "predicted overall volume") is no longer performed.
7. **Company modification via API requires root token** - a company authentication token does not provide sufficient privileges to modify a company itself.

Maintenance Fixes Since the Initial 3.4.1 Release

This section identifies fixes that were introduced in patch releases and hot fixes since the original 3.4.1 release, and are therefore now included in the latest version of the 3.4.1 release. Maintenance versions and associated fixes are listed from earliest to most recent:

1. **Files scored during upload had incorrect summaries of average call duration and silence (3.4.1-1)** - Previously, the average call duration and silence time in the Agent Summary table were calculated using only the files that had scored. Now these averages are calculated based on all of an agent's calls.
2. **Audio player controls now visible in in the Chrome browser (3.4.1-2)** - Chrome Version 63.0.3239.84 changed how the built-in audio player was invoked, causing player controls not to be displayed. This was fixed in V-Spark and in later versions of Chrome.
3. **Enabled Client metadata fields as valid search parameters in /search API (3.4.1-2)** - Searching for Client metadata fields in /search API was not previously supported.
4. **API /search for exclude phrase are now supported (3.4.1-4)** - Previously, no results were returned for API searches with exclude phrases.
5. **setupHostname error (3.4.1-4)** - Prior to this, configuration files without a hostname specified could cause an error.
6. **Added support for tag search with | or & in a single term (3.4.1-4)** - Tag searches that included boolean operators in a single term previously did not return results.
7. **Agent/Client gender and clarity now default to null when speaker does not exist (3.4.1-4)** - In data uploaded prior to this change, Agent/Client gender defaulted to female and Agent/Client clarity defaulted to 0.0 when the speaker did not exist, such as in calls with only one speaker.
8. Improved startup procedures to avoid licensing timing issues.
9. **Updating a company requires the root token (3.4.1-6)** - Previous releases in the 3.4 family allowed companies to be updated by API calls that used the company token for authorization. This has been corrected so that the root token is required for company updates, while the company token can be used to update related organizations and associated users.

10. Custom metadata is displayed regardless of the case specified in folder (3.4.1-7) - Previously, only custom metadata that matched the case specified in the folder were displayed.

Announcements with 3.4.1

The following upcoming changes to V-Spark are being announced. These changes are not present in V-Spark 3.4.1, but you should review your use of V-Spark and prepare for changes to these aspects of its use:

1. In application scoring searches, the underscore character (' _ '), a wildcard matching 0 or more words, and the ' - -> ' character combination, matching phrases in consecutive speaker turns, are being deprecated in favor of the ' ~n ' (search for phrases appearing up to n words apart) and ' ~t ' (search for phrases appearing within the same speaker turn) expressions.

Known Issues in 3.4.1

1. Issues with Dashboard Displays

- Files uploaded to a folder that is linked to a disabled application or that have not been reprocessed will still appear in the Files View on the application dashboard

2. Issues with Search

- Time-based search
 - Searches are turn-based and may return results for phrases that are outside of the requested timeframe if those phrases are within the same speaker turn
 - Ad-hoc Search snippets in the Files View may display results outside the specified time interval
- Exceptions to standard File Details highlighting:
 - Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone". Searches using (~t) and (~n) will only be highlighted if the words matched appear in the order specified in the search tag.
 - Highlighting may be incomplete for matches marked with emotion
 - Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds
 - Matches for "@*" (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
 - Excluded application search phrases may be highlighted on the File Details page. Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.

3. Other

- Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.

V-Spark 3.4.0

V-Spark 3.4.0 was an internal pre-release of V-Spark 3.4.1, used for QA and testing.

V-Spark 3.3.1

This section provides release notes for V-Spark 3.3.1, describing new features and capabilities that were introduced in the 3.3.0 and 3.3.1 versions of the product. This document also provides information about Known Issues in 3.3.1 with V-Spark 3.3.1 that could affect your use of the product.

V-Spark 3.3.1 Release Notes

V-Spark Version 3.3.1 provides performance, usability, and stability improvements, and also provides significant API enhancements.

1. Added new REST APIs to enable programmatic access to V-Spark configuration and settings. The APIs added are `/config`, `/list`, `/logs`, `/search`, and `/status`. See the "Using the V-Spark 3.5.0 API" for more information.
2. Created an online, searchable version of all Voci documentation that can be used by Voci customers, prospects, and employees. PDF documentation can still be referenced or the online documentation can be linked to from the V-Spark Help menu. The online documentation is built from the same content that is used to produce Voci's PDF documentation. Access to the online documentation requires the ability to access [the documentation web site](#) and an account in Voci's portion of Salesforce.com.
3. The Elasticsearch software used by V-Spark was upgraded, enabling easier integration of V-Spark with external reporting tools.
4. Improved definition and display of speaker turns generated from V-Spark's Automatic Speech Recognition (ASR) utterances
5. System-wide announcements can now be set by an administrator so that they are visible to anyone who is using V-Spark
6. V-Spark now supports a read-only mode that enables maintenance and diagnostic tasks to take place while users can still view and explore existing data. No new audio can be added or processed while V-Spark is in read-only mode.
7. The Settings menu's Folders page now provides graphical controls that enable you to pause and continue processing. Pausing file processing enables current processing operations to complete, but no new processing can be done until the continue portion of the pause/continue control is selected.
8. Improved folder processing status, enabling folder views to be sorted by time (most recently processed). V-Spark now also provides access to the folder processing log queue, enabling users to see which folders are waiting to be processed, are currently being processed, are waiting for analysis, or have encountered errors during processing.
9. Reporting has been enhanced to provide access to more data, such as a table that provides agent summary information
10. The File Details page now:

- a. displays overtalk with blue edge highlighting in the transcript and correspondingly displays overtalk scores in blue text
 - b. enables you to collapse client metadata so that the general details are easier to review
 - c. provides information about the processing date of application scoring
11. File metadata can now be displayed and used in Adhoc searches, rather than only being viewable on the **File Details** page
 12. Enhanced the Voci license server to improve robustness and responsiveness

Fixes in 3.3.1

This section identifies specific issues that have been resolved in the 3.3.1 release of V-Spark:

1. Fields within CSV reports that are exported from the dashboard are now quoted
2. Improved reliability of application scoring and reprocessing
3. RequestIDs are now correctly assigned and displayed when reloading existing Voci JSON files and associated MP3 audio
4. Favicons are now correctly shown in V-Spark pages in a web browser
5. Configuration files enable partners to customize system email sender and content
6. Highlighting on the File Details page now correctly considers word boundaries when interpreting regular expression phrases

Known Issues in 3.3.1

1. Exporting large numbers of files (10K or larger) as CSV can take an extremely long time, and may not complete successfully. A suggested workaround when you want to export large file lists in CSV format is to subset those lists into smaller groups that you can subsequently combine in your target spreadsheet application.
2. Calls to the `/search` API will only return custom metadata for matches if a folder has been specified. In other words, suppose that a call to the `/search` API matches files in folders under a company named `Doc Test Co` and an organization named `Doc Testing`, and that custom metadata is associated with a folder under that organization that is named `Test01`. Matching results in that folder that are returned by calling `/search/DocTestCo-DocTesting` will not contain any custom metadata, while matching results that are returned by calling `/search/DocTestCo-DocTesting/Test01` will contain the custom metadata.
3. Downloading JSON and MP3 data simultaneously for large numbers of files (5K or greater) can fail without that failure being displayed in the GUI. The size of the zip file created for downloading depends on the size of the files that are being included in the archive. A suggested workaround when you want to download transcripts and MP3 data for large numbers of files is to subset your download requests into smaller numbers of files.
4. Issues with Dashboard Displays
 - a. Stat summaries for dates in the past are not updated when files with historical timestamps older than the current day are loaded
 - b. Files uploaded to a folder that is linked to a disabled application or that have not been reprocessed will still appear in the Files View on the application dashboard
5. Issues with Search
 - a. Time-based search
 - i. Searches are turn-based and may return results for phrases that are outside of the requested timeframe, if they are within the same speaker turn
 - ii. Ad-hoc Search snippets in the Files View may display results outside the specified time interval

b. Exceptions to standard File Details Highlighting (in order of severity)

- i. Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone". Searches using (~t) and (~n) will only be highlighted if the words matched appear in the order specified in the search tag.
- ii. Highlighting may be incomplete for matches marked with emotion
- iii. Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds
- iv. Matches for "@*" (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
- v. Excluded application search phrases may be highlighted on the File Details page. Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.

6. Other

- a. The Trending View may take some time to fully load, especially for longer date ranges
- b. Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.

V-Spark 3.3.0

V-Spark 3.3.0 was an internal pre-release of V-Spark 3.3.1, used for QA and testing.

V-Spark 3.2.1

This section provides release notes for V-Spark Version 3.2.1, a maintenance release that was delivered to customers to resolve a few minor problems in the V-Spark 3.2.0 release. There are no additional known issues for 3.2.1 - all of the known issues given in the 3.2 release notes still apply to the 3.2.1 release.

V-Spark3.2.1 Release Notes

V-Spark Version 3.2.1 resolves a variety of issues in the V-Spark 3.2.0 release. These are:

1. Added support for WebAPI ASR servers/load balancers.
2. Added support for single file non-zip uploads with no restrictions on filename extension.
3. Added optional support for insecure HTTPS for ASR servers and callbacks.
4. Improved validation of Company, Organization, Folder, and Application names.
5. Display activity logs entries with local server time.
6. Removed client emotion filter from the customer experience application template.
7. Fixes:
 - a. Improved handling of server and model lists when configuring companies and folders.
 - b. Resolved several edge case errors when importing data and reprocessing.
 - c. Resolved error processing metadata keys set to empty string.
 - d. Improved V-Spark service script to better kill old processes when restarting.
 - e. Eliminated need to clear browser cache after upgrades.
 - f. Improved handling of errors produced by versions of the libmediainfo package that are greater than version 0.7.94.
 - g. Fixed folder processing log viewer hangs due to mismatched table headings.
 - h. Ensured that downloaded text filename is consistent with json and mp3 filenames.
 - i. Enabled use of no authentication in SMTP server setup.

- j. Improved handling of Boolean values for custom ASR tag configuration.
- k. Fixed broken link to G729A help document.
- l. Fixed errors when deleting archived files on some remote filesystems.
- m. Fixed errors related to trailing slashes on some callback URLs.

V-Spark 3.2

This section provides release notes for V-Spark 3.2, describing new features and capabilities that were introduced in the 3.2 version of the product. This document also provides information about Known Issues in 3.2 with V-Spark 3.2 that could affect your use of the product.

V-Spark 3.2 Release Notes

V-Spark Version 3.2 provides new language models, significant performance, usability, and stability improvements, and API enhancements to simplify transcript and audio retrieval.

1. Added new `request` API to retrieve Text and JSON transcripts and MP3 audio from V-Spark outside of the traditional callback interface. The `request` API provides 4 verbs to query job status and retrieve results by RequestID: `status`, `summary`, `details`, and `result`. See "Using the V-Spark 3.5.0 API" for more information.
2. Added Support for Non-English Languages
 - a. These language models support the following features (requires Automatic Speech Recognition (ASR) greater than or equal to 5.4.1):
 - i. File loading
 - ii. Playback and Display
 - iii. Ad-hoc Search
 - b. The following features are not supported in these language models for 3.2:
 - i. Applications
 - ii. Application Templates
 - iii. Reporting
3. Internet licensing supported in V-Spark. Internet licensing simplifies license management, making it easy to remotely add, upgrade, or update licenses.
4. Improved responsiveness of application graphs, summary tables, and settings pages
5. Improved timezone handling in the user interface, databases, and JSON transcripts. Date and time now stored in UTC in JSON transcripts.

6. Expanded callback support to include SFTP and Amazon S3, providing the ability to automatically send Text and JSON transcripts and MP3 audio to a URL
7. Additional ASR tags can now be configured in the V-Spark UI
8. Redesigned date range pickers for improved performance and usability
9. Icon added to Settings pages to enable direct navigation to Dashboard view of applications and folders
10. Simplified folder hierarchy is retained in downloaded zip files, preventing filename collisions
11. Increased usability of Activity log: Filter settings are stored as query parameters to enable sharing URLs with other users, and "API" type added to filter log entries related to API calls
12. System administrators now have the ability to restrict users from downloading MP3, JSON, and text transcriptions from the File Details page.
13. Support for Red Hat Enterprise Linux (RHEL) 6
14. Approximate search mode is no longer supported
15. Fixes
 - a. Corrected and improved metadata handling in Voci JSON output. The intermediate `<metadata>` section was removed from the `<client_data>` section to flatten the JSON and simplify client data processing. See the Voci JSON Output Format Guide for more information.
 - b. Application queries using `~n` are now scored correctly even if a turn contains one of the phrase words multiple times

Known Issues in 3.2

The items that are described in this section may affect your use of V-Spark, and are recognized by Voci as items that we expect to address in a future release.

1. Dashboard Displays
 - a. Stat summaries for dates in the past are not updated when files with historical timestamps older than the current day are loaded
 - b. Files uploaded to a folder that is linked to a disabled application or that have not been reprocessed will still appear in the Files View on the application dashboard
2. Unlinking an application from a large folder (greater than 100,000 files) may cause errors where the folder will be unselectable in file editing mode, and where the app is unlinked, but scores may still be visible in certain views
3. Search
 - a. Time-based search
 - i. Searches are turn-based and may return results for phrases that are outside of the requested timeframe, if they are within the same speaker turn
 - ii. Ad-hoc Search snippets in the Files View may display results outside the specified time interval
 - b. The (-) operator for excluding a search tag in ad-hoc search is not supported for file names
 - c. Exceptions to standard File Details Highlighting (in order of severity)
 - i. Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone". Searches using (~t) and (~n) will only be highlighted if the words matched appear in the order specified in the search tag.

- ii. Highlighting may be incomplete for matches marked with emotion
 - iii. Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds
 - iv. Matches for "@"* (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
 - v. Excluded application search phrases may be highlighted on the File Details page. Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.
4. Other
- a. The Trending View may take some time to fully load, especially for longer date ranges
 - b. Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.
5. The transcribe API S3 option can currently only be used in cloud mode

V-Spark 3.1

Limited release with support for Internet licensing

V-Spark 3.0.0

V-Spark version 3.0 features several major updates, especially in the area of application development, including multi-level applications, fully-integrated application scores, and even more powerful search capabilities. Other highlights include the addition of time-based search and text-only uploads of email and chat data.

V-Spark 3.0.0 Release Notes

The capabilities introduced in V-Spark version 3.0 can be organized into several general areas of improvement:

1. Navigation

- a. All pages under Settings will have the same breadcrumb-style navigation as the Dashboard.
- b. Folders and Applications are now displayed on separate pages.
- c. The Application Editor is now accessible from the Applications Page.
- d. Application reprocessing is now controlled directly from the Applications page.

2. Dashboard Graphs & Statistics

- a. Overview and Applications graphs can be set to display historical data (data that is older than thirty days).
- b. Overview Call Statistics table now displays monthly averages.
- c. Call Statistics values will not be updated for data uploaded with historical date times.
- d. Application graphs (both Overview and Daily Chart) now display the number of calls that matched each category.

3. Dashboard Search

- a. Time-based ad-hoc searches can now be conducted.
 - i. Proper format: search phrase followed by ~s or ~e (to indicate the beginning or end of the call), a greater/less than symbol (>/<), and the number of seconds (to indicate which section of the call to search).
 - ii. For example, "thank you < 30" will find all calls where "thank you" was said during the last 30 seconds of the call.
 - iii. Time-based searches terms can be excluded as well, using a minus symbol (-).
 - iv. Time-based search will only be enabled for data uploaded after updating to version 3.0.

- b. New ad-hoc search capabilities have been added so that users can use (^) and (\$) to search for phrases at the very beginning or end of a speaker turn, respectively.

NOTE: These searches will only be enabled for data uploaded after updating to version 3.0.

- c. Application search has been improved so that users can view results for multiple application categories at once, and perform Boolean AND/OR searches on these selections.

4. File Details Page

- a. Application scores are fully integrated into the File Details page. Phrases matching each application category can be highlighted individually.
- b. Silent portions of the call are now displayed as part of the transcript so users can see where silence occurs and for how long. Users will also be able to skip these silent portions when replaying audio.
- c. More efficient display of File Details menu:
 - i. File names are now displayed above the audio player.
 - ii. Controls for turning emotion visualization and silence skipping on/off have been moved above the call metadata for easier access.
- d. The name of the file being viewed will appear within the browser tab for easier navigation between tabs.

5. Enhanced Application Support

a. Key Highlights

- i. Version 3 now includes the following application templates:
 1. Agent Scorecard (improved from version 2)

2. Call Drivers (improved from version 2, formerly Voice of the Customer)
 3. Customer Experience (includes analyses for first call resolution, churn detection, and customer sentiment)
- ii. Multi-level applications now support more in-depth analysis with a category hierarchy of up to 5 categories deep by default.
 1. Currently configured applications will be converted to the new underlying JSON structure.
 2. Application configuration files with the old JSON structure are no longer supported and cannot be uploaded to custom applications.
 3. Ad-hoc application search for data uploaded before updating to version 3 will only work for top-level categories.
 - iii. New applications can be created by copying any existing application within the user's available organizations.

b. Editing

- i. New application search capabilities have been added for more flexible analysis. More detailed explanations of each search character are available in the "V-Spark Application Development Guide".
 1. Underscore (_) now indicates 0 or more words, rather than 1 or more words.
 2. Boolean searches find files that contain both (&) or at least one (|) search phrase.
 3. Distance-based search (~n) finds files containing phrases that appear up to n words apart. Turn-based search (~t) finds files containing phrases that appear within a single speaker turn.
 4. Wildcards for single characters (?) and 0+ characters (*).
 5. Location-based searches find files containing phrases at the very beginning (^) or very end (\$) of the speaker turn.
 6. Searches can be conducted across speaker turns using (-->), e.g. "how can i help you --> billing" will match any calls where Turn A contains "how can i help you" and Turn B contains "billing".
- ii. Custom metadata values may be added to applications as include and exclude filters so that users can better tailor analysis to

their particular company.

- iii. Simultaneous editing of applications by multiple people is now supported.
- iv. All application phrases added using the Application Editor, phrase list upload, or JSON config upload will be validated.
 - 1. Attempts to add invalid phrases will generate an error message.
 - 2. Invalid phrases added using the Application Editor will not be uploaded. Phrases containing errors in syntax will be uploaded, but corrected to the proper syntax.
 - 3. Invalid phrases added using the phrase list upload will not be uploaded, but all remaining phrases in the phrase list file will.
 - 4. Invalid phrases added using a JSON config file will generate an error message listing the first invalid phrase and no portion of the new config file will be uploaded.
- v. Any changes made to an application will be saved automatically as soon as they are entered, rather than at the end of each session.

c. Scoring

- i. Application scores can now be exported for downstream analytics more easily.
 - 1. Application scores for custom date ranges, rather than a single day can be exported from the Agents View.
 - 2. Application scores will automatically be included in JSON transcripts.
 - 3. Reprocessed scores will be available in JSON transcripts downloaded from the File Details page, but not via bulk download from the Settings page.
- ii. Users can select a default scoring type for applications (either Hit/Miss or Coverage) to be displayed on the Dashboard graphs, Agents View, and Summary View).
- iii. Application scoring has changed. Key changes have been summarized below. More detailed information can be found in

"V-Spark 3.5.0 Application Development Guide".

1. Categories at the lowest level of an application (leaf-level categories) will receive a Count score that is equal to the number of phrases that were matched for that category.
2. Upper-level categories will receive both a Coverage score (formerly Average) and a Hit/Miss score (formerly Total).
3. The default value for existing applications will automatically be set to Hit/Miss. Users can switch this to Coverage, but the scores displayed will be equivalent to the old Average scores. For new Coverage scores, users will have to reprocess the application.
4. Exclude phrases in lower-level categories will only cancel out other phrase matches, rather than entire calls.

6. Application Queue

- a. The application queue now includes more detailed status information and a timestamp indicating the time a reprocessing request was received or a folder was unlinked from an application.
- b. The application queue will appear immediately after submitting a reprocessing/delete request.

7. Trending

- a. Trending Editor functionality has been brought in line with Application Editor functionality, with (-) being used to designate exclude phrases, rather than a separate tab.
- b. Users will be able to view a roll-up of trending phrases for custom date ranges, rather than a single day.

8. Logging

- a. Detailed logs of user activity and system status can now be viewed by Admin users, including the following information:
 - i. Creation, deletion, and updates to companies, organizations, folders, and applications
 - ii. Approval of new user accounts as well as updates or deletion of existing user accounts

- iii. Reprocessing of applications
- iv. Uploading of audio/text files for initial transcription User logins and logouts
- b. Users with Admin and Manager permissions can access a log of transcription status for individual audio/text files from the Settings/Folders page.

9. Upload/Download Improvements

- a. Bulk downloads of all file types can now be done for a custom date range.
- b. MP3 files can now be downloaded in bulk from the Settings/Folders page.
- c. JSON representing text data sources such as chat and email can be consumed without audio.
- d. The system has been optimized for greater overall throughput of transcription processing.

10. Fixes

- a. Emotion metadata filters in applications no longer trigger a reprocessing error.
- b. More consistent highlighting in Dashboard search snippet preview and File Details page.
- c. Trending phrases with equivalent frequency counts are now displayed in alphabetical order.

11. For Developers

- a. V-Spark clients who have S3 configured can now pull zipped input directly from S3 via the Voci API.
- b. Backend logging has been improved to rotate logs daily and delete logs older than 30 days.

Known Issues in 3.0

1. Dashboard Displays

- a. Values in the Dashboard Overview Call Statistics table will not be updated for data uploaded with historical datetimes.
- b. Files uploaded to a folder that is linked to a disabled application will still appear in the Files View on the application dashboard.
- c. Files uploaded to a folder before an application has been reprocessed will still appear in the Files View on the application dashboard.

2. Search

- a. Time-based search
 - i. Searches are turn-based and may return results for phrases outside of the specified timeframe, but within the same speaker turn that starts or ends within that timeframe.
 - ii. Search snippets in the Files View may display results outside the specified time interval.
- b. The (-) operator for excluding a search tag in ad-hoc search is not supported for file names.
- c. File Details Highlighting Exceptions* (in order of severity):
 - i. Approximate searches using (~a) will not be highlighted.
 - ii. Searches using (~n) may highlight phrases with the specified number of words in between each of the search terms, rather than between all of the search terms. For example, "get help phone ~4" may highlight phrases with 4 words between "get" and "help" and another 4 words between "help" and "phone"
 - iii. Searches using (~t) and (~n) will only be highlighted if the words matched appear in the order specified in the search tag.
 - iv. Highlighting may be incomplete for matches marked with emotion
 - v. Highlighting may be incomplete for matches with intervening silence tags greater than 2 seconds.

- vi. Matches for "@"* (such as when searching for email addresses) will only be highlighted until the first "." in the email address. Regular expression searches may be used to circumvent this issue.
- vii. Excluded application search phrases may be highlighted on the File Details page.

NOTE: Note that these exceptions only affect highlighting, and that all search results using the above characters will still be accurate.

3. Other

- a. The Trending View may take some time to fully load, especially for longer date ranges.
- b. Users playing longer audio files via Firefox may encounter playback/syncing issues when jumping to different sections of the transcript. This is an external issue caused by the Firefox browser. Users who encounter this issue can simply wait for the browser to finish buffering the file before resuming their search.

V-Spark 2.2.5

V-Spark 2.2.5 was a maintenance release that was given to customers to resolve a few minor problems:

1. Fix "Bookmark not showing up" error
2. Improved audio length limit
3. Resolved ERR_RESPONSE_HEADERS_MULTIPLE_CONTENT_DISPOSITION

V-Spark 2.2.4

V-Spark 2.2.4 was a maintenance release that was given to customers to improve audio handling and to resolve a few problems:

1. (new) Added browser support for Internet Explorer
2. (new) Application configuration files can be uploaded and downloaded for custom applications
3. Enabled upload of metadata filters in application configuration files
4. Emotion-based searches after application reprocessing
5. Improved audio playback
6. Improved password reset notification
7. Improved Trending
8. Increased efficiency of Elasticsearch operations
9. Increased API data flow efficiency

V-Spark 2.2.3 Release Notes

This document provides release notes for V-Spark, describing new features and capabilities that have been introduced in each version of the product. This provides a convenient reference for helping understand the version of the product in which features have been introduced.

V-Spark 2.2.3 was a maintenance release that was given to customers to add several new features and to resolve a few existing problems:

1. (new) Application scores for old audio can be updated/reprocessed when changes are made in the Application Editor.
2. (new) Adds the ability to define custom metadata fields for individual folders and search on custom metadata values.
3. (new) Provides an improved representation of diarization and emotion scores and faster updates during the trending phase.
4. Selection of available models can be changed for a company
5. RequestID assigned to all audio files uploaded in a zip file
6. Admins can select which models are available for use by a company
7. More consistent calculation and display of Diarization scores
8. Dashboard Emotion graphs updated to improve readability
9. Agent/client classifier automatically assigns speaker labels to each channel
10. Increased speed of trending phrase updates
11. JSON and audio can be posted to a user-specified URL (API callback)
12. Ability to configure preferred hostname and port
13. Improved method used to calculate available disk space

V-Spark 2.2.0 Release Notes

This major release introduces V-Spark's new multi-tenant configuration, the ability to customize trending phrases and scorecard applications for personalized analysis, and improved search and emotional intelligence capabilities, processing speed, and ease of use.

1. (new) Multi-tenancy implemented to enable increased privacy
2. (new) New user group classification: Admin (super-admin), Manager (old-admin), User, and Disabled
3. (new) Ability to download all JSON/TEXT files from a folder for a specific date
4. (new) New navigation menu system allows for easy switching between pages
5. (new) User can subscribe to receive reports via email
6. (new) Company usage and limit feature. Admin users can change the limit of each company and monitor company usage from the Accounts tab under Settings
7. (new) Ability to create and edit custom applications using the Application Editor
8. (new) Updated Agent Scorecard template and new Voice of the Customer templates
9. (new) Ability to create custom lists of trending phrases using the Trending Editor
10. (new) Disabled Long/Short Trending phrase lists
11. (new) User Documentation added under Help Menu
12. (new) Improved emotional intelligence calculations
13. (new) Improved Time zone support
14. (fixed) Improved cluster script to run more transcripts faster
15. (fixed) Support callback of JSON data to a URL endpoint via POST

16. Added New database configuration to allow V-Spark to use cloud or remote database
17. Created `/transcribe` API to bypass login for API-based upload of files

Third-party service QA tests

The following table shows which versions of various third-party dependencies were QA tested for a given version of V-Spark. Voci cannot guarantee performance for versions that diverge from those listed here.

V-Spark dependency version tests

V-Spark	Elasticsearch	Database	Redis
4.3.1-1 and later	7.16.2-1	MariaDB 5.5 (EPEL repository)	3.2 (EPEL repository)
4.3.0-1	7.6 and 7.13 (avoid 7.6.2)	MariaDB 5.5 (EPEL repository)	3.2 (EPEL repository)
4.2.0-1	7.6 and 7.13 (avoid 7.6.2)	MariaDB 5.5 (EPEL repository) / AWS RDS MariaDB 10.4.8	3.2 (EPEL repository) / AWS ElastiCache 3.2 and Redis 6.0
4.1.0-1	7.6 and 7.13 (avoid 7.6.2)	MariaDB 5.5 (EPEL repository)	3.2 (EPEL repository)
4.0.2-1	7.6 (avoid 7.6.2)	MariaDB 5.5 (EPEL repository)	3.2 (EPEL repository)