



Private & Confidential

Intelligent Voice 4.14 Release Note

Date: 22nd August 2020

Intelligent Voice Limited
St Clare House
30-33 Minories
London EC3N 1DD
t: +44 20 36 27 26 70
e: info@intelligentvoice.com
w: www.intelligentvoice.com



INTELLIGENT VOICE DISCLAIMER

© 2020 Intelligent Voice. All rights reserved.

Intelligent Voice believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. The information in this publication is provided “as is”. Intelligent Voice makes no representation or warranties of any kind with respect to the information in this publication and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Using, copying, and distribution of any Intelligent Voice software described in this publication requires an applicable software license.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS THE PROPRIETARY AND CONFIDENTIAL INFORMATION OF INTELLIGENT VOICE LIMITED. YOU MAY NOT DISCLOSE, PROVIDE OR MAKE AVAILABLE THIS DOCUMENT, OR ANY INFORMATION CONTAINED IN THIS DOCUMENT, TO ANY THIRD PARTY, WITHOUT THE PRIOR WRITTEN CONSENT OF INTELLIGENT VOICE LIMITED.

The information in this document is provided for use with Intelligent Voice products. No licence, express or implied, to any intellectual property associated with this document or such products is granted by this document.

All Intelligent Voice products described in this document are owned by Intelligent Voice Limited (or those companies that have licensed technology to Intelligent Voice) and are protected by patents, trade secrets, copyrights or other industrial property rights.

The Intelligent Voice products described in this document may still be in development. The final form of each product and release date thereof is at the sole and absolute discretion of Intelligent Voice. Your purchase, license and/or use of Intelligent Voice products shall be subject to Intelligent Voice’s then current sales terms and conditions.

Table of Contents

New Features	4
Inferencing Framework	4
Sentiment detection.....	4
Language Model Builder ships as a separate component.....	4
Dynamic Normalisation	4
Bug Fixes	5
Multilingual Tagger.....	5
Voice Activity Detector	5
SmartTranscript.....	5
Language Model Builder	5
Automatic Speech Recognition	5
Text Chat Import	6
Known Issues – Installation	6
Red Hat Enterprise Linux: Installer fails offline if there is no internet access	6
Installer fails on passwords containing: @&	6
Sporadic issue with permissions on SSL Certificates	6
Red Hat Enterprise Linux: MySQL doesn't come back up after reboot.....	6
Known Issues – Application	7
'Row Size Too Large' in some upgrades	7
Items ingested by URL fill up space as don't get deleted.....	7
Docker containers hanging after upgrade from 4.13.x.....	7
Queries ran by schedulers get stuck.....	8
Psrt Mangler concatenates digits into a single unique word	8

New Features

Inferencing Framework

Release 4.14 is the first Intelligent Voice release to support the use of NVIDIA Triton Inference Server.

Triton is an open source inference serving software that simplifies the deployment of AI models at scale in production.

It supports the deployment of trained AI models from any framework (TensorFlow, TensorRT, PyTorch, ONNX Runtime, or a custom framework).

Models can be loaded from local storage or Google Cloud Platform or AWS S3 on any GPU- or CPU-based infrastructure (cloud, data centre, or edge).

Sentiment detection

The first model to be available for use by the Inferencing Framework is Sentiment. The 'Sentiment Worker' sends inferencing requests to the bespoke trained AI model hosted in the NVIDIA Triton Inference Server.

The model returns a Sentiment score for each TCU (Turn Construction Unit - aka segment or utterance). The score is translated as Sentiment = Positive, Negative or Neutral.

Please note that the Sentiment AI model:

- Is not loaded in Intelligent Voice by default and has to be installed separately.
- Is only made available to customers who have purchased the Sentiment license.

Please contact sales@intelligentvoice.com for details on acquiring this capability.

Language Model Builder ships as a separate component

The Language Model Builder 'LMAWorker' is now shipped as a separate component, with its own lifecycles, separate to the Voice Processing node.

Dynamic Normalisation

Configurable Dynamic Normalisation has been added to the ffmpeg conversion process.

Within out this feature, files with very loud portions can lose segments of speech from the Automatic Speech Recognition, as the very loud portions confused the Voice Activity Detection, leading to the segment not being submitted for Automatic Speech Recognition.

Bug Fixes

Multilingual Tagger

Items	INT-3840, INT-3846, INT-3565, INT-1856
Description	Tags are now correctly produced for <ul style="list-style-type: none"> • Arabic (UAE) [ar-AE] • Catalan • German [de-DE] • Japanese • Korean • Mandarin Chinese [cmn] • Russian • Simplified Chinese [cmn-Hans-TW] • Spanish • Swiss German [de-CH]

Voice Activity Detector

Items	INT-3760
Description	Error Log for Voice Activity Detector Worker fills up if audio file is removed before the Worker has processed the audio file.

SmartTranscript

Items	INT-3529, INT-3704
Description	SmartTranscript now loads correctly on Microsoft Edge browser. Video SmartTranscripts can now play if redactions had been made.

Language Model Builder

Items	INT-3836, INT-3743, INT-3699
Description	The Language Model Builder had not been working for most recently created base models. The Language Model Builder had not been building via API calls. The Language Model Builder could not run if installed on a different server to the Automatic Speech Recognition.

Automatic Speech Recognition

Items	INT-3744
Description	Words were being duplicated in the index when a fusion of two separate models for the same language were being run.

Text Chat Import

Items	INT-3645
Description	The Import Worker for Text Chat was failing.

Known Issues – Installation

Red Hat Enterprise Linux: Installer fails offline if there is no internet access

Items	INT-3732
Description	Offline installation tries to reach some online repository and fails if there is no internet access.
Workaround	None. Offline install without internet access is not supported for Release 4.14

Installer fails on passwords containing: @&

Items	INT-3749, INT-3750
Description	Installer fails if user selects a password with @ vrxServlet will not deploy if user selects a password with &
Workaround	The workaround is to use other non-alphanumeric characters instead of @ and &

Sporadic issue with permissions on SSL Certificates

Items	INT-3762
Description	Sporadic issue with permissions on SSL Certificates after the install completes. jump2web doesn't load after the installation completes, with error visible in catalina.log: SEVERE [main] org.apache.coyote.AbstractProtocol.init Failed to initialize end point associated with ProtocolHandler ["http-apr-8443"] java.lang.Exception: Unable to load certificate key /opt/jumpto/ssl/server-key.pem (error:0200100D:system library:fopen:Permission denied)
Workaround	The workaround is to run the command: chmod og+r /opt/jumpto/ssl/*

Red Hat Enterprise Linux: MySQL doesn't come back up after reboot

Items	INT-3714
Description	MySQL is unable to automatically restart after reboots, including for installation. For an installation this will manifest as Apache starting without MySQL having first started.
Workaround	The workaround is to run the command at the end of the installation: sudo systemctl enable mariadb

Known Issues – Application

'Row Size Too Large' in some upgrades

Items	INT-3605
Description	<p>When upgrading an existing Ubuntu IV System, the installer of the app node throws this error:</p> <pre>database/alter_evite m_table.sqlERROR 1118 (42000) at line 115: Row size too large. The maximum row size for the used table type, not counting BLOBs, is 8126. This inclu des storage overhead, check the manual. You have to change some columns to TEXT or BLOBs</pre> <p>The root cause is that within MariaDB InnoDB has a maximum row size that is roughly equivalent to half of the value of the innodb_page_size system variable.</p>
Workaround	<p>The workaround is to run the following command for each evitem table belonging to each group:</p> <pre>ALTER TABLE tab ROW_FORMAT=DYNAMIC;</pre>

Items ingested by URL fill up space as don't get deleted

Items	INT-3739
Description	<p>Create Item webservice from Http Source fills up space as deletion doesn't occur. When an Item from Http Source is created, the copy is stored within /opt/downloadFiles folder.</p> <p>This folder gets filled up with files that never get removed, even if deleteItem webservice is called, and this might quickly lead to "No Space Left on device", depending on the type of ingestion that the system will have.</p>
Workaround	<p>The workaround is to manually delete the copies stored within the /opt/downloadFiles folder</p>

Docker containers hanging after upgrade from 4.13.x

Items	INT-3852
Description	<p>Sometimes the docker container with the same name doesn't get removed and remains in "docker ps -a" list. For example, message:</p> <pre>docker: Error response from daemon: Conflict. The container name "/TensorRTInferenceServer" is already in use by container</pre>
Workaround	<p>The workaround is to run the commands at the end of the installation:</p> <pre>docker container rm TaggerWorker service TaggerWorker restart</pre>

Queries run by schedulers get stuck

Items	INT-3534
Description	<p>Queries run by a schedule may appear to get stuck as the number of items increase in the system, as the scheduler appears not to be running.</p> <p>The issue was caused by the sql queries not having a LIMIT.</p>
Workaround	Review all the queries ran by the schedulers and add the "LIMIT 1000" instruction. For the queries with UNION, add LIMIT 1000 to each one, wrapping them into () brackets to avoid SQL syntax error.

Psrt Mangler concatenates digits into a single unique word

Items	INT-3606, INT-3703
Description	<p>For example, the statement 'one eighty' would generate both ONE and EIGHTY in the lattice of alternative words, but only "one" would be present in the srt file.</p> <p>The words 'teen' and 'teens' can also be lost between the lattice of alternative words and the srt file when the Psrt Mangler is used.</p>
Workaround	None.