

*Powered by Universal Speech Solutions LLC*



# Azure SS Plugin

## Usage Guide

---

Revision: 11

Created: November 11, 2017

Last updated: August 28, 2020

Author: Arsen Chaloyan

# Table of Contents

1	Overview.....	4
1.1	Installation .....	4
1.2	Applicable Versions.....	4
2	Supported Features.....	5
2.1	MRCP Methods .....	5
2.2	MRCP Events .....	5
2.3	MRCP Header Fields .....	5
2.4	Speech Data .....	5
3	Supported Voices .....	6
4	Configuration Format.....	7
4.1	Document.....	7
4.2	Synthesis Settings .....	8
4.3	Waveform Manager .....	9
4.4	SDR Manager .....	11
4.5	Monitoring Agent .....	12
4.6	Usage Change Handler .....	13
4.7	Usage Refresh Handler .....	13
4.8	License Server.....	14
5	Configuration Steps .....	16
5.1	Using Default Configuration.....	16
5.2	Specifying Synthesis Language .....	16
5.3	Specifying Sampling Rate.....	16
5.4	Specifying Voice Parameters .....	16
5.5	Specifying Prosody Parameters .....	17
5.6	Specifying Speech Data .....	17
5.7	Maintaining Waveforms .....	17
5.8	Maintaining Synthesis Details Records .....	18
5.9	Using Cache.....	19
6	Monitoring Usage Details .....	21
6.1	Log Usage.....	21
6.2	Update Usage.....	21
6.3	Dump Channels.....	22
7	Usage Examples.....	23
7.1	SSML.....	23
7.2	Plain Text.....	23

8 Sequence Diagram ..... 25

9 References ..... 26

    9.1 Microsoft Azure ..... 26

    9.2 Specifications ..... 26

# 1 Overview

This guide describes how to configure and use the Microsoft Azure Speech Synthesis (SS) plugin to the UniMRCP server. The document is intended for users having a certain knowledge of Microsoft Azure Speech API and UniMRCP.



## 1.1 Installation

For installation instructions, use one of the guides below.

- RPM Package Installation (Red Hat / Cent OS)
- Deb Package Installation (Debian / Ubuntu)

## 1.2 Applicable Versions

Instructions provided in this guide are applicable to the following versions.



UniMRCP 1.5.0 and above

UniMRCP Azure SS Plugin 1.0.0 and above

# 2 Supported Features

This is a brief check list of the features currently supported by the UniMRCP server running with the Azure SS plugin.

## 2.1 MRCP Methods

- ✓ SPEAK
- ✓ STOP
- ✓ PAUSE
- ✓ RESUME
- ✓ BARGE-IN-OCCURRED
- ✓ SET-PARAMS
- ✓ GET-PARAMS

## 2.2 MRCP Events

- ✓ SPEECH-MARKER
- ✓ SPEAK-COMPLETE

## 2.3 MRCP Header Fields

- ✓ Kill-On-Barge-In
- ✓ Completion-Cause
- ✓ Voice-Gender
- ✓ Voice-Name
- ✓ Prosody-Rate
- ✓ Prosody-Volume
- ✓ Speech-Language
- ✓ Logging-Tag
- ✓ Cache-Control

## 2.4 Speech Data

- ✓ Plain text (text/plain)
- ✓ SSML (application/ssml+xml or application/synthesis+ssml)

# 3 Supported Voices

All the supported voices are stored in the configuration file *umsazurevoices.xml* located in the directory */opt/unimrcp/conf*. The configuration file is synched with the actual set of voices supported by Microsoft Azure Speech API listed in the following page:

[Supported locales and voice fonts](#)

# 4 Configuration Format

The configuration file of the Azure SS plugin is located in `/opt/unimrcp/conf/umsazuress.xml`. The configuration file is written in XML.

## 4.1 Document

The root element of the XML document must be `<umsazuress>`.

### Attributes

Name	Unit	Description
<b>license-file</b>	File path	Specifies the license file. File name may include patterns containing '*' sign. If multiple files match the pattern, the most recent one gets used.
<b>subscription-key-file</b>	File path	Specifies the Microsoft subscription key file to use. File name may include patterns containing '*' sign. If multiple files match the pattern, the most recent one gets used.

### Parent

None.

### Children

Name	Unit	Description
<b>&lt;synth-settings&gt;</b>	String	Specifies synthesis parameters.
<b>&lt;waveform-manager&gt;</b>	String	Specifies parameters of the waveform manager. Available since Azure SS 1.4.0.
<b>&lt;sdr-manager&gt;</b>	String	Specifies parameters of the Synthesis Details Record (SDR) manager. Available since Azure SS 1.4.0.
<b>&lt;monitoring-agent&gt;</b>	String	Specifies parameters of the monitoring manager.
<b>&lt;license-server&gt;</b>	String	Specifies parameters used to connect to the license server. The use of the license server is optional.

## Example

This is an example of a bare document.

```
<umsazuress license-file="umsazuress_*.lic" subscription-key-  
file="cognitive.subscription.key">  
</umsazuress>
```

## 4.2 Synthesis Settings

This element specifies synthesis parameters.

### Attributes

Name	Unit	Description
<b>language</b>	String	Specifies the default language to use, if not set by the client.
<b>voice-name</b>	String	Specifies the default voice name. Can be overridden by client. Available since Azure SS 1.3.0.
<b>voice-gender</b>	String	Specifies the default voice gender. Can be overridden by client. Available since Azure SS 1.3.0.
<b>prosody-pitch</b>	String	Specifies the default prosody pitch. Use either labels (x-low, low, medium, high, x-high) or an absolute value (600 Hz) or relative changes from the default pitch (+80 Hz or -20 Hz). Can be overridden by client in SSML. Available since Azure SS 1.11.0.
<b>prosody-contour</b>	String	Specifies the default prosody contour. Can be overridden by client in SSML. Available since Azure SS 1.11.0.
<b>prosody-range</b>	String	Specifies the default prosody range. Use either labels (x-low, low, medium, high, x-high) or an absolute value (600 Hz) or relative changes from the default pitch (+80 Hz or -20 Hz). Can be overridden by client in SSML. Available since Azure SS 1.11.0.
<b>prosody-rate</b>	String	Specifies the default prosody rate. Use either labels (x-slow, slow, medium, fast, x-fast) or relative changes from the default rate as a percentage (+10% or -20%) or numeric values as a multiplier where 1 indicates no change. Can be overridden by client. Available since Azure SS 1.11.0.
<b>prosody-duration</b>	String	Specifies the default prosody duration. Can be overridden



		by client in SSML. Available since Azure SS 1.11.0.
<b>prosody-volume</b>	String	Specifies the default prosody volume. Use either labels (silent, x-soft, soft, medium, loud, x-loud) or relative changes from the default volume as a percentage (+10% or -20%) or numeric values in the range from 0 to 100. Can be overridden by client. Available since Azure SS 1.11.0.
<b>bypass-ssml</b>	Boolean	Specifies whether transparently bypass or normalize received SSML content before sending it to the service for synthesis.
<b>auth-validation-period</b>	Integer	Specifies a period in seconds used to re-validate access token based on subscription key. The lifetime of retrieved access token is set to 10 min by Microsoft.
<b>http-proxy</b>	String	Specifies the URI of HTTP proxy, if used. Available since Azure SS 1.6.0.
<b>caching</b>	Boolean	Specifies whether to enable caching of synthesized waveforms. Available since Azure SS 1.11.0.

### Parent

<umsazuress>

### Children

None.

### Example

This is an example of synthesis parameters.

```
<synth-settings
  language="en-US"
  bypass-ssml="false"
  auth-validation-period="480"
/>
```

## 4.3 Waveform Manager

This element specifies parameters of the waveform manager.

### Availability

>= Azure SS 1.4.0.

## Attributes

Name	Unit	Description
<b>save-waveforms</b>	Boolean	Specifies whether to save waveforms or not.
<b>purge-existing</b>	Boolean	Specifies whether to delete existing records on start-up.
<b>max-file-age</b>	Time interval [min]	Specifies a time interval in minutes after expiration of which a waveform is deleted. Set 0 for infinite.
<b>max-file-count</b>	Integer	Specifies the max number of waveforms to store. If reached, the oldest waveform is deleted. Set 0 for infinite.
<b>waveform-folder</b>	Dir path	Specifies a folder the waveforms should be stored in.
<b>file-prefix</b>	String	Specifies a prefix used to compose the name of the file to be stored. Defaults to 'umsazuress-', if not specified.
<b>use-logging-tag</b>	Boolean	Specifies whether to use the MRCP header field Logging-Tag, if present, to compose the name of the file to be stored. Available since Azure SS 1.11.0.

## Parent

<umsazuress>

## Children

None.

## Example

The example below defines a typical utterance manager having the default parameters set.

```
<waveform-manager
  save-waveforms="false"
  purge-existing="false"
  max-file-age="60"
  max-file-count="100"
  waveform-folder=""
/>
```

## 4.4 SDR Manager

This element specifies parameters of the Synthesis Details Record (SDR) manager.

### Availability

>= Azure SS 1.4.0.

### Attributes

Name	Unit	Description
<b>save-records</b>	Boolean	Specifies whether to save recognition details records or not.
<b>purge-existing</b>	Boolean	Specifies whether to delete existing records on start-up.
<b>max-file-age</b>	Time interval [min]	Specifies a time interval in minutes after expiration of which a record is deleted. Set 0 for infinite.
<b>max-file-count</b>	Integer	Specifies the max number of records to store. If reached, the oldest record is deleted. Set 0 for infinite.
<b>record-folder</b>	Dir path	Specifies a folder to store recognition details records in. Defaults to <code>\${UniMRCPIInstallDir}/var</code> .
<b>file-prefix</b>	String	Specifies a prefix used to compose the name of the file to be stored. Defaults to 'umsazuress-', if not specified.
<b>use-logging-tag</b>	Boolean	Specifies whether to use the MRCP header field Logging-Tag, if present, to compose the name of the file to be stored. Available since Azure SS 1.11.0.

### Parent

<umsazuress>

### Children

None.

## Example

The example below defines a typical utterance manager having the default parameters set.

```
<sdr-manager
  save-records="false"
  purge-existing="false"
  max-file-age="60"
  max-file-count="100"
  waveform-folder=""
/>
```

## 4.5 Monitoring Agent

This element specifies parameters of the monitoring agent.

### Attributes

Name	Unit	Description
<b>refresh-period</b>	Time interval [sec]	Specifies a time interval in seconds used to periodically refresh usage details. See <usage-refresh-handler>.

### Parent

```
<umsazuress>
```

### Children

```
<usage-change-handler>
<usage-refresh-handler>
```

## Example

The example below defines a monitoring agent with usage change and refresh handlers.

```
<monitoring-agent refresh-period="60">
  <usage-change-handler>
    <log-usage enable="true" priority="NOTICE"/>
  </usage-change-handler>
  <usage-refresh-handler>
    <dump-channels enable="true" status-file="umsazuress-channels.status"/>
  </usage-refresh-handler >
```

```
</monitoring-agent>
```

## 4.6 Usage Change Handler

This element specifies an event handler called on every usage change.

### Attributes

None.

### Parent

```
<monitoring-agent>
```

### Children

```
<log-usage>  
<update-usage>  
<dump-channels>
```

### Example

This is an example of the usage change event handler.

```
<usage-change-handler>  
  <log-usage enable="true" priority="NOTICE"/>  
  <update-usage enable="false" status-file="umsazuress-usage.status"/>  
  <dump-channels enable="false" status-file="umsazuress-channels.status"/>  
</usage-change-handler>
```

## 4.7 Usage Refresh Handler

This element specifies an event handler called periodically to update usage details.

### Attributes

None.

### Parent

```
<monitoring-agent>
```

### Children

```
<log-usage>  
<update-usage>  
<dump-channels>
```

## Example

This is an example of the usage change event handler.

```
<usage-refresh-handler>  
  <log-usage enable="true" priority="NOTICE"/>  
  <update-usage enable="false" status-file="umsazuress-usage.status"/>  
  <dump-channels enable="false" status-file="umsazuress-channels.status"/>  
</usage-refresh-handler>
```

## 4.8 License Server

This element specifies parameters used to connect to the license server.

### Attributes

Name	Unit	Description
<b>enable</b>	Boolean	Specifies whether the use of license server is enabled or not. If enabled, the license-file attribute is not honored.
<b>server-address</b>	String	Specifies the IP address or host name of the license server.
<b>certificate-file</b>	File path	Specifies the client certificate used to connect to the license server. File name may include patterns containing a '*' sign. If multiple files match the pattern, the most recent one gets used.
<b>ca-file</b>	File path	Specifies the certificate authority used to validate the license server.
<b>channel-count</b>	Integer	Specifies the number of channels to check out from the license server. If not specified or set to 0, either all available channels or a pool of channels will be checked based on the configuration of the license server.
<b>http-proxy-address</b>	String	Specifies the IP address or host name of the HTTP proxy server, if used. Available since 1.10.0.
<b>http-proxy-port</b>	Integer	Specifies the port number of the HTTP proxy server, if used. Available since 1.10.0.

## Parent

<umsazuress>

## Children

None.

## Example

The example below defines a typical configuration which can be used to connect to a license server located, for example, at 10.0.0.1.

```
<license-server
  enable="true"
  server-address="10.0.0.1"
  certificate-file="unilic_client_*.cert"
  ca-file="unilic_ca.crt"
/>
```

For further reference to the license server, visit

<http://unimrcp.org/licserver>

# 5 Configuration Steps

This section outlines common configuration steps.

## 5.1 Using Default Configuration

The default configuration should be sufficient for the general use.

## 5.2 Specifying Synthesis Language

Synthesis language can be specified by the client per MRCP session by means of the header field *Speech-Language* set in a *SET-PARAMS* or *SPEAK* request, or inline in the SSML data. Otherwise, the parameter *language* set in the configuration file *umsazureess.xml* is used. The parameter defaults to *en-US*.

For supported languages and their corresponding codes, visit the following link.

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech/api-reference-rest/supportedlanguages>

## 5.3 Specifying Sampling Rate

Sampling rate is determined based on the SDP negotiation. Refer to the configuration guide of the UniMRCP server on how to specify supported encodings and sampling rates to be used in communication between the client and server. Either 8 or 16 kHz can be used by Microsoft Azure Speech API for synthesis.

## 5.4 Specifying Voice Parameters

### Global Settings

The default voice name and gender can be specified from the configuration file *umsazureess.xml* using the *voice-name* and *voice-gender* attributes of the *synth-settings* element. This functionality is available since Azure SS 1.3.0 release.

### MRCP Header Fields

The voice name and gender can be specified by the MRCP client in *SET-PARAMS* and *SPEAK* requests.

- Voice-Name

This is an optional parameter indicating the name of the voice to use for synthesis.

- Voice-Gender

This is an optional parameter indicating the preferred gender of the voice to use for synthesis, which can be set to either *male* or *female*.



## 5.5 Specifying Prosody Parameters

The following prosody parameters can be specified by the MRCP client in *SET-PARAMS* and *SPEAK* requests.

- Prosody-Rate

This is an optional parameter indicating the speaking rate, which can be set to one of the following labels: *x-slow*, *slow*, *medium*, *fast*, *x-fast*, *default*.

- Prosody-Volume

This is an optional parameter indicating the speaking volume, which can be set to one of the following labels: *silent*, *x-soft*, *soft*, *medium*, *loud*, *x-loud*, *default*.

## 5.6 Specifying Speech Data

Speech data can be specified by the MRCP client in *SPEAK* requests using one of the following content types:

- plain/text
- application/ssml+xml (or application/synthesis+ssml)

## 5.7 Maintaining Waveforms

Collection of waveforms is not required for regular operation and is disabled by default. However, enabling this functionality allows to save synthesized speech received from the Microsoft Azure Speech service and later listen to them offline.

The relevant settings can be specified via the element *waveform-manager*.

- save-waveforms

Utterances can optionally be recorded and stored if the configuration parameter *save-waveforms* is set to true.

- purge-existing

This parameter specifies whether to delete existing waveforms on start-up.

- max-file-age

This parameter specifies a time interval in minutes after expiration of which a waveform is deleted. If set to 0, there is no expiration time specified.

- max-file-count

This parameter specifies the maximum number of waveforms to store. If the specified number is reached, the oldest waveform is deleted. If set to 0, there is no limit specified.

- waveform-folder

This parameter specifies a path to the directory used to store waveforms in. The directory defaults to `${UniMRCPIInstallDir}/var`.

## 5.8 Maintaining Synthesis Details Records

Collection of synthesis details records (SDR) is not required for regular operation and is disabled by default. However, enabling this functionality allows to store details of each synthesis attempt in a separate file and analyze them later offline. The SDRs are stored in the JSON format.

The relevant settings can be specified via the element `sdr-manager`.

- `save-records`

This parameter specifies whether to save synthesis details records or not.

- `purge-existing`

This parameter specifies whether to delete existing records on start-up.

- `max-file-age`

This parameter specifies a time interval in minutes after expiration of which a record is deleted. If set to 0, there is no expiration time specified.

- `max-file-count`

This parameter specifies the maximum number of records to store. If the specified number is reached, the oldest record is deleted. If set to 0, there is no limit specified.

- `record-folder`

This parameter specifies a path to the directory used to store records in. The directory defaults to `${UniMRCPIInstallDir}/var`.

The following is the content of a sample SDR.

```
{ "synth-details-record": {  
  "datetime": "2019-01-19 12:57:44",  
  "language": "en-US",  
  "voice-name": "",  
  "sampling-rate": "8000 Hz",  
  "codec": "PCMU",  
  "data-length": "18150 bytes",  
  "start-of-streaming-ts": "282 ms",  
  "completion-ts": "2550 ms",  
  "completion-cause": "normal",  
}}
```

where the stored attributes are:

- `datetime`

This attribute denotes the date and time captured when the corresponding MRCP SPEAK request is received.

- language

This attribute denotes the speech language used with the request.

- voice-name

This attribute denotes the voice name used with the request.

- sampling-rate

This attribute denotes the sampling rate used with the request.

- codec

This attribute denotes the codec of synthesized speech received from the service.

- data-length

This attribute denotes the number of bytes of synthesized speech received from the service.

- start-of-streaming-ts

This attribute denotes a time interval in milliseconds, elapsed since initiation of the request, captured when streaming to the MRCP client is started. This attribute also denotes the HTTP response time of the service.

- completion-ts

This attribute denotes a time interval in milliseconds, elapsed since initiation of the request, captured upon completion of the request (SPEAK-COMplete is sent).

- completion-cause

This attribute denotes the completion cause of the request.

## 5.9 Using Cache

Since Azure SS 1.11.0, synthesized waveforms can be stored and re-used for consecutive speech synthesis requests, when applicable. In order to use this functionality, the attribute *caching* of the element *synth-settings* must be set to *true*. The attribute defaults to *false*.

The lifetime and size of cached records are controlled by the attributes *max-file-age* and *max-file-count* of the element *waveform-manager*.

The cached records are persistent and populated on initial loading, unless the attribute *purge-existing* of the element *waveform-manager* is set to *true*.

The following speech synthesis parameters are observed while searching for a cached record.

- language
- voice-name

- voice-gender
- sampling-rate
- prosody-pitch
- prosody-contour
- prosody-range
- prosody-rate
- prosody-duration
- prosody-volume
- content

The following cache control directives are observed while searching for a cached record.

- max-age
- min-fresh

The cache control directives can be specified by the client per individual speech synthesis request via the MRCP header field *Cache-Control*. By default, no cache control directives are applied.

# 6 Monitoring Usage Details

The number of in-use and total licensed channels can be monitored in several alternate ways. There is a set of actions which can take place on certain events. The behavior is configurable via the element *monitoring-agent*, which contains two event handlers: *usage-change-handler* and *usage-refresh-handler*.

While the *usage-change-handler* is invoked on every acquisition and release of a licensed channel, the *usage-refresh-handler* is invoked periodically on expiration of a timeout specified by the attribute *refresh-period*.

The following actions can be specified for either of the two handlers.

## 6.1 Log Usage

The action *log-usage* logs the following data in the order specified.

- The number of currently in-use channels.
- The maximum number of channels used concurrently. Available since Azure SS 1.4.0.
- The total number of licensed channels.

The following is a sample log statement, indicating 0 in-use, 0 max-used and 2 total channels.

```
[NOTICE] BINGSS Usage: 0/0/2
```

## 6.2 Update Usage

The action *update-usage* writes the following data to a status file *umsazure-usage.status*, located by default in the directory `${UniMRCPIInstallDir}/var/status`.

- The number of currently in-use channels.
- The maximum number of channels used concurrently. Available since Azure SS 1.4.0.
- The total number of licensed channels.
- The current status of the license permit.
- The license server alarm. Set to *on*, if the license server is not available for more than one hour; otherwise, set to *off*. This parameter is maintained only if the license server is used. Available since Azure SS 1.5.0.

The following is a sample content of the status file.

```
in-use channels: 0
```

```
max used channels: 0
total channels: 2
license permit: true
licserver alarm: off
```

## 6.3 Dump Channels

The action *dump-channel* writes the identifiers of in-use channels to a status file *umsazuress-channels.status*, located by default in the directory `${UniMRCPIInstallDir}/var/status`.

# 7 Usage Examples

## 7.1 SSML

This examples demonstrates how to perform speech synthesis by using a SPEAK request with an SSML content.

C->S:

```
MRCP/2.0 309 SPEAK 1
Channel-Identifier: 4dde51f37d1a9546@speechsynth
Content-Type: application/ssml+xml
Voice-Age: 28
Content-Length: 163

<?xml version="1.0"?>
<speak version="1.0" xml:lang="en-US" xmlns="http://www.w3.org/2001/10/synthesis">
  <p>
    <s>Welcome to Uni MRCP.</s>
  </p>
</speak>
```

S->C:

```
MRCP/2.0 83 1 200 IN-PROGRESS
Channel-Identifier: 4dde51f37d1a9546@speechsynth
```

S->C:

```
MRCP/2.0 122 SPEAK-COMPLETE 1 COMPLETE
Channel-Identifier: 4dde51f37d1a9546@speechsynth
Completion-Cause: 000 normal
```

## 7.2 Plain Text

This examples demonstrates how to perform speech synthesis by using a SPEAK request with a plain text content.

C->S:

MRCP/2.0 155 SPEAK 1  
Channel-Identifier: 85667d0efbf95345@speechsynth  
Content-Type: text/plain  
Voice-Age: 28  
Content-Length: 20

Welcome to Uni MRCP.

S->C:

MRCP/2.0 83 1 200 IN-PROGRESS  
Channel-Identifier: 85667d0efbf95345@speechsynth

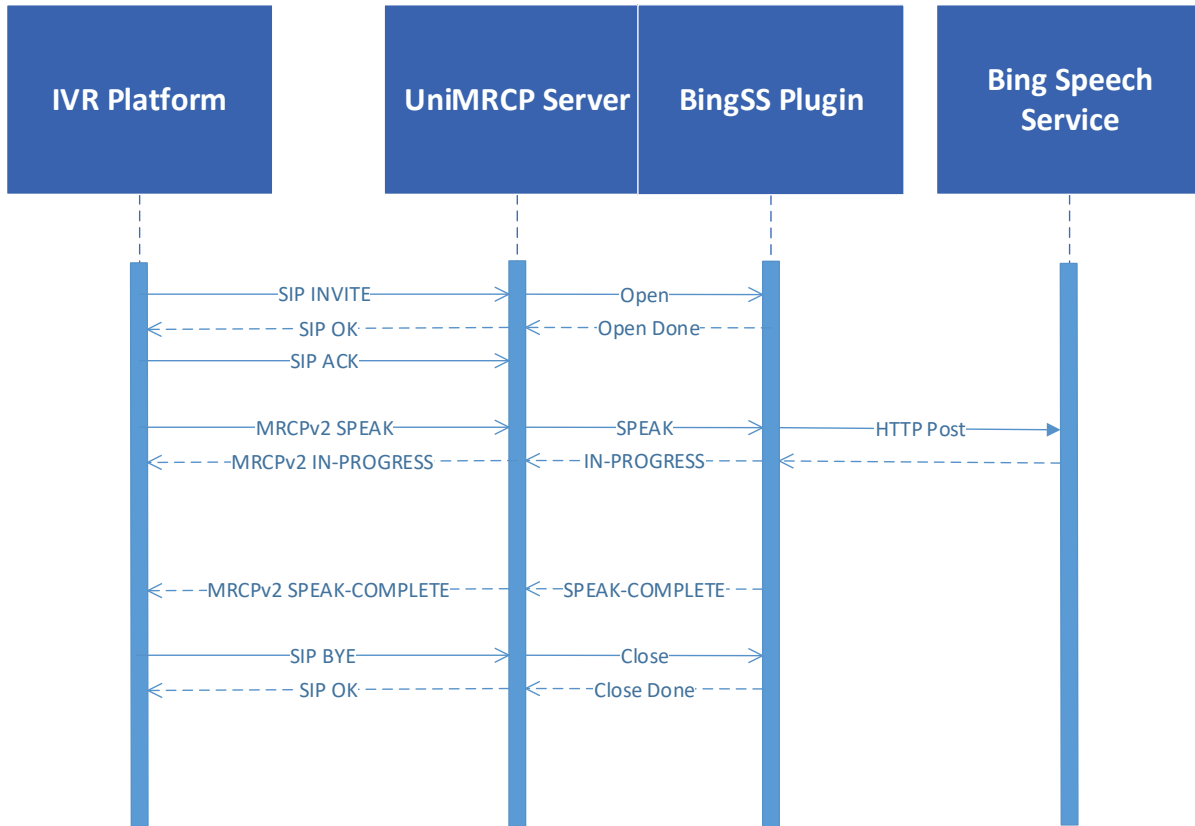
S->C:

MRCP/2.0 122 SPEAK-COMplete 1 COMPLETE  
Channel-Identifier: 85667d0efbf95345@speechsynth  
Completion-Cause: 000 normal



# 8 Sequence Diagram

The following sequence diagram outlines common interactions between all the main components involved in a typical synthesis session performed over MRCPv2.



# 9 References

## 9.1 Microsoft Azure

- [Text to Speech API](#)
- [Authentication](#)

## 9.2 Specifications

- [Speech Synthesizer Resource](#)
- [SSML](#)