

IN3035UG: DPLMX ISAAC API

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Intona products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Intona hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Intona shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Intona had been advised of the possibility of the same. Intona assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Intona products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance.

© Copyright Intona Technology GmbH, Germany.

Intona and other designated brands included herein are trademarks of Intona in Germany and other countries. All other trademarks are the property of their respective owners.

Website: <https://intona.eu>

Contents

1	UDP Network API	4
2	Introduction	4
2.1	Getting started	4
2.1.1	Finding the device on the network	4
2.1.2	Sending and receiving commands (via CLI on Linux/macOS/msys)	4
2.1.3	Setting up audio streaming (via CLI)	5
3	Persistence and reboots	6
4	AES67 support	6
5	Firmware updates	7
6	UDP API protocol definition and reference	8
6.1	Basics	8
6.1.1	Network layer	8
6.1.2	JSON	8
6.1.3	Commands	8
6.1.4	Reponses	9
6.1.5	Extensibility	10
6.1.6	Packet drops and reordering	10
6.1.7	Retrieving updates and events from a device	10
7	Reference	11
7.1	device_info	11
7.1.1	Parameters	11
7.1.2	Response	11
7.1.3	Product IDs	11
7.1.4	DSP settings	12
7.1.5	Bank/Preset mapping	12
7.1.6	Audio output levels	14
7.1.7	EQ fields	14
7.1.8	EQ type	15
7.1.9	Network settings/state	15
7.1.10	UI fields	15
7.1.11	AES67 settings	16
7.1.12	AES67 session list	16
7.1.13	RTP state	16
7.1.14	Logging state	17
7.1.15	Description	18
7.1.16	Example	18
7.2	set_params	21
7.2.1	Parameters	21
7.2.2	Response	21
7.2.3	Description	21
7.2.4	Example	21
7.3	show_rtp_status	21
7.3.1	Parameters	22
7.3.2	Response	22
7.3.3	Description	22
7.3.4	Example	22
7.4	rtc_get	23
7.4.1	Parameters	23
7.4.2	Response	23

7.4.3	Description	23
7.4.4	Example	23
7.5	rtc_set	24
7.5.1	Parameters	24
7.5.2	Response	24
7.5.3	Description	24
7.5.4	Example	24
7.6	rtc_ntp	24
7.6.1	Parameters	24
7.6.2	Response	25
7.6.3	Description	25
7.6.4	Example	25
7.7	sap_purge	25
7.7.1	Parameters	25
7.7.2	Response	25
7.7.3	Description	25
7.8	reset_settings	26
7.8.1	Parameters	26
7.8.2	Response	26
7.8.3	Description	26
7.9	reboot	26
7.9.1	Parameters	26
7.9.2	Response	26
7.9.3	Description	26

1 UDP Network API

2 Introduction

The device can be configured over network. It exchanges commands and responses as JSON objects in UDP packets. The main intention is that this is used as API by control programs, although you can also send and receive raw commands manually.

Each UDP packet contains one or multiple commands as JSON objects, formatted as single line of text. These UDP packets are sent to port 7054. The device will respond with an UDP packet containing JSON object, sent back to the sender's IP, MAC and UDP port.

Here is an example:

User => Device

```
{"command": "set_params", "api_version": 6, "seq": 123, "dsp": {"mute": true}}
```

Device reply

```
{"seq": 123}
```

Here the "set_params" command is used to mute audio output. The reply indicates success, because it does not contain an "error" field. The "seq" field is for free use by the sender (but must be a JSON number), and will be sent back with the reply. It can be used to associate responses with commands.

Erroneous requests (including JSON syntax errors etc.) will result in a response with the "error" field set.

2.1 Getting started

This section shows how to initialize the device and how to start audio streaming.

2.1.1 Finding the device on the network

Unless configured otherwise, each device uses an automatically assigned link-local IP address (RFC 3927)¹. Addresses start at **169.254.1.0**, but are generally unpredictable and unstable. For example, if a device is restarted, it may pick a different address. Usually, it will attempt to reuse the previous address, but you cannot rely on this. (You could use the set_params command to assign a static address to a device to avoid this.)

You should perform device discovery by sending **device_info** commands to the broadcast address (**169.254.255.255**). These commands can be sent periodically to provide live updates on settings, and whether the device is still up. To avoid overloading the device, an update period of no shorter than 500 ms is recommended. This is especially important if several clients are on the network.

Use the **device_id** field (queried with **device_info**) as unique, stable device identifier.

2.1.2 Sending and receiving commands (via CLI on Linux/macOS/msys)

Since the protocol simply sends and receives UDP packets containing JSON text, you can send and receive commands with common open source tools:

```
r1wrap -S '>' nc -u 169.254.1.0 7054 | jq
```

¹<https://tools.ietf.org/html/rfc3927>

nc sends and receives UDP packets (as text in this case), **r1wrap** provides line editing and a history, and **jq** pretty prints JSON responses.

2.1.3 Setting up audio streaming (via CLI)

The device supports SAP (Session Announcement Protocol, RFC 2974)² AES67 session discovery. Passive discovery is always active. The device does by default not play any audio from network. The DPLMX network API can be used to play a specific stream.

All important settings and information can be queried with:

User => Device

```
{"command": "device_info"}
```

Which results in this reply (only AES67 relevant parts shown):

Device reply

```
{
  ...
  "stream": {
    "name": "BKLYN-II-0d0c9a : 31",
    "link_offset": 48,
    "nominal_level_dbu": 0,
    "ch0": 0
  },
  "streams": {
    "list": [
      {
        "n": "BKLYN-II-0d0c9a : 31",
        "i": "1 channels: Analog R",
        "c": 1
      },
      ...
    ]
  },
  "rtp": {
    "lock": 3
  },
  ...
}
```

This example run returned two AES67 sessions. The device is streaming the first one (according to the stream.name field).

A specific session can be selected with:

User => device

```
{"command": "set_params", "stream": {"name": "BKLYN-II-0d0c9a : 32"}}
```

²<https://tools.ietf.org/html/rfc2974>

The `set_params` command has a large number of parameters, which can be used to control every detail about the device. The `stream.name` parameter controls which session should be selected. Internally, the name is matched against the internal list of sessions, and the first matching session is selected. By default, the name is empty, and no session matches.

The streaming status is indicated by the `rtp.lock` field.

Using this command persists all set parameters to the device flash memory (see section below).

3 Persistence and reboots

Most device settings are persistent. They are written to flash memory. After boot, the settings are restored from flash. There is a delay of about 1 second before the device starts writing the changed settings to the flash. This delay counts since the most recent change. For this reason, the device should not be power cycled immediately after a change. (The device should survive power cycles at any possible moment, but recently changed settings will not be restored.)

It can take a while until a rebooted device recovers and streams audio. The following processes incur additional delays:

- Waiting for SAP announcements. (Unless they are cached on flash; if they are not cached, not replicated by other devices, and the sender is a Dante device, it can take up to 30 seconds.)
- Waiting for PTP clock stabilization. (Requires receiving at least 2 PTP announcements, streaming may start before; can take about 2 seconds.)

The processes listed above happen in parallel. For network API access, the device needs to negotiate a link layer IP address, which takes about about 7 seconds on average. If static IP configuration is used, the API is available as soon as the network link is up.

Firmware updates require a reboot. Firmware updates can reset all or some settings. Never apply firmware updates in a production setting.

4 AES67 support

The device follows the AES67 standard. The following design choices were made:

- Only 48000 Hz is supported.
- SAP is required, which is optional on AES67. (Other session management protocols may be supported in the future.)
- Only Dante devices were available as sources for testing and development, and there are some workarounds that may not agree with the standards.
- The stream name (SDP "s" field) is used to uniquely identify sessions. Duplicate names are considered an invalid misconfiguration.
- ISAAC devices send out cached SAPs at a relatively high rate. The SAP was specified for old network technology and prioritized reducing traffic. ISAAC devices are likely used in settings where every second can count. This "SAP replication" can help if devices need to be setup very quickly.
- Likewise, SAP timeouts were reduced to avoid listing stale sessions.
- There is a workaround that is supposed to help against badly implemented IGMP-snooping in consumer switches. (See `igmp_hack`.)

5 Firmware updates

The device runs a TFTP server for firmware updates. It accepts firmware binaries under the name "firmware.bin". All other files are rejected. Update can be done with any TFTP client, for example curl:

```
curl -T path/to/firmware.bin tftp://169.254.1.0
```

6 UDP API protocol definition and reference

6.1 Basics

6.1.1 Network layer

The protocol uses JSON encoded in UTF-8 text, encapsulated in UDP. Each UDP packet payload contains exactly 1 JSON object (a list of fields enclosed with "{" and "}"). A JSON object cannot span multiple UDP packets. A JSON object represents a single command. The convention is to add a trailing newline character ("\n", ASCII 0x0A) at the end of the JSON text.

UDP packets are exchanged between devices and API clients. Clients send commands as UDP packets to port 7054. The device will attempt to send a response packet back to the client immediately, using the command packet's source IP/port as destination IP/port for the response packet.

The device does not support receiving fragmented IP packets. All commands are assumed to fit within 1472 bytes of UDP payload. On the other hand, responses by the device may be sent as fragmented UDP packets if they exceed the MTU.

A MTU of 1500 bytes is assumed.

Be aware that UDP is an unreliable protocol. Packet loss and reordering is possible. The device sends exactly one response for each request, and this response could get dropped by the network. If you do not receive any response, you cannot know whether the request or response was dropped (or whether the device went offline). You may have to repeat the request until you get a response (assuming the command is idempotent - most are). Since the network has to reliably transport AES67 streams, it is assumed that everything operates with some headroom below the maximum possible network bandwidth, and packet dropping does not actually happen in practice.

6.1.2 JSON

JSON, as used by the DPLMX network protocol, is specified by RFC 8259³ and RFC 7493⁴. The protocol requires that all top-level JSON values are JSON objects, and that each UDP packet contains a single JSON object. In particular, it is assumed that JSON numbers use double floats (binary64 in IEEE 754-2008⁵), which limits their integer precision to about 53 bits (see RFC 7493 section 2.2⁶ for details).

The command reference will use the term "integer" for numbers which are representable as signed 32 bit integers (sometimes unsigned 32 bit integers).

6.1.3 Commands

The UDP payload consists of text that can be parsed as a single JSON object. There must be no 0 bytes or other non-text data. Whitespace can be present, but should be avoided to reduce the packet size. A final newline should be added at the end of the JSON text. (It's currently not required, but this is a threat that this may change.)

On the IP level, the command must be a single IP fragment.

Each JSON object sent to the device represents a single command. It can contain the following fields:

³<https://tools.ietf.org/html/rfc8259>

⁴<https://tools.ietf.org/html/rfc7493>

⁵https://en.wikipedia.org/wiki/IEEE_754-2008

⁶<https://tools.ietf.org/html/rfc7493#section-2.2>

Name	Type	Meaning
command	String	Required. Name of the command to run.
seq	Number	Sent back with responses, and is otherwise not interpreted or changed by the device.
api_version	Integer	Optional, but passing this is highly recommended. The current protocol version is 6.
add_crc	Boolean	If true, a CRC will be appended to the JSON data of the response. (See Responses section below.)
(other)	...	Command specific parameters are located in the same top-level JSON object.

All fields except "command" are optional. The "api_version" should also be passed.

The most important commands are "device_info" and "set_params". All other commands are obscure and rarely needed.

6.1.4 Responses

Responses indicate success or failure of a previously sent command. There is exactly one response per command. (Minus dropped network packets.)

The payload contains a single line of JSON text, with redundant whitespace trimmed (clients shouldn't rely on this) and a single newline character at the end (clients can rely on this).

On the IP level, large responses are sent as multiple IP fragments. These are reassembled to a large UDP packet, whose payload follows the described format. Large responses usually happen with "device_info" commands, and can be mostly avoided by using the "select" parameter.

The responses JSON objects contain the following fields:

Name	Type	Meaning
seq	Number	Always present. The same value as the command packet, or 0 if it was absent.
error	String	If present, a protocol error of some kind has happened. This is usually used for malformed commands only. It's a string value, that contains an error message. It is not intended to be machine-readable.
warning	String	If present, warnings or errors. This is a string value and contains newline characters (one message per line). This is not necessarily an error. It is not intended to be machine-readable. The main use of this field and the "error" field is to help with debugging.
(other)	...	Command specific response fields are located in the same top-level JSON object.

If "add_crc" was set to true in the command, the response packet uses a slightly modified format. The four byte sequence "\n/#" (0a 2f 2f 23 in hex) is appended after the JSON payload (instead of a final "\n" byte), followed by 8 ASCII digits representing the CRC32 of the payload before the four byte header. For example, the byte sequence " 0a 2f 2f 23 61 39 34 64 61 31 65 61" at the end of the packet indicates the CRC32 value 0x a94d a1ea. If present, this byte sequence will always be 12 bytes long and end with the packet's payload. The CRC32 uses the same polynomial as Ethernet and zlib. This can be useful if fragmented UDP packets occur, and the UDP checksum is not trusted. (In single fragment cases, the packet is protected sufficiently by the Ethernet CRC.) The format was chosen such that we can claim it's still a text based protocol. Naive clients can just not set this field. Then the "\n" byte is always the last byte of the payload, and the only one in the payload.

6.1.5 Extensibility

JSON was picked as base of the ISAAC protocol to make it easier to extend the protocol in the future. For this reason, both device and clients must make the following assumptions:

- The ordering of object members is arbitrary. Both the device and the client are required to deal with any ordering of members in an object. It is possible that members change their order on firmware updates, or randomly.
- Unknown object members should be ignored. The device will add a warning to the response object if an object member is unknown, but will otherwise ignore it.
- If new parameters are added to a command, it must be done in a backward compatible way.
- Any code handling commands or responses must be forward compatible, and ignore unknown parameters.

6.1.6 Packet drops and reordering

UDP can drop or reorder packets. Clients can use the "seq" field to avoid confusion due to reordering. It can also be used to detect packet drops (for example, if no response for a command is received, the client can issue a redundant command to check whether the device is still operating). If a response is lost, the client does not know whether the command was executed, or the result status. In this case the client needs to check the device state, or send the command again to be sure.

6.1.7 Retrieving updates and events from a device

There is no mechanism for receiving any kind of events with the API. Poll the device by sending **device_info** commands. A polling period of 500 ms or higher is recommended.

7 Reference

All commands can include standard parameter fields described in the section above. The same applies to responses. Only command-specific request/response fields are listed in the tables below.

7.1 device_info

Retrieve information about device firmware, DSP settings, AES67 settings/status/session list, and more.

7.1.1 Parameters

Member	Type	Meaning
select	JSON array	Optional. Each array item must be a string. Each string can be the key of a sub-object to include in the response. If missing, all sub-objects are returned.

7.1.2 Response

Member	Type	Meaning
product	String	General product ID. Always "ISAAC" (read-only).
firmware_version	String	Firmware version identifier, usually a release version string (read-only). Note: never parse this in program code! Some firmware binaries use a git hash instead of a version number.
api_version	Integer	API revision used by device, may be different from client's version in "api_version" field (read-only). Note: older firmware versions do not have this field.
fw_date	Integer	UNIX-time of firmware build date (read-only). Note: older firmware versions do not have this field.
device_id	String	128 bit globally unique hardware identifier in hexadecimal, for example "00313553504e52433033303436303535". This is the MCU's builtin UUID. It never changes, unless the MCU is physically replaced. (Read-only.)
firmware_update_error	String	Optional. If present, a human readable error description on firmware update errors. Always cleared on reboots. (Read-only.)
product_id	Integer	Product ID (see table below). (Read-only, normally.)
dsp	Object	DSP settings (see table below).
net	Object	Network settings/state (see table below).
ui	Object	Informational fields for Seeburg Network Manager (see table below).
stream	Object	AES67 parameters (see table below).
streams	Object	AES67 stream list (see table below).
rtp	Object	AES67 state (see table below).
logging	Object	Logging control (see table below).

7.1.3 Product IDs

The "product_id" field is encoded as follows:

Raw value	Name
0	(Never used, might happen on major device flash corruption.)
1	Seeburg G Sub 1201 dp++

Raw value	Name
2	Seeburg G Sub 1501 dp++
3	Seeburg Trisonus 10 dp
4	Seeburg X2 dp
5	Seeburg X4 dp
6	Seeburg X6 dp
7	Seeburg X8 dp
8+	Future products

7.1.4 DSP settings

The response "dsp" field is an object with the following fields:

Member	Type	Meaning
mute	Boolean	Audio output muted setting.
preset	Integer	Preset ID (starting from 0). Meaning and range depend on product_id and bank fields (see table below).
npresets	Integer	Preset count (read-only).
bank	Integer	Bank ID (starting from 1). Meaning and range depends on product_id (see table below).
nbanks	Integer	Bank count (read-only).
level	Integer	Audio output level (see table below).
delay	Float	Audio output delay in meters (range 0-1.8, out of range undefined).
eq_en	Boolean	EQ processing enabled.
eq0	Object	Each of these 5 fields contains information for each EQ band.
eq1		See table below for EQ fields.
eq2		
eq3		
eq4		

7.1.5 Bank/Preset mapping

The dsp bank and preset fields depend on the product_id as follows:

product_id	bank	preset list
1	1	"LP 120Hz/HP 120Hz" "LP 160Hz/HP 160Hz" "LP 160Hz/HP 120Hz" "LP 120Hz/HP 120Hz + Aux" "LP 160Hz/HP 160Hz + Aux" "LP 160Hz/HP 120Hz + Aux"
1	2	"LP 120Hz/L-Series HP 180Hz" "LP 160Hz/L-Series HP 180Hz" "LP 120Hz/L-Serie +Low" "LP 120Hz/L-Series HP 180Hz + Aux" "LP 160Hz/L-Series HP 180Hz + Aux" "LP 120Hz/L-Serie +Low + Aux"

product_id	bank	preset list
1	3	"LP 120Hz + i5" "LP 160Hz + i4" "LP 120Hz + GL-Serie xov" "LP 120Hz + i5 + Aux" "LP 160Hz + i4 + Aux" "LP 120Hz + GL-Serie xov + Aux"
2	1	"LP 100Hz/HP 100Hz" "LP 140Hz/HP 140Hz" "LP 140Hz/HP 100Hz" "LP 100Hz/HP 100Hz + Aux" "LP 140Hz/HP 140Hz + Aux" "LP 140Hz/HP 100Hz + Aux"
2	2	"LP 120Hz/GL-Series xov Flat" "LP 120Hz/GL-Series xov -Low" "LP 120Hz/GL-Series xov +Low" "LP 120Hz/GL-Series xov Flat + Aux" "LP 120Hz/GL-Series xov -Low + Aux" "LP 120Hz/GL-Series xov +Low + Aux"
2	3	"LP 100Hz/K24 xov Flat" "LP 120Hz/K20 60° Flat" "LP 120Hz/K20 90° Flat" "LP 100Hz/K24 xov Flat + Aux" "LP 120Hz/K20 60° Flat + Aux" "LP 120Hz/K20 90° Flat + Aux"
3	1	"Flat" "HP 100Hz" "Low-Boost" "Flat (soft)" "HP 100Hz (soft)" "Low-Boost (soft)"
3	2	"Flat on-wall" "HP 100Hz on-wall" "Low-Boost on-wall" "Flat on-wall (soft)" "HP 100Hz on-wall (soft)" "Low-Boost on-wall (soft)"
3	3	"Flat in-wall" "HP 100Hz in-wall" "Low-Boost in-wall" "Flat in-wall (soft)" "HP 100Hz in-wall (soft)" "Low-Boost in-wall (soft)"
4	1	"Flat" "HP 120Hz" "Low-Boost" "Flat (soft)" "HP 120Hz (soft)" "Low-Boost (soft)"

product_id	bank	preset list
5	1	"Flat"
		"HP 100Hz"
		"Low-Boost"
		"Flat (soft)"
		"HP 100Hz (soft)"
6	1	"Flat"
		"HP 100Hz"
		"Low-Boost"
		"Flat (soft)"
		"HP 100Hz (soft)"
7	1	"Flat"
		"HP 100Hz"
		"Low-Boost"
		"Flat (soft)"
		"HP 100Hz (soft)"
		"Low-Boost (soft)"

In the table above, the third column specifies the list of presets (starting from preset ID 0) for a specific product_id and bank field ID. For example, with product_id=3 and bank=2, preset=0 selects "Flat on-wall", preset=4 selects "HP 100Hz on-wall (soft)", etc. For product_id=3, the range of the bank field is 1-3, and the range of the preset field is 0-5 (each bank has always the same number of presets for a given product). Setting out of range values lead to undefined behavior.

7.1.6 Audio output levels

The "level" field is encoded as follows:

Raw value	Gain
0	-12 dB
1	-9 dB
2	-6 dB
3	-3 dB
4	0 dB
5	+3 dB
6	+6 dB

Out of range values when setting the field lead to undefined behavior.

7.1.7 EQ fields

Each of the eqN fields is an object with the following fields:

Member	Type	Meaning
en	Boolean	Enable this band. (If the dsp.eq_en field is false, the EQ is disabled regardless of the value of the en field.)
type	Integer	EQ type ID (see table below.)
freq	Float	EQ band frequency (range 10-24000, out of range undefined).
q	Float	EQ Q value (range 0.1-100, out of range undefined).
gain	Float	EQ gain value in dB (range: -25-25, out of range undefined).

7.1.8 EQ type

The EQ "type" field is encoded as follows:

Raw value	Name	Meaning
0	None	Same function as "en" field set to false.
1	LSHELF	
2	PEQ	
3	PEQ2	
4	HSHELF	

Out of range values when setting the field will lead to an error response.

7.1.9 Network settings/state

The response "net" field is an object with the following fields:

Member	Type	Meaning
mac	String	Current MAC address, formatted as 6 groups of 2 hexadecimal numbers separated by ":", for example "AB:CD:12:34:56:7E". (Read-only.)
ip	String	Current IP address formatted as quad-dotted IP address, for example "169.254.1.0". (Read-only.)
static_ip	String	Configured static IP address (the same format as the "ip" field), or "0.0.0.0" if the IP address should be dynamically allocated. Applied after device reboot.
igmp_hack	Boolean	If true, enable an IGMP workaround, that is supposed to help with broken IGMP-snooping switches. (Enabled by default.)

7.1.10 UI fields

The response "ui" field is an object with the following fields:

Member	Type	Meaning
order	Integer	Order ID (used for GUI device list sort order).
name	String	User-assigned device name (max. 127 bytes).
loc	String	User-assigned device location text (max. 127 bytes).
memo	String	Text field for free use (max. 127 bytes).

The strings are supposed to be encoded as UTF-8 (usually enforced by JSON), though the device will accept any encoding in JSON or these fields.

None of the values in this object are actually interpreted by the firmware. They are for use by GUI tools.

7.1.11 AES67 settings

The response "stream" field is an object with the following fields:

Member	Type	Meaning
name	String	Stream name, used to select session.
link_offset	Integer	Link offset in samples.
nominal_level_dbu	Float	Nominal level in dBu.
ch0	Integer	Source channel to play to stream, counting from 0.

All current devices support mono only, and ch0 is the only channel. Should more channels be supported in the future, fields like ch1, ch2, ... will be added.

7.1.12 AES67 session list

The response "streams" field is an object with the following fields:

Member	Type	Meaning
list	JSON array	Actual session array, see below.

Each item in the "list" array is an object with the following fields:

Member	Type	Meaning
n	String	Session name
i	String	Session info
c	Integer	Channel count

This is the list of discovered sessions. The session name is the "s" field in the SDP, and is compared to stream.name to select a session. The session info is the "i" field in the SDP. The channel count is also as indicated in the SDP, and gives the upper range of the stream.ch0 field (exclusive).

7.1.13 RTP state

The response "rtp" field is an object with the following fields:

Member	Type	Meaning
lock	Integer	AES67 streaming state

The lock field can have the following values:

Raw value	Meaning
0	No PTP, no RTP data
1	PTP locked, but no RTP data
2	No PTP, but RTP locked
3	PTP and RTP locked

“Locked” in this context means “synchronized to”. Streaming is possible at state 2 or 3. State 2 is indicated by a blinking yellow Ethernet port LED, at state 3 it glows normally. RTP is for Real-Time Transfer Protocol, the part of AES67 that transports actual audio data over network.

7.1.14 Logging state

The response “logging” field is an object with the following fields:

Member	Type	Meaning
en	Boolean	Whether logging is enabled

Logging is a debug feature. The log is useful to firmware developers only. Technically inclined people may be able to interpret parts of the log, but it’s not recommended. Logging makes the device broadcast UDP packets to IP 255.255.255.255 on port 7055. The log state is off by default, and is not restored across firmware reboots.

Log packets have a format that is slightly different from normal protocol messages. The first part of the UDP payload consists of JSON text, like in normal messages. This is followed by a 0 byte, a binary length field, and the log text.

Position	Length	Field
0	N	JSON text of length N bytes
N + 0	1	Literal 0 byte (N found by searching for the first 0 byte)
N + 1	2	Little endian uint16_t log text length
N + 2	length	ASCII log text, using length field above

This binary length field and log payload is found by searching for a 0 byte. JSON cannot legally contain a 0 byte.

This format was chosen because we wanted to continue using JSON for its positive properties, all while reducing device MCU load by not having to turn the log message into a legal JSON string value (escaping).

The JSON text is an object with the following fields:

Member	Type	Meaning
command	String	Contains the string “raw_log”, otherwise it’s not a log message
pos	Integer	Log byte position, can be used to detect dropped log messages
overflow	Boolean	Internal overflow flag (not always present)

The “pos” field is the number of log bytes output so far. The log receiver can add the packet’s log text length to the “pos” field to know the “pos” field of the next log packet. If the values mismatch, a log packet

must have been dropped or reordered. The value of this field is an unsigned 32 bit integer and rolls over. "overflow" indicates that the device log buffer (which is of limited size) overflowed before the contents could be sent to the network. The "pos" field is also discontinuous in this case.

7.1.15 Description

This command reports general information about device, network status, and DSP settings. This command should be used for device discovery.

The "set_params" command is an important companion command. It mirrors most fields in the "device_info" command, and can be used to set the corresponding fields. All references to setting fields in the parameter description above is in context of the "set_params" command.

The "select" field can be used to retrieve only some information. It is an array of strings, where each string is the name of a JSON sub-object. Only sub-objects mentioned in this list are sent in the response. For example, if you send **select: ["dsp"]**, then other sub-objects like "streams" are not sent, only the "dsp" sub-object. If the field is missing, everything is sent.

7.1.16 Example

User => Device

```
{"command": "device_info", "seq": 123, "api_version": 6}
```

Device reply (prettified)

```

{
  "seq": 123,
  "product": "ISAAC",
  "firmware_version": "v1.5.2",
  "api_version": 6,
  "fw_date": 1654077358,
  "device_id": "00313753384d37373038303432303139",
  "product_id": 3,
  "net": {
    "mac": "1A:7D:9C:4F:26:3A",
    "ip": "169.254.225.237",
    "static_ip": "0.0.0.0",
    "igmp_hack": true
  },
  "stream": {
    "name": "AVIOUSB-5346cb : 2",
    "link_offset": 48,
    "nominal_level_dbu": 0,
    "ch0": 0
  },
  "streams": {
    "list": [
      {
        "n": "AVIOUSB-5346cb : 2",
        "i": "1 channels: Left",
        "c": 1
      }
    ]
  },
  "rtp": {
    "lock": 0
  },
  "ui": {
    "order": 0,
    "name": "",
    "loc": "",
    "memo": ""
  },
  "logging": {
    "en": false
  },
  ...

```

Device reply (continued)

```
...
"dsp": {
  "mute": false,
  "preset": 1,
  "npresets": 2,
  "nbanks": 1,
  "bank": 1,
  "level": 4,
  "peak": 0,
  "gr": 0,
  "eq_en": true,
  "delay": 0,
  "eq0": {
    "en": true,
    "type": 2,
    "freq": 100,
    "q": 1,
    "gain": 0
  },
  "eq1": {
    "en": true,
    "type": 2,
    "freq": 200,
    "q": 1,
    "gain": 0
  },
  "eq2": {
    "en": true,
    "type": 2,
    "freq": 500,
    "q": 1,
    "gain": 0
  },
  "eq3": {
    "en": true,
    "type": 2,
    "freq": 2000,
    "q": 1,
    "gain": 0
  },
  "eq4": {
    "en": true,
    "type": 2,
    "freq": 4000,
    "q": 1,
    "gain": 0
  }
}
}
```

7.2 set_params

Write various device settings.

7.2.1 Parameters

This command mirrors a number of fields in the device_info response. All fields are optional, and only parameters sent with the command are actually set.

See the device_info command for type and meaning of each field, and whether it can be written.

7.2.2 Response

Success/error only. If an error happens, and the request contains multiple parameters, it's possible that the request is partially applied. In this case, it's unpredictable which parameters were applied and which not. Use the device_info command to confirm the new values.

7.2.3 Description

This command can set most fields which appear in device_info. Name, location, JSON data type, and meaning are exactly the same as the fields in device_info.

If the client needs to change multiple parameters, it's recommended to set all parameters at once, using a single "set_params" command.

7.2.4 Example

This mutes the DSP, and leaves all other parameters untouched:

User => Device

```
{"command": "set_params", "dsp": { "mute": true }}
```

Device reply

```
{"seq": 0}
```

This sets multiple fields (enable mute, enable EQ 3, set memo to "hello"), and leaves all other parameters untouched:

User => Device

```
{"command": "set_params", "dsp": { "mute": true, "eq3": {"en": true} },
  "ui": { "memo": "hello" } }
```

Device reply

```
{"seq": 0}
```

7.3 show_rtp_status

Show information about current AES67 streaming (network to device).

7.3.1 Parameters

None.

7.3.2 Response

Member	Type	Meaning
ip	String	IP address of source stream.
port	Integer	UDP port of source stream.
clock_offset	Integer	SDP mediack parameter (0 if unset)
link_offset	Integer	Offset chosen with set_session command.
samplesize	Integer	Audio sample size in bits per channel
samplerate	Integer	Audio sample rate
channels	Integer	Number of channels sent by source stream.
output_channels	Array	Channels used, as chosen with set_session command.
packet_drops	Integer	New packet drops since last command.
packet_drop_last_ms	Integer	Absolute device time of last packet drop in ms. Can wrap around.
rtp_received_last_ms	Integer	Absolute device time of last RTP audio packet received. Can wrap around.
clock_locked	Integer	Complete lock. (device_info rtp.lock has more detailed information.)

7.3.3 Description

This shows the status of AES67 streaming, as well as some of the selected parameters.

This command is for debugging, and you should probably not rely on it being present or compatible in the next firmware version.

7.3.4 Example

User => Device

```
{"command": "show_rtp_status"}
```

Device reply

```

{
  "seq": 0,
  "ip": "239.69.214.91",
  "port": 5004,
  "clock_offset": -1479386051,
  "link_offset": 64,
  "samplesize": 24,
  "samplerate": 48000,
  "channels": 1,
  "output_channels": [
    0
  ],
  "packet_drops": 2147483647,
  "packet_drop_last_ms": 693922,
  "rtp_received_last_ms": 823747,
  "clock_locked": true,
  "ptp_sync_last_ms": 823543
}

```

7.4 rtc_get

Read RTC time.

7.4.1 Parameters

None.

7.4.2 Response

Member	Type	Meaning
utc	Integer	Current time in UTC

7.4.3 Description

Report the current time in UTC⁷ from the internal RTC (Real Time Clock). This is maintained independently from the PTP time, and is less precise, but continues progressing while the device is turned off.

7.4.4 Example

User => Device

```
{"command": "rtc_get"}
```

Device reply

⁷https://en.wikipedia.org/wiki/Coordinated_Universal_Time

```
{
  "seq": 0,
  "utc": 1544017789
}
```

7.5 rtc_set

Set RTC time.

7.5.1 Parameters

Member	Type	Meaning
utc	Integer	New current time in UTC

7.5.2 Response

Success/error only.

7.5.3 Description

Counter part of utc_get.

7.5.4 Example

User => Device

```
{
  "command": "rtc_set",
  "utc": 1544017789
}
```

Device reply

```
{
  "seq": 0,
}
```

7.6 rtc_ntp

Initiate NTP time synchronization.

7.6.1 Parameters

Member	Type	Meaning
ntp_server_ip	String	NTP server, formatted as IP address.

7.6.2 Response

Success/error only (this is sent before the NTP server is contacted - it indicates correctness of the command formatting only).

7.6.3 Description

Synchronize with an explicitly specified NTP time server once. This is not like normal NTP operation. The firmware won't attempt to periodically synchronize the time. Instead, the operation ends once a single NTP server packet has been received. Whether the time was correctly set should be checked with `rtc_get`, and there is no event or any other indication whether the time server could even be reached.

7.6.4 Example

User => Device

```
{
  "rtc_ntp": "rtc_ntp",
  "ntp_server_ip": "192.168.5.132"
}
```

Device reply

```
{
  "seq": 0,
}
```

7.7 sap_purge

Remove sessions from the device's session cache.

7.7.1 Parameters

Member	Type	Meaning
age	Float	Maximum session age in seconds that should be preserved (older sessions are deleted). (Default: 60)
blocktime	Float	Time in seconds during which new SAPs should be ignored. (Default: 0)

7.7.2 Response

Success/error only.

7.7.3 Description

The device caches previously discovered sessions (using the SAP protocol) and hold them in RAM and on the flash. There is currently a session timeout of 100 seconds, after which a session is automatically deleted. This command can be used to remove a session immediately. Sometimes, it can help resolving playback problems by removing broken cache entries.

The "blocktime" parameter can be set if `sap_purge` is supposed to be run on all devices in the network. The problem is that the devices will never execute the command at exactly the same time, and bogus SAP

packets still could be propagating through the network. In particular, devices will send their own bogus cached/replicated SAPs to all other devices. To avoid that devices immediately readd bogus sessions, this parameter can disable SAP processing temporarily.

7.8 reset_settings

Wipe all settings.

7.8.1 Parameters

None.

7.8.2 Response

None.

7.8.3 Description

This resets all user settings, and reboots. Like with the "reboot" command, no response is sent. The device_id is not changed.

7.9 reboot

Reboots the device.

7.9.1 Parameters

None.

7.9.2 Response

None.

7.9.3 Description

The device is reset, as if power-cycled. There is no response, because the hardware is reset before a response can be sent out.

Document version: 19 / Jun 23, 2022 11:08