

# Construct API



**Advanced Simulation Technology inc.**

---

500 A Huntmar Park Drive • Herndon, Virginia 20170 U.S.A.

Tel. (703)471-2104 • Fax. (703)471-2108

[www.asti-usa.com](http://www.asti-usa.com)

# Contents

<b>Contents</b>	<b>1</b>
<b>1 Application Examples</b>	<b>2</b>
<b>2 Scenarios</b>	<b>2</b>
2.1 Entities . . . . .	3
2.2 Interactions . . . . .	4
2.3 Sounds . . . . .	5
2.4 Radio Effects . . . . .	5
2.5 Language Models . . . . .	5
<b>3 Construct Settings</b>	<b>7</b>
<b>4 Natural Language Parsers</b>	<b>7</b>
<b>5 Artificial Intelligence Markup Language (AIML)</b>	<b>7</b>
<b>6 Runtime Entities</b>	<b>7</b>
<b>7 Utterances</b>	<b>8</b>
<b>8 Streaming Status and Events</b>	<b>8</b>
8.1 Entity Update . . . . .	10
8.2 Entity Delete . . . . .	10
8.3 Speech Event . . . . .	10
8.4 Speech-Finished Event . . . . .	10
8.5 Speech Recognition Event . . . . .	10

The Construct API allows developers to dynamically add sounds and voiced radio transmissions to a simulation environment. HTTP clients are able to interact with Construct by making standard HTTP requests to the Voisus server. Simulation entities, interactions, and sounds are represented as JSON resources that can be manipulated by clients to create highly customized audio and radio environments.

The Construct API is part of the larger Voisus Server API. Visit the Voisus Server API Documentation<sup>1</sup> first for an overview of the API including information on REST, HTTP, and JSON.

Developers have access to all Construct functionality using the HTTP API:

- Create entities representing vehicles or human characters
- Attach to entities on the simulation network
- Trigger voice transmissions on simulated radios
- Script sequences of radio calls between entities
- Add realistic sound effects to radios
- Position entity speech in 3D game environments
- Create face-to-face conversations between game avatars and human players
- Synthesize variable speech using text-to-speech (TTS)
- Add automated speech recognition and intelligent behaviors to entities

## 1 Application Examples

- To create a radio interaction between two Entities, create two Entity resources and an Interaction resource that specifies the speech for each radio transmission.
- Trigger dynamic, on-the-fly radio calls from Entities by creating new Utterance resources.
- Add jet engine noise to radio calls from aircraft by setting the `background` attribute on an Utterance.
- Receive speech recognition events from Entities using the streaming status feed.

## 2 Scenarios

Most Construct resources are contained within a Voisus Scenario. This includes resources for Entities, Interactions, Sounds, Radio Effects, and Language Models. When the Scenario runs, Entities are spawned and Interactions start running. Scenarios themselves are also created, deleted, and run using the API. Consult the Voisus Server API Documentation<sup>2</sup> for more information on scenario management.

---

<sup>1</sup>voisus\_server\_api.pdf

<sup>2</sup>voisus\_server\_api.pdf

## 2.1 Entities

`api/scenarios/scenario_id/entities/`

Construct Entities are used to add speech and radio capabilities to objects in simulation environments. These Entities may correspond and be synchronized with Entities in external simulation systems, or they may modeled solely by this Voisus Server.

Entities are created using either the web interface or this HTTP API. The Entity resource shown below is stored in a Voisus Scenario and will cause a runtime Entity to be created and destroyed each time the Scenario is installed and uninstalled, respectively. A separate section of the HTTP API is available for interacting with the runtime representation of an Entity, including retrieving speech recognition results and triggering text-to-speech transmissions.

When creating an Entity in a Scenario, there are a number of optional parameters and only a small number of required parameters. Some parameters are links to other Scenario resources, including `net`, `domain`, `language_model`, and `radio_effects`. In these cases, the parameter value is the URI of the referenced resource.

To create an Entity with basic radio functionality, the `net` and `domain` properties must be set. A TTS voice will be chosen randomly if one is not assigned.

The JSON below represents an Entity named “Broadsword 11”.

Variable	Definition
<code>domain</code>	Link to a Radio domain resource (controls DIS exercise)
<code>net</code>	Link to Net resource for to enable radio communication
<code>behavior</code>	Behavior that will control the entity (optional)
<code>language_model</code>	Language model for speech recognition (optional)
<code>radio_effects</code>	Link to a Radio Effect resource (optional)
<code>aiml</code>	Name of Artificial Intelligence Markup Language (AIML) file to use (optional)
<code>parser</code>	Name of speech parser to use for Natural Language Understanding (NLU)
<code>attributes</code>	Arbitrary attributes for use in speech and behaviors
<code>listen</code>	Enables/disables listening with speech recognition
<code>radio_enabled</code>	Enables/disables radio communications
<code>earshot_enabled</code>	Enables/disables face-to-face communication in 3D environments
<code>marking</code>	String used for DIS entity-attach (optional)
<code>pitch_shift</code>	Pitch shift amount; 0.9 low, 1 normal, 1.1 high
<code>tts_voice</code>	Text-to-speech (TTS) voice
<code>tts_rate</code>	Speech rate (1-9)
<code>tts_volume</code>	Speech volume (1-9)
<code>breakpoints</code>	Breakpoints for inspecting Entity behavior execution
<code>default_position</code>	Initial (X,Y,Z) position of the Entity in geocentric coordinates

## 2.2 Interactions

`api/scenarios/<scenario_id>/interactions/`

Interactions cause one or more Entities to speak. Each Interaction specifies a list of speech actions to execute, one after the other. As an example, a “IED Attack” interaction might involve a series of radio calls between the attacked vehicle and command personnel.

The JSON below represents an interaction named “Blackjack Under Attack” with two TTS radio calls from different Entities. Since `enabled` is true and `time` is 0, the interaction will start immediately. There will be a three second delay between radio calls because `action_delay` is 3.

Variable	Definition
<code>action_delay</code>	Default delay between actions (seconds)
<code>enabled</code>	Enables/disables the interaction
<code>variables</code>	Arbitrary variables; used during speech synthesis
<code>holdoff</code>	Delay between loops (seconds)
<code>limit</code>	Number of times the interaction should run (0=forever)
<code>time</code>	Interaction start time (seconds)
<code>time_type</code>	Enumeration controlling the start time [“rel”, “utc”]
<code>vbs2_trigger</code>	Optional VBS scripting expression to trigger this interaction
<code>actions</code>	List of actions to be executed (see table below)

The `time_type` property determines the meaning of the specified start `time`. A relative time type of “rel” means the `time` is relative to when the scenario begins or when the Interaction is created. The UTC time type of “utc” means the `time` is a wall clock time, specified in seconds since the Unix epoch. If an interaction is created with a UTC time in the past, the interaction will not execute. The Visus server system clock is used when determining when an Interaction should begin. Confirm this clock is synchronized with external simulation system clocks to ensure proper time-alignment of simulation events.

### Actions

Each action in an Interaction has the following format. The action `type` defines whether TTS or prerecorded speech will be used. In either case, the `background` parameter can be used to add environmental ambiance or sound effects like gunfire into the speech stream.

Variable	Definition
<code>entity</code>	Link to an Entity resource (required)
<code>type</code>	Starting offset in sound file (samples)
<code>predelay</code>	Delay before this action executes (seconds)
<code>text</code>	Action type enumeration [“tts”, “sound”]
<code>sound</code>	Link to a Sound resource
<code>background</code>	Link to a Sound resource to play in the background

If a `type` of “tts” is specified, the `sound` property is ignored. If `text` is specified for a “sound” action, the text is understood to be a transcript of the sound file, and is shown as such on the Construct status webpage.

## 2.3 Sounds

`api/scenarios/<scenario_id>/sounds/`

Sound resources are used to play a sound effect or recorded speech in Construct. Each Sound selects a sound file by name and a segment within the file to play. Sounds may be reused across multiple Entities and Interactions. Sounds are separate from, but reference, sound files. Upload sound files to reference using the web interface.

The JSON below represents a looping engine sound named “Diesel Engine”. The `play_count` value of 0 means the sound should loop indefinitely.

Variable	Definition
<code>file</code>	Sound file name (e.g. “explosion.wav”)
<code>offset</code>	Starting offset in sound file (samples)
<code>length</code>	Length of section to play; 0=all (samples)
<code>gain</code>	Playback gain; 1.0 means no volume adjustment
<code>play_count</code>	Number of times to play when triggered (0=forever)
<code>play_all</code>	If true, Sound will always play to the end before stopping
<code>random</code>	If true, each playback begins at a random position
<code>loop_start</code>	Start position of loop [0-1] (0=beginning, 1=end)
<code>loop_end</code>	End position of loop [0-1] (0=beginning, 1=end)
<code>loop_delay</code>	Delay between loops (seconds)
<code>transcript</code>	Optional text transcript for a sound file containing speech. Transcript is shown when an Entity “speaks” using this Sound.

In many cases setting only the `file` property and using the defaults for the other properties will yield the expected results. The Sound defaults to playing the entire sound file once at normal volume.

## 2.4 Radio Effects

`api/scenarios/<scenario_id>/radio_effects/`

Radio Effects customize the sound of Entity radio transmissions by adding distortion, noise, and band-limiting audio effects. Link a Radio Effect to one or more Entities by setting the `radio_effects` property in an Entity to the Radio Effect “self” URI.

The JSON below represents a Radio Effect named “High Distortion”.

## 2.5 Language Models

`api/scenarios/<scenario_id>/language_models/`

Variable	Definition
<code>distortion_mode</code>	Switches between distortion, overdrive, and none (0=none, 1=overdrive, 2=distortion)
<code>distortion_mix</code>	Adjusts the ratio of distorted to clean signal to be transmitted (0=clean only, 1=distorted only)
<code>distortion_gain</code>	Gain applied to the signal before entering the distortion effect. Higher gains yield a more distorted sound.
<code>limiter_threshold</code>	Threshold in decibels (dB) that limits the peak signal strength. More negative values yield quieter transmissions.
<code>noise_gain</code>	Gain controlling the volume of noise added to the transmission
<code>gain</code>	Gain applied to the signal just before transmission
<code>noise_color</code>	Type of added noise: “white”, “pink”, or “brown”
<code>hp_freq</code>	Highpass filter frequency in Hertz (Hz)
<code>lp_freq</code>	Lowpass filter frequency in Hertz (Hz)

Language Model resources select a speech recognition model and specify custom parameters for the recognition engine. To use a Language Model, set the `language_model` property of an Entity to the Language Model “self” URI.

Statistical language models or grammars must be pre-installed on the Voisus server for Speech Recognition resources to be functional. Installed models are referenced by name in this resource. In the example below, “atc11” is the name of a model provided by ASTi. Currently installed models are listed in the model dropdown on the Language Models webpage. Contact ASTi for more information on available language models.

The JSON below represents an ASR configuration named “ATC-ASR”.

Variable	Definition
<code>asr_model</code>	Name of the language model on disk
<code>grammar</code>	Grammar text in BNF format
<code>parameters</code>	Newline-separated low-level parameters for the recognition engine
<code>root</code>	Name of the grammar rule to enable at start
<code>shared</code>	Boolean controlling whether a single instance of this language model should be shared across all entities to conserve memory

When `asr_model` points to a statistical language model, the `grammar` and `root` variables are ignored.

Shared language models have one important limitation: only one entity may use it to process speech at a time. If another entity starts listening at the same time with the same shared Language

Model, the speech it hears will not be transcribed with speech recognition.

### 3 Construct Settings

#### `api/construct/settings/`

These settings affect Construct behavior overall and are not Scenario specific. The settings resource can only be modified (it cannot be created or deleted). Issue a HTTP PUT request to the settings URI to update settings.

Variable	Definition
<code>log_audio</code>	Boolean controlling whether speech recognition audio is logged
<code>log_events</code>	Boolean controlling whether speech events are logged
<code>dm</code>	Optional network settings for a Discovery Machine behavior engine
<code>vbs2</code>	Optional settings for interfacing with a VBS instance
<code>tts_substitutions</code>	Optional text transformations for use during TTS

Frequently these settings are easily managed through the web interface, but they are also available in the API for dynamic modification from external systems.

We recommend enabling both `log_audio` and `log_events` in order to capture the complete history of events on this Construct instance for later review and tuning.

### 4 Natural Language Parsers

#### `api/construct/parsers/`

This collection resource lists all available Natural Language Parsers installed on the Voisus Server. Parser plugins provided by ASTi will be shown here. This list may be empty if no plugins are installed.

This resource is read-only and is meant to support the selection of a parser in a Construct Entity. The `parser` Entity attribute should be set to the `name` value of a parser in this collection.

### 5 Artificial Intelligence Markup Language (AIML)

#### `api/construct/aiml/`

AIML describes voice interactions and natural language processing capabilities for a Construct Entity. AIML resources created here are referenced by ID with an Entity's `aiml` attribute. The `text` attribute of the AIML resource should contain the complete XML specification of the AIML behavior.

### 6 Runtime Entities

#### `api/construct/entities/`



Variable	Definition
<b>name</b>	Name of the AIML definition
<b>text</b>	XML formatted AIML specification

Construct Entities that are actively running exist here in the runtime Entity API. Runtime Entities are separate from, but related to, Entity resources defined in the Scenario. The key distinction is that a Scenario and the Entity resources within can be edited without the Scenario currently running. When a Scenario runs, its Entities are created and become accessible through this API endpoint. Alternatively, if an application calls for Entities that exist only temporarily, and should not be saved in the Scenario, this endpoint is used to manage those Entities directly.

Runtime entities are synchronized with their corresponding Scenario resource, if one exists.

The attributes shown below track the current runtime Entity state and are updated dynamically as the system runs. For example, the **speaking** value updates when the Entity speaking state changes.

This runtime Entity state is also available via the Construct status streaming API endpoint to eliminate the need to repeatedly poll for this information. See below for documentation on the streaming endpoint.

Note that references to other resources are in the form of IDs or names here instead of by URI like in the Scenario resources.

## 7 Utterances

### **api/construct/utterances/**

An utterance is a single speech event for a Construct Entity. Entity speech can be triggered indirectly from Interactions and Behaviors, or directly from a remote system via this Utterances API endpoint.

This endpoint is only useful when a Scenario is actively running.

An HTTP client may POST to **/api/construct/utterances/to** create a new utterance. The client must provide a JSON object with an **entity** parameter specifying the ID of the entity that should speak. Either **text** or **sound** should be provided, depending on whether the client wishes to use TTS or a prerecorded speech file, respectively. The server response includes the **self** URI which can be polled for status, including the **error** and **complete** flags.

## 8 Streaming Status and Events

### **api/construct/status/**

The Construct status endpoint is useful when a client wishes to receive continuous state updates and event notifications from Construct Entities. An HTTP GET request on this API endpoint will block until the running Scenario is stopped. While the Scenario is running, JSON messages, separated by newline characters, will be sent from the server. Due to the blocking behavior of this endpoint, clients may wish to spawn a separate thread to issue this request in order to not block the client application's main thread.

When a Scenario is running and a client makes a request, a connected message will be sent:

Variable	Definition
<code>id</code>	Entity ID
<code>domain</code>	Name of the Entity's selected DIS exercise
<code>net</code>	Net ID
<code>net_name</code>	Net name from the Commplan
<code>language_model</code>	Language Model ID
<code>language_model_name</code>	Language Model name
<code>behavior</code>	Behavior ID
<code>behavior_name</code>	Behavior name
<code>behavior_state</code>	Enumeration: ['inactive', 'active', 'error', 'finished']
<code>behavior_node</code>	Last executed behavior node
<code>blackboard</code>	URI of Entity blackboard containing current state variables
<code>aiml</code>	AIML definition ID
<code>aiml_name</code>	AIML name
<code>parser</code>	Selected natural language parser name
<code>state</code>	Enumeration: ['initializing', 'active']
<code>listen</code>	true if the Entity is actively listening for speech
<code>speaking</code>	true if the Entity is currently speaking
<code>speech_state</code>	Enumeration: ['idle', 'waiting', 'ptt', 'speaking', 'finished']
<code>speech_text</code>	Text of the current utterance, if there is one
<code>asr_state</code>	Enumeration: ['inactive', 'initializing', 'active', 'error']
<code>radio_enabled</code>	Radio enabled/disabled state
<code>earshot_enabled</code>	Earshot enabled/disabled state
<code>pitch_shift</code>	Voice pitch shift value
<code>tts_voice</code>	Selected TTS voice
<code>tts_rate</code>	Current TTS speech rate (1-9)
<code>tts_volume</code>	Current TTS speech volume (1-9)
<code>marking</code>	Marking field string used to link with an Entity on the network
<code>attributes</code>	Arbitrary key:value pairs for use in behaviors and speech
<code>position</code>	Current geocentric position
<code>default_position</code>	Initial geocentric position
<code>breakpoints</code>	List of currently set behavior breakpoints
<code>breakpoint</code>	Active breakpoint information
<code>radio_state</code>	<code>tx_active</code> and <code>rx_active</code> state for the radio
<code>earshot_state</code>	<code>tx_active</code> and <code>rx_active</code> state for 3D communications

Variable	Definition
<b>id</b>	Utterance ID
<b>entity</b>	URI or ID of the Entity that should speak
<b>text</b>	Speech text for TTS
<b>sound</b>	Sound ID for prerecorded speech playback
<b>background</b>	Sound ID for a background sound during the transmission
<b>complete</b>	Status returned from the server; true when utterance speech has finished
<b>error</b>	Error status returned from the server

All messages from the server will include a **type** which is either “connected”, “update”, “delete”, or “event” depending on the situation. Clients should ignore received messages that have an unknown type.

### 8.1 Entity Update

An update message notifies the client that an Entity was created or updated.

### 8.2 Entity Delete

A delete message notifies the client that an Entity was deleted.

### 8.3 Speech Event

The speech event type notifies the client that an Entity has started speaking.

### 8.4 Speech-Finished Event

The speech-finished event type notifies the client that an Entity has stopped speaking.

### 8.5 Speech Recognition Event

The message event type notifies the client of a speech recognition event.

Variable	Definition
<code>text</code>	Recognized speech text after optional substitutions
<code>original_text</code>	Raw recognized speech from the recognition engine
<code>confidence</code>	Speech recognition confidence level (0–100)
<code>recording</code>	URI for recording of recognized speech
<code>meaning</code>	Extracted speech meaning from parser or AIML
<code>net_id</code>	ID of radio net that received the speech
<code>net_name</code>	Name of the radio net
<code>max_sample</code>	Peak sample value in the received speech
<code>voice</code>	true if this was a voice message (as opposed to a text message)
<code>time</code>	Unix timestamp when speech was recognized