

Monitor by 2N5

Documentation

The user guide, the product, the telephony reference, the guides and the legal terms in effect, in one document.

Version	October 25, 2026
Published by	2N5
Website	monitor.2n5global.com

Contents

User guide

[Getting started](#)
[The dashboard](#)
[A voice circuit's page](#)
[Data circuits and the agent](#)
[Websites and email](#)

[Incidents and alerts](#)
[Reports](#)
[People and organizations](#)
[Settings](#)
[FAQ and help](#)

The product

[How managed VoIP monitoring works](#)
[SIP trunk and VoIP monitoring with real test calls](#)
[Toll-free \(0800\) and IVR testing and monitoring](#)
[VoIP voice quality measured with MOS-LQO](#)
[Line-down alerts and escalation](#)
[Carrier SLA report, monthly and yearly](#)

[Internet link monitoring for every branch](#)
[Website monitoring: downtime and content](#)
[Email monitoring: are messages arriving?](#)
[Security and LGPD in VoIP monitoring](#)
[Plans and pricing for 0800, IVR and SIP trunk monitoring](#)

Reference

[SIP codes](#)
[Q.850 causes](#)

[Glossary](#)

Guides

[0800 not connecting? How to monitor 0800 and IVR](#)
[Failing calls: diagnosis by SIP code and Q.850](#)
[How to enforce a carrier SLA with proof](#)
[Line-down alerts without false alarms](#)
[Passive monitoring vs active test calls](#)
[SAC decree and RGC: how to prove your 0800 was up](#)

[Silent calls and one-way audio in VoIP](#)
[SIP trunk, 0800 and IVR monitoring: a guide](#)
[Telephony SLA: yearly and monthly calculation](#)
[Voice quality: MOS and MOS-LQO in practice](#)
[What is synthetic call monitoring](#)
[Why an 0800 number won't connect from a mobile phone](#)

Legal

[Terms of Use](#)

[Privacy Policy](#)

User guide

USER GUIDE · CHAPTER 1

Getting started

To start using Monitor by 2N5, open the invitation 2N5 emailed you, create your password and sign in. This chapter covers your first sign-in, two-step verification, how to recover your password and where to change the language and the theme.


By [2N5 team](#) Updated October 14, 2026

Signing in for the first time

You do not create an account yourself: 2N5 sets up your organization and emails you an invitation.


1. Open the invitation email and click its link. The **Manage** page shows your organization's name and your name, already filled in.
2. Type your password in **Password** and again in **Confirm password**.
3. Tick the box that accepts the Terms of Use and confirms you have read the Privacy Policy.
4. Click **Start managing**. Monitor opens your organization's dashboard.

The first time you sign in, a short introduction, **Welcome to Monitor**, explains in six steps how Monitor works. Use **Next** to move on or **Skip** to close it. On the last step you can start the guided walkthrough, which stops once on each page.



After that, sign in at Monitor's address:


1. Type your **Email** and **Password**.
2. Click **Sign in**.



Protecting your account with two-step verification

Two-step verification asks, besides your password, for a six-digit code from an authenticator app on your phone. We recommend turning it on. Your organization can make it required for every manager; Monitor then asks you to set it up right after your password.

1. In the bottom-left corner, click your name and choose **Account settings**.
2. In the **Two-step verification** box, click **Set up**, then **Set up** again in the window that opens.
3. Scan the QR code with your authenticator app. If you prefer, copy the **Setup key** with **Copy key** and paste it into the app.
4. Type the code the app shows in **Code from your authenticator app** and click **Turn on**.
5. Save your recovery codes with **Copy** or **Download** and click **Done**. Each code stands in for the app once, if you lose your phone.



From then on, after your password, Monitor asks for the code from the app. Type it and click **Verify**. Without your phone, choose **Use a recovery code**.

If you forget your password


1. On the sign-in screen, type your email in **Email**.
2. Click **Forgot your password?**. If that address has an account, a reset link arrives by email.
3. Open the link. On the **Set a new password** page, type the new password in **New password** and **Confirm new password**.
4. Click **Set password**, then **Go to sign in**.

Changing your password signs out every session that was open on your account, on any computer. To change a password you still know, use the **Password** box in **Account settings**.

Language

Monitor speaks Brazilian Portuguese and English. The language you choose is saved on your account.


1. In the sidebar, click the current language, just above your name.
2. Choose **English** or **Português (Brasil)**.



The sign-in screen has the language menu too, in the top-right corner.

Light or dark theme

1. Click your name in the bottom-left corner.
2. Choose **Light**, **Dark** or **System**. **System** follows your computer's theme.



In the same menu, **Sign out** ends your session on this computer. To end it everywhere, use **Sign out of all devices** in **Account settings**.

Frequently asked questions

I never received the invitation. What do I do?

Check your spam folder and any promotions tab for a message from 2N5. If it is not there, ask 2N5 to send the invitation again: the link is valid for seven days, and a new invitation replaces the old one. A person who already manages your organization can also invite you from the Managers screen, with the Invite manager button.

I lost the phone with my authenticator app. How do I sign in?

On the code screen, choose Use a recovery code and type one of the codes you saved when you turned two-step verification on. Each code works once. If you have none left, ask 2N5 or another manager of your organization to reset your two-step verification, then set it up again on the new phone the next time you sign in.

Does the language I choose also apply to the alerts?

No. The language in the sidebar applies to the screens you see, and it is saved on your account, so another computer opens in the same language. Each alert by email goes out in the language of the contact receiving it, set on that contact on the Escalation screen, and a report's language is chosen when you send it.

USER GUIDE · CHAPTER 2

The dashboard

The dashboard is the first screen after you sign in, and it answers one question: is everything working right now? The **Circuits** tab shows your voice and data circuits, and the **Online services** tab shows your websites and mailboxes.

By [2N5 team](#) Updated October 21, 2026


What the dashboard shows

1. In the sidebar, click **Dashboard**.
2. At the top, choose the **Circuits** or the **Online services** tab. The number on each tab is how many items it shows; a red dot means something is down.

On the **Circuits** tab, the **Circuit overview** page opens with five figures:

- **Monitored circuits:** how many voice circuits your organization watches.
- **Needs attention:** circuits that are down, degraded or suspect, and how many incidents are open.
- **Lowest 30-day availability:** the circuit with the lowest availability over the last month, and its value.
- **Monitoring blind:** circuits Monitor could not test. That is a problem on 2N5's side, not an outage of yours.
- **Hourly call budget:** how much of this hour's allowance of test calls has been used.

Below them, the **Incidents over time** chart shows the last 30 days of incidents, with one color for each kind of failure.



States and colors

The **Voice and data circuits** table lists every circuit with its current state:

State	What it means
Up	The latest test passed.
Suspect	A test failed and is being retried before anyone is told.
Degraded	The circuit answers, but its quality fell below the limit in tests in a row.
Down	The outage was confirmed after four failed tests in a row, and an incident is open.
Monitoring blind	Monitor could not test it; the time does not count as an outage.

The **Last 24 hours** column shows one bar per hour. Green is **Up throughout**, orange is **Degraded**, red-orange is **Down briefly**, red is **Down for much of it**, purple is **Monitoring blind** and light gray is **Not observed**. The **Quality** column shows the latest measured MOS-LQO of the circuits answered by a recording, and **30-day availability** shows the share of the last month the circuit was up.

Click a circuit's name to open its page, covered in the chapter [A voice circuit's page](#).

Filtering and sorting

1. In the table, click **Filter**.
2. On the left, choose what to filter by: **State**, **Entity** or **Provider**.
3. On the right, tick the options. The number next to each one says how many circuits it has.

4. Click **Done**. To see everything again, use **Clear all**.

Voice and data circuits

Filter

Clear all

MONITORED

KIND

STATE

0800 main number

Aurora Clinic

Operadora Exemplo

Voice circuit

Up

Booking IVR

Paulista branch

Operadora Exemplo

Voice circuit

Degrad

Campinas front desk

Campinas branch

Operadora Exemplo

Voice circuit

Down

Filter

Clear all

State

Entity

Provider

Up

3

Suspect

0

Degraded

1

Down

1

Monitoring blind

0

Clear all

Done

scored

Paulista front desk

Paulista branch

Operadora Exemplo

Voice circuit

Up

—

not scored

99.80%

30 min ago

Paulista fiber

Paulista branch

Fibra Exemplo Telecom

Data circuit

Up

14 ms · 1.8 ms jitter · 0.0% loss

latest round


—

6 min ago

To sort, click a column heading that has the triangle, such as **Monitored**, **State** or **Last tested**. A second click reverses the order.

The Online services tab

The **Online services** tab does the same for the websites and mailboxes 2N5 monitors for you. The figures at the top are **Monitored services**, **Needs attention**, the websites' **Average response** over the last 24 hours, **Email reaching the inbox** and **Certificates expiring** in the next 14 days.



Further down are the **Websites** and **Email** tables, with the same filter and sorting. The chapter [Websites and email](#) explains each one.

Frequently asked questions

Why is a circuit shown as Suspect when nobody was alerted?

Suspect means one test failed and Monitor is retrying before it tells anyone. A single failure never becomes an alert. An incident opens only after four failed tests in a row, the first one and three retries. If a retry passes, the circuit goes back to Up and nobody receives a message about it.

What does Monitoring blind mean?

It means Monitor itself could not test the circuit, for example because 2N5's own connection failed at that moment. That time is not counted as an outage of your circuit, and it stays out of the availability in your reports. 2N5 is alerted and takes care of it; there is nothing you need to do.

Does the dashboard update by itself?

Yes. The dashboard receives changes live, without reloading the page, and shows in its top-right corner how long ago it was updated. At the top of the screen, Monitoring live means the tests are running normally. If the connection drops, the dashboard says so and keeps the last good data on screen while it reconnects.

USER GUIDE · CHAPTER 3


A voice circuit's page

Every voice circuit has a page with everything Monitor knows about it: its state now, the last 100 hours, the measured voice quality, its incidents and every test call. It is also where you tell 2N5 about planned maintenance.

By 2N5 team Updated October 14, 2026

Opening the circuit's page

1. On the **Dashboard**, in the **Voice and data circuits** table, click the circuit's name.
2. To go back, use **Circuits** in the path above the name.



At the top of the page you find:

- The circuit's **state**, such as **Degraded**, and how much of the call was proven, such as **Audio one way** or **Audio both ways**.
- The **measured MOS-LQO** of the latest scored call, in red when it is below the limit.

- The circuit's organization, entity and provider, the number Monitor **Dials** and **What answers it: A greeting or recorded message, A test tone or echo service** or **Other: it may ring or be answered**.
- **Last tested:** when the latest test call was made.
- **Planned maintenance:** the next window on record, if any.

The last 100 hours

The **Last 100 hours** box has one bar per hour and shows what the run of tests concluded in each hour, in the same colors as the dashboard. Hover over a bar to see the hour and what happened; a bar with an incident leads to it.

Voice quality: MOS and MOS-LQO

MOS (Mean Opinion Score) is the 1-to-5 scale telephony uses to say how a call sounds to the person listening: close to 4.5 is excellent, and below 3 a conversation gets hard. MOS-LQO is the same scale calculated from the audio, with no listeners: Monitor compares the audio each test call received with the expected recording, using a reference-based perceptual listening-quality model that returns MOS-LQO (ITU-T P.800.1).

The **Voice quality** chart shows the score of each answered call over the last hours, with two lines:

- **degraded below 3.0:** calls below this line count as degraded. When that repeats in tests in a row, the circuit becomes **Degraded** and an incident opens.
- **fully healthy from 3.8:** above this line a call counts as fully healthy.

These limits belong to your organization and are under **Settings**, on the **Quality thresholds** tab. Circuits not answered by a recording show **not scored**. The glossary explains [MOS](#) and [MOS-LQO](#) in more detail.

State history and incidents

- **State history** lists every change of state, newest first, with its reason and time. An outage shows first as **Suspect**, when the first test fails, then as **Down** or **Degraded**, once it is confirmed after four failed tests in a row.
- **Incident record** shows confirmed failures only. Click an incident to open its page, covered in [Incidents and alerts](#).

The observation log: every test call

Further down, the **Observation log** keeps every test call, newest first:

- **Tested:** when the call was made, and when it was due.
- **Verdict:** what the call concluded.
- **SIP:** the provider's response code, such as 200 for answered.
- **Answer:** how long the call took to be answered.
- **RTP sent / received:** the audio packets sent and received.
- **MOS-LQO** and **Heard:** the call's score and whether the expected recording was heard.

Observation log ⓘ							
Filter Download CSV							
TESTED ⌵	VERDICT ⌵	SIP ⌵	ANSWER ⌵	RTP SENT / RECEIVED	MOS-LQO ⌵	HEARD	TONES SENT / BACK
14 Oct 2026, 20:30:06 due 14 Oct 2026, 20:30:00	Up from 2N5	200	1405 ms	493 / 493	4.13	Expected recording heard	—
14 Oct 2026, 20:00:04 due 14 Oct 2026, 20:00:00	Up from 2N5	200	3238 ms	485 / 479	4.18	Expected recording heard	—
14 Oct 2026, 19:30:06 due 14 Oct 2026, 19:30:00	Up from 2N5	200	1972 ms	508 / 504	4.37	Expected recording heard	—
14 Oct 2026, 19:00:08 due 14 Oct 2026, 19:00:00	Up from 2N5	200	1405 ms	513 / 508	4.10	Expected recording heard	—
1–20 of 2106 observations < 1 of 106 >							

Use **Filter** to see only one kind of result and **Download CSV** to take the log into a spreadsheet. The arrows at the foot turn the pages.

Adding planned maintenance

If your provider or your own team is going to work on the circuit, add a maintenance window. Failures inside the window open no incident, and the time stays out of availability.

1. On the circuit's page, in the **Planned maintenance** box, click **+ (Add a window)**.
2. Under **What the work is on**, check the circuit. You can choose other circuits at the same entity.
3. In **Start date** and **Start time**, enter when the work starts. Times are in your organization's time zone, shown beside them.
4. Choose how long it **Lasts** (1 h, 2 h, 4 h, 8 h or 24 h), or fill in **End date** and **End time**.
5. Write the **Reason**, for example replacing a piece of equipment.
6. Click **Save window**.



Add planned maintenance

Time inside a window is kept out of the denominator, and reported separately.



Kept at

Aurora Clinic · Organization

What the work is on ⓘ

Paulista front desk

WINDOW

Start date

14/10/2026



Start time

09:00 PM

Sao Paulo, UTC-3



Lasts

1 h

2 h

4 h

8 h

24 h

End date

14/10/2026



End time

11:00 PM



Reason

Carrier card swap

Cancel

Save window

The window appears in the **Planned maintenance** box and can be deleted there, as long as it has not happened yet.

Frequently asked questions

Why does my circuit show not scored under quality?

The MOS-LQO score is measured only when the circuit is answered by a greeting or recorded message, because it compares the audio the call received with the expected recording. A circuit that rings or is answered by a person is tested by the call itself, without a score. If you want the score, ask 2N5 to set up the circuit's reference recording.

Does a failure during planned maintenance count against availability?

No. Time inside a maintenance window is kept out of the availability figure and reported separately in your reports, and a failure in that time opens no incident and alerts nobody. Add the window before the work starts: a window does not change what was recorded before it existed.

Can I change how often the circuit is tested?

Each circuit keeps the values it was added with: the interval between tests, the interval between retries and how many retries to make before an outage is confirmed. To change them on an existing circuit, ask 2N5. The values on the Settings screen apply to circuits added afterwards.

USER GUIDE · CHAPTER 4


Data circuits and the agent

A data circuit is a site's internet link, measured by an agent: a small service that runs in the background on a computer at that site. The circuit's page shows whether the link is up, the quality of the connection and how much of the contracted speed actually arrives.

By [2N5 team](#) Updated October 14, 2026

What you see of a data circuit

1. On the **Dashboard**, **Circuits** tab, data circuits appear in the same table as voice circuits, with the **Kind Data circuit**. The **Quality** column shows the latest round's latency, jitter and loss.
2. Click the circuit's name to open its page.



On the circuit's page:

- **Reporting from** shows the computer the agent runs on, its system and when it was last heard from. **Reporting** in green means readings are arriving.
- **State** is **Up**, **Degraded** or **Down**. A link outage is confirmed by further measurements, as with voice circuits, and only then does it open an incident and alert the escalation ladder.
- **Connection latency** is the round trip to the internet, in milliseconds.
- **Download speed** and **Upload speed** show the latest speed test and how much of the **Contract** it reached.

- **Connection details** shows the public address, the site's router (and whether it answered) and the provider.
- **Last 100 hours** has one bar per hour, in the dashboard's colors.
- **Connection quality** holds the **Latency**, **Packet loss** and **Jitter** charts, and **Connection speed** the **Download** and **Upload** tests.
- The **Observation log**, at the foot of the page, keeps every measurement round, with **Filter** and **Download CSV**.

Because the agent also measures the site's router, Monitor can tell a router fault apart from a provider fault.

Installing the agent with 2N5's help

2N5 usually installs the agent with you: it agrees a time, confirms which computer will run the agent and records your written authorization. Choose a computer at the site that stays on and uses the link to be measured.

Until an agent reports, the circuit's page shows **Waiting for its agent**, with two choices: **Install on another computer** and **Set up this computer**. To install it together with 2N5:

1. Open the circuit's page on the computer at the site and click **Set up this computer**. The **Install the agent for** page shows the circuit's name.
2. Check **Authorized by** and **Their email**, which carry your name, and write in **Computer** which machine it is, for example the reception PC.
3. Tick **The customer has authorized this installation in writing**.
4. Choose the computer's system and click **Get the command for**. Run the command on that computer, as 2N5 guides you.
5. Watch **2. What we have seen**: its steps move on by themselves as Monitor sees the agent register and send its first readings.

[Dashboard](#) > [Paulista fiber](#) > Install the agent

Install the agent for Paulista fiber

A web page cannot install software on the machine you are reading it on. What it can do is give you the right download for this machine, with everything it needs to find us already in it, and then tell you what we have actually seen.

1. Download it on the machine at the site

Run this on the machine whose internet connection you want measured — not on this one, unless they are the same machine.

Authorized by

Marina Alves

Their email

marina.alves@example.com

Computer

Reception PC, São Paulo office

☐ The customer has authorized this installation in writing.

Ubuntu

macOS

Windows

↓ Get the command for macOS

Once the agent reports, the circuit's page shows its measurements. The first hours of the **Last 100 hours** box stay light gray, as **Not observed**, because they come before the install.

What only 2N5 changes

How often the link is measured and its speed tested, the limits that make it count as degraded or down, how long a silence counts as a quiet agent, and the check from outside are set by 2N5, because 2N5 answers the link's alerts. To change any of them, ask 2N5.

Frequently asked questions

Does the agent leave a port open on the computer?

No. The agent makes outbound connections only: it measures the site's link and sends the readings to 2N5, without opening any incoming port on the computer or the network. It sends link measurements only, from a fixed list. If your security team wants the details, ask 2N5 for the list of what the agent measures.

What happens if the agent's computer is switched off?

The readings stop arriving and the agent shows as no longer reporting. That alone does not mean the link is down, because a quiet agent may just be a computer that was switched off. 2N5 is alerted when an agent goes quiet and agrees with you what to do. That is why the agent belongs on a computer that stays on.

What if the site has no computer for the agent?

Ask 2N5 for a check from outside. 2N5 then tests from outside whether the link answers, up or down, without measuring speed, latency, loss or DNS. It needs a fixed public address; a link behind shared addressing (CGNAT) cannot be checked this way.

USER GUIDE · CHAPTER 5


Websites and email

Besides your circuits, 2N5 can monitor your websites and your mailboxes. For a website, Monitor checks that it answers, that it is fast, that its certificate is valid and whether its content changed; for a mailbox, whether a test message reaches the inbox in time.

By [2N5 team](#) Updated October 21, 2026

A website's page

1. On the **Dashboard**, open the **Online services** tab.
2. In the **Websites** table, click the website's name.



At the top are the website's state, its address and three figures: **Answer** (the code the server returned, such as **200 OK**), **Tested every** and **Slow above**. The website is:

- **Up** when it answers normally.

- **Degraded** when it answers but more slowly than the slow limit, or when its security certificate expires within 14 days.
- **Down** when it does not answer, answers with an error, or its certificate has expired or is not trusted. The outage is confirmed after retries, and only then does an incident open.

Below that come **Last 100 hours**, **Response time** over the last 24 hours, **State history** and **Incident record**. Further down, **Answers** sums up the last 30 days of codes, **Certificate** shows **Valid until**, **Days left** and **Issued by**, and the **Observation log** keeps every test.

Content changes

When the **Also check this website's content** switch is on, Monitor compares the page with a reference: the text and structure the page had when the switch was turned on. There are no content settings: the check uses fixed limits and the website's own test schedule. A page built entirely by JavaScript cannot be checked and reads as unreadable. That is how it notices, for example, a menu that disappeared, a page that was swapped or text somebody changed without telling anyone.

1. In the **Websites** table, the website with the changed page shows as **Content changed**, a state worse than **Up** and better than **Down**. Click its name.
2. On the website's page, below **Answers** and **Certificate**, two sections show the content: the state and what changed, with the **This change is expected** button and the full comparison, and the content checks, one per round.

[illegible]

When the page changes more than allowed, Monitor checks again before it concludes. If the change persists, it opens a **Content changed** incident and alerts the escalation ladder. If the website is down at the same time, there is a single incident, the outage, and the content change stays as a line in the observation log. A slow website plus a content change are two incidents.

This change is expected

If the change was made on purpose, for example a holiday opening-hours notice, close the incident like this:

1. Open the incident from the website page's **Incident record**, or from **Incidents** in the sidebar.
2. Read **Was this change made on purpose?:** how many lines and tags changed, whether a new script appeared and **Where on the page** the change is.
3. To see exactly what changed, click **See the full comparison**. The **Removed** and **Added** text appear side by side.
4. If the change is expected, click **This change is expected**. Today's page becomes the new reference, the incident closes and your name is recorded.

Incidents > INC-0005

CONTENT CHANGED

INC-0005

Patient portal

34.6% of the text changed, above the 5% limit, and the structure is 93.3% the same.

Open the website
History

Open for
1 h 01 min
since 14 Oct, 19:55

Severity
Critical
inside working hours

Acknowledged
Not yet
waiting since 19:57

Detected
19:55

Confirmed
19:57

Team told
19:57

Acknowledged
not yet

Recovered
not yet

Was this change made on purpose?
34.6% of the text changed, above the 5% limit, and the structure is 93.3% the same.

✓ This change is expected

2
lines changed

3
tags changed

0
new scripts

Where on the page
Main content

Script changes
2 lines changed
No script changed


See the full comparison

If nobody made the change, do not click: treat it as a problem with the website and talk to whoever looks after it. What the page must show, and the parts that change by themselves, such as a clock or the day's news, are agreed with 2N5.

A mailbox's page

An email check sends a test message to your mailbox's address on a schedule, and checks that it arrived, where it landed and how long it took.

1. On the **Dashboard**, **Online services** tab, in the **Email** table, click the check's name.
2. At the top are **Tested every** and **Waits up to**, the longest the message may take before it counts as late.



The check is **Up** when the message reaches the inbox in time, **Degraded** when it lands in Spam or arrives late, and **Down** when it does not arrive. The **Delivery time** chart shows the latest tests. The log shows, for each test, when it was sent, how long it took to arrive, where it landed, whether it was deleted and the sender checks (SPF, DKIM and DMARC). In the dashboard's **Email** table, the **Inbox** bar shows how full the mailbox is.

For the mailbox to be checked, 2N5 sends you a link to connect it. If the link expires before you use it, ask 2N5 for a new one.

Frequently asked questions

Does a content change count against the website's availability?

No. A content change shows the website as Content changed and opens an incident of its own, and never counts as downtime. If the website goes down at the same time, there is a single incident, the outage. Your reports show availability and content changes separately.

Who can click This change is expected?

A manager of the organization, or of the entity the website sits at. Clicking it makes today's page the new reference, closes the incident and records your name. If the change was not made on purpose, do not click it: treat the website as an incident and talk to whoever looks after it.

Does the test message stay in my mailbox?

Usually not. Once Monitor has checked where the message landed, it deletes it, and the log shows Deleted for every test. If the mailbox connection may only read, the message stays and the log says so. For the mailbox to be checked, 2N5 sends you a link to connect it.

USER GUIDE · CHAPTER 6

Incidents and alerts

An incident is a confirmed failure: it opens only after four failed tests in a row, the first one and three retries. Monitor then alerts the people on your escalation ladder, layer by layer, until somebody acknowledges the incident from the link in the message.

By [2N5 team](#) Updated October 14, 2026

How a failure becomes an incident

1. A test call fails. The circuit becomes **Suspect** and nobody is told yet.
2. Monitor repeats the test, usually every 2 minutes.
3. If the first test and its three retries fail, that is four failed tests in a row, the failure is confirmed: the circuit becomes **Down** or **Degraded** and an incident opens.
4. The first layer of the escalation ladder is alerted.
5. Once the circuit passes two clean tests in a row, the incident closes as recovered and everybody who was alerted gets the message that it is back.

The incident list

1. In the sidebar, click **Incidents**.
2. Use the **All**, **Open**, **Not acknowledged** and **Resolved** tabs, or **Filter**.
3. Click **View** to open an incident.

Incidents

3
open now

All 5

Open 3

Not acknowledged 2

Resolved 2

Filter

TICKET	STATUS	MONITORED	PROVIDER	ACK	STARTED	DUR...	
INC-0005	Open	Patient portal Website content changed		Not acknowledged	14 Oct 19:55	ongoing 1 h 01 min	View
INC-0004	Open	Campinas front desk Refused	Operadora Exemplo	Not acknowledged	14 Oct 20:00	ongoing 56 min	View
INC-0003	Open	Booking IVR Poor audio	Operadora Exemplo	Marina Alves, 15 min	14 Oct 17:30	ongoing 3 h 26 min	View
INC-0002	Resolved	Campinas front desk Refused	Operadora Exemplo	Marina Alves, 10 min	02 Oct 14:00	56 min	View
INC-0001	Resolved	Paulista front desk Silent call	Operadora Exemplo	Rafael Souza, 12 min	23 Sep 10:00	1 h 26 min	View

1-5 of 5 incidents

Every incident has a number, such as **INC-0004**. The **Ack** column shows who acknowledged it and how quickly, or **Not acknowledged**.

An incident's page

Incidents > INC-0004

△ ACTIVE OUTAGE INC-0004

Campinas front desk

Operadora Exemplo · The call reached the provider, but nothing at the destination took it: the phone or phone system there may be switched off, disconnected or not registered (SIP 480).

📞 Open circuit
🕒 History

Down for
56 min
since 14 Oct, 20:00

Severity
Critical
inside working hours

Acknowledged
Not yet
waiting since 20:06

Test calls since
13
all failing the same way

Detected
20:00

Confirmed
20:06

Team told
20:06

Acknowledged
not yet

Recovered
0 of 2 clean

Provider-ready evidence ⓘ
📄 Copy

Incident: INC-0004
Circuit: Campinas front desk
Provider: Operadora Exemplo
First symptom: 2026-10-14 20:00:00 -03:00
Confirmed: 2026-10-14 20:06:00 -03:00 after 3 failed calls
State: down · Severity: critical · still open

Observed: Confirmed after 4 consecutive failed tests. The call reached the provider, but nothing at the destination took it: the phone or phone system there may be switched off, disconnected or not registered (SIP 480).
13 test calls since the first failure, all failed.

Supporting observations ⓘ

Up

SIP 200
The call was answered and audio flowed in both directions.
14 Oct 2026, 19:56:03

Suspect

SIP 480
The call reached the provider, but nothing at the destination took it: the phone or phone system there may be switched off, disconnected or not registered (SIP 480).

- At the top, the circuit, its provider and the symptom, with **Open circuit** and **History**, which shows every message sent about the incident and to whom.
- **Down for**, **Severity**, **Acknowledged** and **Test calls since**. The severity is **Critical** when the incident opened inside the organization's working hours, and **Minor** when it opened outside them.
- **Where it stands** shows the timeline: **Detected**, **Confirmed**, **Team told**, **Acknowledged** and **Recovered**.
- **Provider-ready evidence** is a block of text with the circuit, the provider, the first symptom, the confirmation and the test calls. Click **Copy** and paste it into the provider's ticket, without rewriting it.
- **Supporting observations** lists the test calls from the incident's own time.

The escalation ladder: who gets told

1. In the sidebar, click **Escalation**. The **Who gets told** page shows **The ladder**.
2. Each layer has a name and its people. A confirmed failure goes to layer 1 first. If nobody acknowledges, the alert is repeated and then climbs to the next layer.
3. To see or change how long each layer waits and how many times it repeats the alert before climbing, click **Timings**.
4. To add somebody, use **Add someone** on the layer, or **Add contact** at the top, and enter their name, role, language and addresses: email and mobile.

Who gets told ⓘ

[+ Add contact](#)

The ladder ⓘ [Timings](#)

1 IT team 1 person

RS Rafael Souza
 Email · avg 6 min
[+ Add someone](#)

2 Coordination 1 person

MA Marina Alves
 Email
 no acknowledgments yet
[+ Add someone](#)

3 Directors 1 person

BL Beatriz Lima
 Email
 no acknowledgments yet
[+ Add someone](#)

Beside the ladder ⓘ

Client contacts · 1 **Administrators · 0** **Providers · 2**


CP Carlos Pereira
 Medical director
 Email
 [>](#)

Beside the ladder holds the **Client contacts**: people who get the message that a circuit is down and being handled, and the message that it is back, without having to acknowledge anything. If no layer acknowledges an incident, 2N5 is alerted to act.

Acknowledging from an email link

Every alert from the ladder carries a link to acknowledge the incident. Acknowledging means: I have it, you can stop alerting.

1. Tap the link in the message. The page opens on your phone or computer, without signing in to Monitor.
2. Tap **Acknowledge incident**.
3. The page confirms it: the incident is acknowledged and no further alerts will be sent for it.



If somebody else already acknowledged it, the page says who and that nothing more is needed. If the circuit had already recovered, the page says so and records your response.

What each message means

Alerts arrive by email. The examples below use the circuits of this guide's fictitious clinic.

The outage is confirmed

Alert for Aurora Clinic: circuit Campinas front desk is DOWN since 10/14 19:36. Acknowledge at .../a/... to stop further alerts.

It goes to the ladder's first layer. Tap the link and acknowledge; then check the circuit or open a ticket with the provider.

Nobody has acknowledged yet

Alert for Aurora Clinic: Campinas front desk still DOWN since 10/14 19:36, not yet acknowledged.

Acknowledge at .../a/... to stop alerts.

It goes to the next layer, when nobody acknowledged in time. Acknowledge from the link if you are taking it.

An update for the client

Update for Aurora Clinic: Campinas front desk is down since 10/14 19:36. We are working on it and will tell you when it is back.

It goes to the client contacts. There is nothing to do; it is for information.

The circuit is back

Update for Aurora Clinic: Campinas front desk is back UP at 10/14 20:42 after 1 h 6 min down. No action needed.

It goes to everybody who was told about the outage. The incident is closed.

The emails

By email, the subject says what it is about in a plain sentence, the same one the email opens with: *Main line is down* for an outage, *Main line has poor call quality* when calls connect but sound bad, and *Main line is working normally again* when it is back. Each email ends by saying why you receive it and who to ask to stop it. The email gives the problem in plain words, when it started, when it was confirmed and, for people on the ladder, the **Acknowledge** link. Data circuits, websites and mailboxes follow the same pattern, with the circuit's or the service's name in the message.

Frequently asked questions

Why did I get an alert when the circuit was already back?

The alert goes out the moment the failure is confirmed. If the circuit comes back soon after, you also get the message saying it is working again. Opening the link of an alert that is already resolved shows that the circuit had recovered and records your response; no new alert is sent about that incident.

Do I need to acknowledge an incident somebody else already acknowledged?

No. One person acknowledging is enough to stop the alerts for everyone. If you open the link afterwards, the page shows who acknowledged it first and says nothing more is needed from you. Who acknowledged and how long it took stay on the incident and go into your reports.

Through which channels do I get alerts?

Monitor alerts by email. Each contact can have more than one address, and the order they are used in is set on the contact.

USER GUIDE · CHAPTER 7

Reports

A service report is the proof of how your circuits behaved over a period: availability against target, every incident and how long it took to be answered. You generate as many as you like, email them as PDFs and can schedule the monthly one to go out on its own.

By [2N5 team](#) Updated October 14, 2026

Reports and the SLA are for Enterprise customers.

Generating a report

1. In the sidebar, click **Reports**. The **Service reports** page lists the reports already generated under **History**, newest first.
2. Click **Generate report**.
3. Under **Coverage**, choose in **Covers** the whole organization or one entity, and in **Period** one of these:
 - **This month, automatically**: the usual report, for the current month.
 - **A month I choose**: to close an earlier month, such as September.
 - **A custom period**: enter the **First day** and **Last day**, at most one year in total. That is how you make the yearly SLA report.
4. Click **Next**. Under **Contents**, write a **Title** if you like, and choose the **Pages to include**, the **Severities to include** and the **Circuits**. Leaving everything ticked gives the full report.
5. Click **Next**, check the **Review** and click **Generate**.



Generate a report

Leave everything as it is for the usual report: this month, every circuit, all pages.



1 Coverage

2 Contents

3 Review

Covers ⓘ

Aurora Clinic · Organization

Period

This month, automatically

Cancel

Next >

The new report appears at the top of **History**. Click **View** to open it.

Service reports

Delivery

Off


Generate report

History ⓘ

PERIOD ⌵	TITLE ⌵	GENERATED ⌵	BY ⌵	
2026-09	Service report — 1-30 September 2026 Aurora Clinic	14 Oct 2026, 20:56 America/Sao_Paulo	marina.alves@example.com	View

Reading the report

On screen, the report appears as the pages of the PDF. At the top are **Back to history**, **CSV**, **Download PDF** and **Send**.



- **Summary:** the **Voice availability** for the **Report period** and the **Last 12 months**, with the target and whether it was met; the number of **Incidents**, the **Total downtime** and the **Average response** to the first acknowledgment; **The month, day by day**, with one row per circuit and a mark on each day with a confirmed fault; and **In short**, a paragraph that tells the month in words.
- **What happened:** every incident, with its date, the circuit, what Monitor observed, how long it was down and how quickly it was answered, and the total at the end.
- **Circuit comparison:** every circuit side by side, with its call quality.
- **Provider scorecard:** each provider, what was promised against what was delivered.

The availability target is under **Settings**, on the **Reports & clock** tab, and each provider's contract terms on the **Contracts** tab. **CSV** downloads the figures as a spreadsheet; **Download PDF** downloads the document.

Sending a report

1. Open the report and click **Send**.
2. Under **Recipients**, choose client contacts or type an address. For each person, choose the pages and the language of the PDF.
3. Under **Message**, write a **Note** if you like. It appears after the greeting, above the figures.
4. Under **Review and send**, check **How the email reads** and send it.

Each person receives a PDF holding only the pages chosen for them, with the stored report's figures exactly as they are. The same report can be sent a few times an hour; after that, Monitor asks you to wait.

Scheduling the monthly report

1. On the **Service reports** page, click **Delivery**.
2. Under **Automatic delivery**, turn on **Send a report every month** and choose **Produced on** (the day) and **At** (the time).
3. Choose the email's language, adjust its text if you like, and tick **Include the headline figures** to put availability and incidents in the body of the email.
4. Save. **Next report** shows when the next one goes out and the period it covers: one month back, up to the end of the day before.

Frequently asked questions

Does the report's availability include maintenance windows?

No. Time inside a planned maintenance window is kept out of the figure and shown separately. Time when Monitor itself could not test, Monitoring blind, is kept out too, because it is not an outage of your circuit. Availability is the time up divided by the time measured in the period.

Does a report change if I generate it again later?

Each report is kept exactly as it was produced, with the figures and names of that moment, and always opens the same from History. Generating the same period again creates a new report next to the old one. That lets you send a provider or your directors a document that does not change afterwards.

How do I make the yearly SLA report?

In Generate report, choose A custom period and enter the first and last day of the year, for example from January 1 to December 31. A period can be at most one year. The first page gives the period's availability and that of the last 12 months, compared with your organization's availability target.

USER GUIDE · CHAPTER 8

People and organizations

Monitor has two kinds of people: managers, who sign in and see the screens, and contacts, who receive alerts and reports without signing in. The organization's structure, in entities, decides what each manager sees.

By 2N5 team Updated October 14, 2026

Managers and contacts

- **Managers** sign in to Monitor with an email and password. They see the dashboard, incidents and reports for the part of the organization they manage, and look after the escalation ladder, the entities and the other managers.
- **Contacts** do not sign in. They are the people on the escalation ladder, who receive alerts with the link to acknowledge, and the **Client contacts**, who receive the down and back-up updates and can receive reports. Contacts are added under **Escalation**, as the chapter [Incidents and alerts](#) explains.

Entities: the organization's structure

The organization sits at level 0, and below it come companies, branches, offices or departments, in as many levels as you need. Every circuit, website and mailbox sits at one entity.

1. In the sidebar, under **Community**, click **Entities**.
2. The table shows each entity with its **Level**, **Type**, how many **Circuits** it has, the **Storage** used, when it was **Created** and its **Managers**. Use **Name or type** to search.
3. To create an entity, click **Add entity**, give it a name, choose one of the organization's types or type a new one, and choose under **Placed below** where it sits.
4. To download the structure as a spreadsheet, click **Export CSV**.

Entities

2 entity types in this organization

Export CSV

Add entity

Q Name or type

NAME	TYPE	CIRCUITS	STORAGE	CREATED	MANAGERS
Aurora Clinic Level 0	Organization	4 of 10 voice circuits	14.1 KiB of 1 GiB	25 Aug 2026	Marina Alves
Campinas branch Level 1	Branch	1		25 Aug 2026	No manager
Paulista branch Level 1	Branch	2		25 Aug 2026	No manager

Moving an entity carries everything below it, and every circuit keeps its history. An entity can be deleted only when nothing is below it and no circuit is in it.

Inviting a manager

1. Under **Community**, click **Managers**. The list shows every manager with their **Email**, **Entity**, status and **Last signed in**, and the invitations still waiting for an answer.
2. Click **Invite manager**.
3. Fill in the **Name** and email and choose the entity the person will manage.
4. Click **Send invitation**. The person receives the email, sets their own password from the link and signs in.

Managers ⓘ

Require two-step

Entity structure

Export CSV

+ Invite manager

Filter

MANAGER ⌵	EMAIL ⌵	ENTITY ⌵	STATUS ⌵	LAST ACCESSED ⌵
Marina Alves	marina.alves@example.com	Aurora Clinic	Active	14 Oct 2026, 20:56:08

1–1 of 1 managers

While an invitation waits, the list shows until when its link is valid, with **Send again** and **Withdraw**. To take somebody's access away, use **Remove manager**: they are signed out at once. **Require two-step** makes two-step verification required for the organization's managers.

Who sees what

A manager sees the entity they manage and everything below it, and nothing beside it:

Manager of	Sees
The organization (level 0)	Every entity, circuit, incident and report.
One branch	That branch only, and whatever is below it.

The dashboard, the incident list, the reports and the entities adjust by themselves to what each manager may see. A report generated by a branch's manager covers only that branch.

Providers

Under **Community**, **Providers** lists the providers of your circuits, with each one's technical support contact. That is where the address for an incident's evidence is kept, for when a ticket has to be opened.

Frequently asked questions

Does a branch's manager see the other branches' circuits?

No. A manager sees the entity they manage and everything below it, and nothing beside it. The Campinas branch's manager sees only Campinas's circuits, incidents and reports; the organization's manager sees every branch. An entity can have several managers, and whoever manages an entity can invite managers to it and to the entities below it.

What is the difference between a manager and a contact?

A manager has a sign-in: they open Monitor, see the dashboard, generate reports and change what their account allows. A contact does not sign in; they receive the escalation ladder's alerts, with the link to acknowledge, or the updates and reports meant for the client. The same person can be a manager and a contact at once.

What happens to the history when a manager leaves the company?

Remove the manager on the Managers screen: they are signed out at once and can no longer see the organization. The circuits, incidents, reports and history stay with the organization. The acknowledgments they made stay on the incidents, under their name, as the record of what happened.

USER GUIDE · CHAPTER 9

Settings

Monitor is a managed service: 2N5 adds your circuits, sets up the tests and runs the test calls. You adjust what belongs to your organization, such as who gets told, your reports and the defaults on the Settings screen, and ask 2N5 for the rest.

By [2N5 team](#) Updated October 14, 2026

The Settings screen

1. In the sidebar, click **Settings**.
2. In **Settings for**, choose the whole organization or one entity. An entity can have values of its own, different from the organization's.
3. Choose a tab: **Voice circuits**, **Data circuits**, **Online services**, **Quality thresholds**, **Reports & clock** or **Contracts**.
4. Change the values. The **What this adds up to** box shows the effect straight away, such as how many calls a day and how long from the first failure to a confirmed outage.
5. Click **Save changes**, or **Discard** to drop them.

Settings

Settings for **Aurora Clinic · Organization**

Voice circuits | Data circuits | Online services | Quality thresholds | Reports & clock | Contracts

Routine tests

Voice test interval ⓘ minutes

Accept ringback ⓘ ☒ On ☐ Off

Accept busy ⓘ ☒ On ☐ Off

When a test fails

Voice retry interval ⓘ minutes

Voice retry count ⓘ retries

Before it counts as recovered

WHAT THIS ADDS UP TO
For these 4 circuits ⓘ

Calls a day, a circuit added now	96
Calls an hour, all circuits ⓘ	8
First failure → outage confirmed	6 min
First clean test → recovered	2 min

No unsaved changes
Nothing changed yet.

What each tab holds:

- **Voice circuits:** the interval between tests, whether ringing out unanswered or a busy signal count as the circuit working, the interval and number of retries before an outage is confirmed, and how many clean tests confirm recovery.

- **Data circuits** and **Online services**: the defaults for links, websites and mailboxes.
- **Quality thresholds**: the measured MOS-LQO scores below which a call counts as degraded and from which it counts as fully healthy.
- **Reports & clock**: the organization's time zone, used on every screen, alert and report, and the reports' availability target.
- **Contracts**: the providers' contract terms, which the report compares with what was delivered.

Each circuit keeps the values it was added with: a change here applies to circuits added afterwards.

What you change

Where	What
Account settings	Password, two-step verification, signing out of all devices and deleting your account.
Sidebar and account menu	Language and theme.
Escalation	Who gets told, in what order, at which addresses, and each layer's Timings .
Entities and Managers	The organization's structure and who manages it.
A circuit's page	Planned maintenance windows.
Reports	Generating, sending and scheduling the monthly delivery.
A content incident's page	This change is expected.
Settings	The defaults your account may change.

What only 2N5 changes

Ask 2N5 to:

- Add a kind of item to your package, or raise the limit of circuits, websites or mailboxes.
- Add a new circuit and run its verification call, or change an existing circuit's number, test interval or retries.
- Set up or replace the reference recording of a circuit answered by a recorded message.
- Change how often a data circuit is measured and its speed tested, and the limits that make a link count as degraded or down.
- Change how often a website or mailbox is checked and when it counts as down, and what a website's page must show.
- Install the agent, or set up a check from outside, for a data circuit.
- Send a new link to connect a mailbox.

When a screen says a value is set by 2N5, or that a problem is on 2N5's side, there is nothing for you to do: 2N5 has already been alerted. The chapter [FAQ and help](#) says how to reach 2N5.

Frequently asked questions

I changed a value in Settings. Why did an older circuit not change?

Each circuit keeps the values it was added with. The values on the Settings screen are the organization's or the entity's defaults, and they apply to circuits added after the change. To change the test interval, the retries or the recovery of a circuit that already exists, ask 2N5.

How do I raise the number of circuits I can monitor?

Your organization's package sets how many voice circuits, data circuits, websites and mailboxes it can have. The Add asset window shows how much of each is already in use. To add a kind that is not in your package or to raise a limit, talk to 2N5.

Can I add a circuit myself?

The Add asset button at the top of the screen opens the choice of voice circuit, data circuit, website and email, within your package's limits. In practice, 2N5 usually adds and verifies circuits for you, because every new circuit goes through a real verification call and depends on 2N5's calling account.

USER GUIDE · CHAPTER 10

FAQ and help


This chapter gathers the questions customers ask most day to day and says where to find help. Inside Monitor, the help button opens a guided walkthrough; outside it, the 2N5 team helps you through the channel agreed when you signed up.

By [2N5 team](#) Updated October 14, 2026

The guided walkthrough

Monitor has a walkthrough that stops once on each page and explains what it shows.

1. At the top of the screen, click the **?** button (**Help and walkthrough**).
2. Click **Take the guided walkthrough**. It is nine short stops, one on each page.



Help is also on the screens themselves: the **i** icon next to a heading explains what the box shows and how its figure is calculated.

How to ask 2N5 for help

Monitor is a managed service, and 2N5 looks after your circuits together with you.

1. Before you write, note the incident's number, such as **INC-0004**, or the circuit's name, and what you saw on the screen.
2. Contact your 2N5 representative through the channel agreed when you signed up.
3. For an incident with your provider, copy the **Provider-ready evidence** block from the incident's page and send it with the ticket.

Requests about personal data, such as access, correction or deletion, go to privacidade@2n5.com.br. [Security and LGPD](#) explains how data is handled.

Further reading

- The [glossary](#) explains the terms that appear on the screens, such as [MOS-LQO](#), [jitter](#) and [maintenance window](#).
- The [Docs](#) gather this guide, the product pages and the reference for SIP codes and Q.850 causes, which appear in the observation log.

Frequently asked questions

Does Monitor call my numbers all the time? Does that get in the way?

Monitor makes one short test call to each circuit at the agreed interval, for example every 15 or 30 minutes, and repeats it only when a test fails. Whoever answers hears that it is an automatic test and the call ends. The hourly call budget on the dashboard shows how much of this hour's capacity has been used.

Why did Monitor not alert me on the first failure?

Because a single failure never becomes an alert. Monitor repeats the test and opens an incident only after four failed tests in a row, the first one and three retries, usually within a few minutes. That way you get no false alarm for a few seconds of trouble, and when an alert arrives the failure is real.

Where do the times on the screens and in the alerts come from?

Every time on the screens, in the alerts and in the reports uses your organization's time zone, shown at the top of the screen next to the clock, such as São Paulo, UTC-3. It is under Settings, on the Reports & clock tab. Your computer's own clock does not change what Monitor shows.

How long is call and contact data kept?

Each kind of data has a retention period and is deleted when it ends. A call's audio is kept only when the call is answered by a recording, never by a person. The periods and your rights are in the Privacy Policy; requests about personal data go to privacidade@2n5.com.br.

Can I use Monitor on my phone?

Yes. The screens fit a phone, with the sidebar opened from the button at the top. Day to day on a phone, the usual thing is to receive the alert by email and tap the link to acknowledge the incident, without signing in. The full dashboard is more comfortable on a bigger screen.

The product

HOW IT WORKS

How managed VoIP monitoring of voice circuits works

VoIP monitoring from Monitor by 2N5 is a managed service: you tell 2N5 which numbers to watch and who needs to know about a problem, and it proves each circuit with a real call and calls it day and night, in Brazil. You get the incident with evidence, after a retest that avoids the false alarm, and, for Enterprise customers, an SLA report every month.

By [2N5 team](#) Updated October 20, 2026

What do I need to do to get started?

Fill out a form. You send 2N5 the numbers to monitor, where each one belongs (company, office, department), and who should be notified when something fails. 2N5 does everything else: nobody at your company creates a service account, installs software on the PBX, or sets up alerts. That is how managed VoIP monitoring works in practice.

What are the steps?

1. **You send your circuits.** Numbers, sites, and who needs to know about a problem.
2. **2N5 verifies each one.** Every circuit is proven with a real call before it goes into monitoring, and the expected greeting is recorded as a reference.
3. **2N5 calls all night long.** At the agreed pace, every call is assessed: did it connect, did audio flow both ways, was it the right message, what was the measured quality. Failures are retried and false alarms are filtered out.
4. **You hear about it first.** An incident with evidence, the escalation ladder by email, and an SLA report every month.

Who does what?

You	2N5
Send the list of circuits and contacts	Adds each circuit and proves it with a real call
Say what should answer each number	Records the reference greeting and tunes the quality thresholds
Choose who gets notified, and in what order	Builds the escalation ladder and the reports
Receive the alerts and take ownership of incidents	Keeps the calls, the platform, and the monitoring itself running

What watches the monitoring itself?

If the fault is on 2N5's side, for example our calling platform not responding, the test is marked **blind**: it does not count against your circuit, does not go into the SLA, and does not wake your team. 2N5 is the one notified. A periodic heartbeat also alerts 2N5 if the monitoring stops running.

Does it work with my PBX and my carrier?

Yes. Monitor calls numbers, the way a customer does. That is why it works with any PBX or contact center platform and with any carrier, with no integration. It does not depend on what your equipment reports about itself. See how each platform tends to fail silently on the pages for [3CX](#), [Asterisk](#), [Genesys Cloud](#), [Amazon Connect](#) and [Microsoft Teams](#).

Where can I read more?

The [guide to SIP trunk, 0800 and IVR monitoring](#) explains the technique, the [synthetic monitoring](#) entry defines the term, and the [guide to VoIP monitoring tools in Brazil](#) compares Monitor with other tools. For Zabbix, see the [comparison](#).

Frequently asked questions

What is a managed voice monitoring service?

You give 2N5 the list of numbers, where each one belongs and who should be notified. 2N5 adds and proves each circuit with a real call, places the test calls, filters out false alarms, notifies your contact ladder and, for Enterprise customers, sends the SLA report every month. Nobody at your company installs software or sets up alerts.

How do I monitor a 3CX / Asterisk / Issabel PBX?

Without installing anything on the PBX: Monitor by 2N5 calls the numbers that come in through it, like a customer, and checks whether the call connects, whether audio flows both ways and whether the voice sounds clear. It works with 3CX, Asterisk, Issabel and others, any version, on-premises or hosted, because it does not depend on what the PBX reports about itself.

How do I monitor Genesys / Amazon Connect / Teams Direct Routing phone lines?

The same way, from the outside, by calling the numbers that reach the platform. Monitor by 2N5 does not integrate with Genesys Cloud, Amazon Connect or Microsoft Teams: it calls the numbers and checks the result like a customer, with the SIP code when a call fails. That is why it works with any contact center platform.

Do I need to install anything?

Not for voice circuits. 2N5 places the calls from its own platform. Only to track a site's internet link does 2N5 install a small agent there, over remote access and with your written authorization. No software goes on the PBX, the SBC or the voice network.

Do the test calls interfere with real calls?

Each test is a short call to a number you choose, such as a test tone, an echo service, or your greeting, within a call limit. It holds a channel for a few seconds, like a customer call, and your customers' calls are never touched.

How long does it take to get started?

As long as it takes you to send the list and 2N5 to verify each circuit with a real call. A circuit only goes into monitoring after it passes that check, and the expected greeting is recorded as the reference at that step. 2N5 agrees with you on the test pace and on who receives each alert.

Can my team see the dashboard?

Yes. Each manager signs in to Monitor and sees their part of the organization (company, office, or department), with the status of every circuit, the incidents, and the reports. People who are not managers still receive the alerts and the reports by email, without signing in to the dashboard.

What if I have an outsourced NOC?

Monitor does not replace whoever fixes the fault; it makes sure someone finds out about it, with the evidence. The NOC can be one of the tiers of the escalation ladder, receiving the confirmed alert with the SIP code and the Q.850 cause, and taking ownership of the incident with a tap.

Which tool monitors SIP trunks, internet links and websites together?

Monitor by 2N5 brings voice circuits (SIP trunks, 0800, 0300 and 4004), internet links, websites and email into one managed service, with the same alert ladder and the same dashboard. Other tools cover parts of this separately; the comparison pages show what each one does, with the source for each statement.

SIP TRUNK AND VOIP MONITORING

SIP trunk and VoIP monitoring with real test calls

Monitor by 2N5 does SIP trunk and VoIP monitoring by calling each trunk the way a customer does, 24 hours a day, in Brazil. When the call does not connect, does not carry audio both ways, or sounds bad, your team hears about it first, with the SIP code the carrier returned. Every failure is retested before the alert, so there are no false alarms.

By [2N5 team](#) Updated October 20, 2026

How do I know if a SIP trunk has stopped completing calls?

The direct way is to call it. Monitor places a real test call to each **voice circuit** (SIP trunk, branch extension, or contact center number) at the agreed pace and checks three things: whether the call connected, whether audio flowed both ways, and whether the voice sounded clear. It is SIP call monitoring done from the outside, the way the customer lives it. When any of the three fails and the failure is confirmed, your team gets the alert with the SIP code as proof.

Conventional monitoring asks the equipment whether it is up. The PBX says yes, the trunk shows as registered, and the network is green, while callers hear silence. A real call does not depend on what the equipment says about itself. The technique is called [synthetic monitoring](#), and the [guide to synthetic call monitoring](#) explains when to use it alongside passive monitoring, as the guide on [passive monitoring vs active test calls](#) compares in detail.

What does each test call check?

Check	What counts as a failure
The call connects	A carrier rejection (for example SIP 403, 404, 480, or 503), or no response at all
The right thing answers	The expected greeting, test tone, or echo service does not answer
Audio both ways	One-way audio or no audio (a silent call)
Voice quality	Measured MOS-LQO below the agreed threshold: the call counts as degraded

Each result carries the SIP code and the Q.850 cause. To read the most common ones, see [SIP 503](#), [SIP 486](#), [SIP 480](#), [SIP 408](#) and Q.850 causes [16](#), [34](#) and [127](#). If the call connects and then goes silent, the guide to [silent calls and one-way audio](#) shows the causes. For the time to the first ring, see [post dial delay](#), and for the share of calls answered, [ASR and ACD](#).

How does Monitor avoid false alarms?

A failed call is retried at your pace, a retest before alerting. Only consecutive failures, four by default (the first plus three retries), open an incident. A one-minute blip does not wake anyone up in the middle of the night, and an alert at 2 a.m. is always real.

If the problem is on 2N5's side, for example our calling platform not responding, the test is marked **blind**: it does not count against the trunk, does not go into the SLA, and does not notify the customer. Maintenance windows silence alerts while the carrier or your team works on the circuit.

What evidence goes to the carrier?

Each incident includes a block ready to paste into the ticket: the circuit, the carrier, the first symptom, the SIP code and Q.850 cause of every call, and the last healthy call. It is the answer to “no fault found”: the time, the number called, and what the route returned.

Sample evidence, with demo data: *SIP 503 Service Unavailable, Q.850 cause 34 (no circuit/channel available), confirmed after 4 failed calls.*

Who gets notified, and how?

Your company's escalation ladder, by email, one tier at a time, until someone acknowledges the incident with a tap on their phone. See [alerts and escalation](#).

How does Monitor compare with network and capture tools?

Network and capture tools answer different questions: what the equipment reports and what passed through the network. Monitor answers whether a real call connects. See the comparisons with [VoIPmonitor](#), [Zabbix](#) and [PRTG](#), and the [voice circuit monitoring guide](#).

Frequently asked questions

How do I monitor a SIP trunk 24/7?

With real test calls at a fixed interval, day and night, not just the registration status. Monitor by 2N5 calls a number answered by that trunk, checks whether the call connected, whether audio flowed both ways and whether the voice sounded clear, retries a failed call, and only then alerts your team, with the SIP code and the Q.850 cause.

Passive vs active VoIP monitoring: which do I need?

Passive monitoring analyzes the calls customers have already made, through packet capture or call records; it shows what happened, but sees nothing when nobody calls. Active test calls dial on a schedule and find the failure before the customer does, overnight included. If you need to know first, you need active, and the two complement each other.

What is synthetic call monitoring?

Synthetic monitoring, or synthetic call monitoring, is testing in which a system places real calls to your number on a schedule and records the result of each one: whether it connected, what answered, whether there was audio, and how the quality was. It is what Monitor by 2N5 does on every voice circuit, as opposed to passive monitoring.

What do SIP errors 503, 486, 480 and 408 and Q.850 causes 16, 34 and 127 mean?

SIP 503 says the server cannot handle the call right now, because of overload, maintenance or no free channel. 486 is busy, 480 is destination temporarily unavailable, and 408 is a timeout with no answer. In Q.850, cause 16 is normal clearing, 34 is no circuit or channel available, and 127 is an unspecified interworking failure. Monitor keeps both codes for every failure.

Does Monitor need access to my PBX or my SBC?

No. 2N5 places the test calls from its own platform, to the numbers you specify, just as anyone would. Nothing is installed on your PBX, your SBC, or your voice network, and the monitoring does not depend on what the equipment reports about itself. That is why it works with any PBX or contact center platform, with no integration.

How often is each trunk tested?

At the pace agreed with you, for example every 15 minutes, or at fixed times, including overnight and on weekends. A failed call is retried before any alert, and only four consecutive failed calls, the first plus three retries, open an incident. So a one-minute blip wakes nobody, and an alert in the middle of the night is always real.

Do the test calls take up channels on my trunk?

Each test is a short call to a number you choose, within an agreed call limit. It holds one channel on the trunk for a few seconds, like a customer call, and never touches your customers' calls. The test pace is agreed with you precisely so that the tests do not compete with your customer service for channels.

Does it work with any carrier?

Yes. The test dials the number, so it works with SIP and E1 trunks from any carrier and with any PBX, with no integration. The result of every call includes the SIP code and the Q.850 cause the route returned, which is exactly what the carrier needs to open or close a ticket.

What happens when the fault is on 2N5's side?

The test is marked as blind, never as a failure of your trunk. It does not count against your SLA, does not open an incident for you and does not wake anyone on your team; 2N5 is the one notified, if it lasts more than 30 minutes. A fault in the monitoring itself is never charged to your circuit.

0800 AND IVR TESTING AND MONITORING

0800 and phone IVR testing and monitoring with real calls

Monitor by 2N5 does toll-free number and IVR testing and monitoring with real calls, 24 hours a day, in Brazil. It calls your 0800, 0300, 4004, or landline number, checks that the phone IVR answers with the right message, and alerts your team before your customers notice, after retrying the call that failed.

By [2N5 team](#) Updated October 20, 2026

How do I know my 0800 number is working 24 hours a day?

By calling it, around the clock, like a customer. Monitor places a real test call to your 0800 (toll-free), 0300, 4004, or landline number at the agreed pace, overnight included, and checks that your **phone IVR** answered with the right message. A green dashboard on the PBX does not prove that; only a call coming in from the outside does.

Each monitored number is a **voice circuit** in Monitor, with its own schedule, its own quality threshold, and the people who need to know when it fails. The [0800 and IVR guide](#) walks through what to do when the number stops, and the one on [0800 calls that fail from mobile phones](#) covers a problem that only an outside test reveals. The terms are in the glossary: [0800](#), [IVR](#) and [numbers 4004, 4003 and 3003](#).

What can answer the number, and how is each case judged?

During onboarding, 2N5 defines what should answer each number. That choice decides what counts as working:

What should answer	Passes when
A greeting or recorded message (the IVR)	The call is answered and the chosen greeting is heard. A carrier's error announcement does not pass
A test tone or echo service	The call is answered and the tones sent come back
Other: may ring or be answered	Ringling out or being answered with audio, as agreed. Whoever answers hears that it is an automated test

Busy, no answer, or ringing out is a failure for an IVR: it answers every call on its own.

Why does the IVR need a reference recording?

Because "it answered" is not enough. A number can connect and play the carrier's announcement saying it does not exist, or an IVR can be swapped out without notice. 2N5 records the reference greeting for each number during onboarding, and every call compares what it heard with it. That same comparison measures audio quality, the [MOS-LQO](#).

What does Decree 11,034 (the SAC decree) require of phone support?

Decree 11,034/2022, in force since October 3, 2022 and until November 1, 2026, requires telephone access to the SAC, access 24 hours a day, seven days a week, through at least one integrated channel, which is not necessarily the phone, and human attendance for at least 8 hours a day. From November 2, 2026, Decree 13,135/2026 applies, which revokes it and keeps phone service mandatory, with live agents from 8 a.m. to 8 p.m. every day (art. 14), and requires at least one electronic channel available 24 hours (art. 12). Neither one says the 0800 must run 24 hours. For regulated industries, showing that the

0800 and the IVR were up is the hard part. Monitor's [SLA reports](#) show the availability of each number over the year and the month, with every failure and the moment someone took ownership. The [guide to the SAC decree and the RGC](#) lists the articles.

Who finds out when the 0800 goes down?

Your company's escalation ladder, by email, until someone acknowledges the incident. See [alerts and escalation](#).

How does Monitor compare with other 0800 and IVR testing tools?

The comparison is on the pages for [Cyara alternatives](#) and [Monitor by 2N5 vs Klearcom](#), with the source for every claim about the competitors and what was not verified.

Frequently asked questions

How do I know my 0800 number is down?

By calling it from the outside, at a fixed interval, and checking what answered. The PBX dashboard can be green while the carrier is not delivering calls. Monitor by 2N5 calls the number the way a customer does, compares the message with the reference recording and, if the failure repeats, alerts your team with the SIP code the route returned.

Is there a company that tests 0800 and IVR automatically in Brazil?

Yes. 2N5 offers Monitor, a managed service that calls, from Brazil, 0800, 0300, 4004, and landline numbers, listens to the phone IVR and compares the message with a reference recording. International toll-free testing services also exist; the site's comparisons show the differences, with the source for every claim.

How can I test a toll-free number automatically?

With an automated test call: a system dials the number on a schedule, listens to what answers and compares it with the recording of the expected greeting. If the message changes, or the carrier plays a number-not-found announcement, the test fails. In Monitor by 2N5 the reference is recorded during onboarding, and you do not have to call or program anything.

How often should IVR test calls run?

The shorter the interval, the less time an outage goes unseen. Each number has its own schedule in Monitor, at intervals such as 15 minutes or at fixed times, including overnight and on weekends, when nobody on the team is watching. A failed call is retried before any alert, so there are no false alarms.

Do Brazilian companies need a 24/7 SAC phone line?

Until November 1, 2026, Decree 11,034/2022 requires the SAC (customer service) to offer access 24 hours a day, seven days a week, through at least one integrated channel, not necessarily the phone, and live-agent phone service for at least 8 hours a day. From November 2, 2026, Decree 13,135/2026 requires an electronic channel available 24 hours and live-agent phone service from 8 a.m. to 8 p.m. The decree guide lists the articles.

What does Decree 11,034 (the SAC decree) require of phone service?

Until November 1, 2026, Decree 11,034/2022 makes phone service mandatory (art. 4, § 2), with at least 8 hours a day of live-agent service and complaint and cancellation options in the first menu (art. 5). From November 2, 2026, Decree 13,135/2026 applies and revokes it: live-agent phone service from 8 a.m. to 8 p.m., every day (art. 14). Neither one mentions the 0800.

How do I know if my company's phone number is working?

Call it from outside the company and from another network, because an internal call does not go through the carrier. To do that around the clock, use automated test calls: Monitor by 2N5 calls the number, 0800, 0300, 4004 or landline, checks whether it connected, whether there was audio and what answered, and alerts when the failure is confirmed.

What is the best IVR monitoring service?

It depends on what you need to test. Look for a service that really calls, from the number's own country, and does not just test the network; that compares the message with a reference recording, to catch a swapped IVR or a carrier announcement; that retries the call before alerting; and that gives the SIP code as proof. Monitor by 2N5 does all four.

How do I test 0800 numbers in Brazil from abroad?

An 0800 belongs to the Brazilian numbering plan, and a call coming from abroad may not connect, depending on the carrier, even when the number works. That is why a faithful test is one that starts in Brazil. Monitor by 2N5's calls leave from carriers in Brazil, like a local customer's, and you do not have to do anything for that.

Does the test call ever reach a live agent?

Not when the number is answered by the IVR or a recording: the call listens for the expected greeting and hangs up. If the circuit can be answered by a person, they first hear a message saying it is an automated test, and the call ends. A person's voice is never kept, out of respect for the privacy of whoever answers.

Do you keep the call audio?

Only when what answers is the recording or IVR you specified, to prove the line works and to measure quality. A person's voice is never kept. Stored audio is encrypted and kept for the period 2N5 sets, then deleted together with its record.

How do you know the IVR message has changed?

During onboarding, 2N5 records the reference greeting for each number. Every test call compares what it heard with that recording. A different message, such as a carrier announcement saying the number does not exist, or an IVR swapped without notice, does not pass, and the failure goes through the same retest and alert routine.

Do the tests cost my company call charges?

The test calls originate from 2N5's platform and are part of the service. For an 0800 number, as for any inbound call, the inbound rate follows your contract with the carrier. Monitor's plan neither includes nor replaces what your carrier charges for receiving the calls.

VOICE QUALITY

VoIP voice quality measured on every call, with MOS-LQO

Monitor by 2N5 measures the voice quality of every test call from the audio it received, compared with the expected recording, not estimated from network figures. The result is MOS-LQO, a score from 1 to 5, and your team is alerted when the call's audio quality drops below the agreed threshold and the drop is confirmed.

By [2N5 team](#) Updated October 14, 2026

How do you measure the voice quality of a SIP trunk or an 0800 number?

By listening to a real call. Monitor records the audio the test call received and compares it with the expected recording (your IVR greeting or the test tone) using a perceptual model that gives a score from 1 to 5, the **MOS-LQO**. It measures what a customer would hear, not a calculation made from network figures. That is how call quality monitoring catches a call that connects and sounds bad.

Measured or estimated: what's the difference?

	Estimated MOS (E-model)	Measured MOS-LQO (Monitor)
Where it comes from	Packet loss, jitter, and delay reported by the network	The audio the call received, compared with the reference
Standard	ITU-T G.107 (R-factor)	ITU-T P.800.1 (objective listening quality, reference-based)
Sees audio cut at the source	No	Yes
Sees a wrong message	No	Yes: the comparison with the reference fails

Both measures have their place. A network estimate is useful for looking at the network; a call that is actually listened to tells you whether the voice got through. See the glossary for [MOS](#), [MOS-LQO](#) and [POLQA and PESQ](#), and the guide to [MOS and MOS-LQO in practice](#).

What is a good score?

The MOS scale runs from 1 (bad) to 5 (excellent), as in ITU-T P.800. In practice, around 4 the voice sounds clear, and below 3 people start to complain. Monitor's default thresholds are:

Range	Reading
3.8 or higher	Healthy
From 3.0 to 3.8	Working, with the score flagged as below ideal
Below 3.0	Degraded: the test call counts as degraded

Where does quality show up?

On each circuit's page, with the score of every test call and the threshold lines; on the dashboard next to the circuit's status; and in the [SLA reports](#). A drop below the threshold goes through the same retries as any failure before it becomes an alert.

Why is the reference recording needed?

Because a perceptual model compares two signals: what should have been heard and what was heard. During onboarding, 2N5 records the reference for each number. The stored audio is that of the recording or the IVR, never a person's voice, and it is kept for the period 2N5 sets.

What about jitter, packet loss and delay?

The measured score shows the effect on the voice; [jitter](#) and [packet loss](#) help find the cause in the network. To see whether the degradation came from the site's internet link, see [internet link monitoring](#). If the problem is a silent call or one-way audio, the guide to [silent calls](#) lists the most common causes.

How does Monitor compare with tools that estimate MOS from the network?

See the comparison with [PRTG](#), which estimates the score from the network, and with [VoIPmonitor](#), which analyzes captured calls, each with the source for its claims.

Frequently asked questions

What is MOS in VoIP and what is a good score?

MOS (Mean Opinion Score) is the 1-to-5 score that summarizes the quality of a voice call, on the ITU-T P.800 scale. Around 4 the voice sounds clear, and below 3 people start to complain. A good score for your circuit is one that stays steady and well away from the alert threshold you agreed.

What is the difference between network MOS and MOS-LQO?

Network MOS is estimated: a model calculates the score from packet loss, jitter and delay, without listening to anything (ITU-T G.107). MOS-LQO is measured: a perceptual model compares the audio that arrived with the reference recording (ITU-T P.800.1). POLQA and PESQ are models of that kind, and scores are only comparable within the same model.

How do I measure the voice quality of a VoIP call?

By listening to the call, not just looking at the network. Place a test call, keep the audio that arrived and compare it with the expected recording using a perceptual model, which gives a score from 1 to 5. It is what Monitor by 2N5 does on every answered test call, around the clock, and it alerts when the score drops below the agreed threshold.

Which measurement model do you use?

A reference-based perceptual listening-quality model that returns MOS-LQO (ITU-T P.800.1). It compares the received audio with the reference audio and returns a score on the MOS scale. Because it uses a reference recorded during onboarding, the result also reveals a wrong message, audio cut off at the source and a silent call.

What are the thresholds?

By default, a circuit is healthy at a MOS-LQO of 3.8 or higher and degraded below 3.0; in between, the call passes with a note. 2N5 tunes the thresholds with you, per circuit, during onboarding, because a customer service line demands more of the voice than an internal test line does.

What happens when quality drops?

A score below the threshold marks the test call as degraded, and it goes through the same retries as any failure, the retest before alerting. If it persists, a quality incident opens and the escalation ladder is notified, with the score of each call as proof for the carrier.

LINE-DOWN ALERTS AND ESCALATION

Voice line-down alerts, escalated until someone takes ownership

Monitor by 2N5's line-down alert fires only for a confirmed failure, with no false alarms: a failed call is retried first. When a voice circuit really goes down, the first tier of your ladder is notified by email and, if nobody acknowledges in time, the alert moves up until someone takes ownership of the incident with a tap on their phone.

By [2N5 team](#) · Updated October 20, 2026

How do I make sure someone finds out when the line goes down overnight?

With an escalation ladder that does not stop until someone takes ownership. In Monitor, every confirmed incident notifies **tier 1** by email. If nobody acknowledges within the agreed time, it notifies tier 2, then tier 3. A single acknowledgment, one tap on the email link, stops the escalation and records who took ownership. It is the line-down alert that arrives before the customer's complaint.

What happens between the failure and the alert?

Moment	What Monitor does
A test call fails	Retries the call at your pace; nobody is notified yet
Consecutive failures (four by default)	Opens an incident with the evidence: SIP codes, Q.850 causes, every call
Incident opened	Notifies tier 1 and the customer contacts, by email
Nobody acknowledged in time	Notifies the next tier, up to the last one
Someone acknowledges with a tap	Stops the escalation; the incident shows who and when
The circuit recovers	Closes the incident and tells the customer contacts it has been resolved

The time to first acknowledgment goes into the [SLA reports](#) as the average response, to show how long the team takes to take ownership (see [MTTR](#) and [MTTA](#)).

Who is outside the ladder but still notified?

- **Customer contacts:** notified when a service has a problem, including poor voice quality, and when it is resolved. They are not part of the escalation.
- **Administrators:** notified when no tier acknowledges.

Why do only confirmed failures become alerts?

Because an alert at 2 a.m. has to be real, or the team stops paying attention. A one-minute blip, a route that failed once, or a problem on 2N5's own side never wakes anyone up. When the fault is 2N5's, the test is marked blind, does not count against the circuit, and 2N5 is the one that gets the notice. The [guide to line-down alerts without false alarms](#) explains the retest step by step.

Can alerts be silenced during maintenance?

Yes. A [maintenance window](#) covers the circuits you choose, at a scheduled time. During the window, failures on those circuits do not open an incident.

How does Monitor compare with other continuous testing services?

Retesting before alerting also exists in other continuous contact center testing services. See the comparison with the [Hammer VoiceWatch alternative](#), with the source for every claim.

Frequently asked questions

How do I get alerts when a phone line goes down?

Register the person's email as a contact in Monitor by 2N5. When an incident is confirmed, a message arrives with the evidence and a link, and one tap on it confirms you have taken ownership. Each person can have more than one address, and the alert is sent only after the retest confirms the failure, never for a single call that failed.

How do I reduce false alerts in call monitoring?

By confirming the failure before alerting. In Monitor by 2N5, a failed call is retried, a retest before alerting, and only four consecutive failed calls, the first plus three retries, open an incident. A one-minute blip, a route that failed once and a problem on 2N5's own side never wake anyone up.

How do I know my phone lines are down before customers do?

By calling the number from the outside, around the clock, instead of waiting for a customer to complain. Monitor by 2N5 places a test call per circuit at the agreed pace, overnight included, and when the failure is confirmed it notifies the first tier of your ladder by email, with the SIP code, before the first complaint arrives.

Which channels do alerts come through?

By email, with the incident's evidence and codes and a link: one tap on it confirms you have taken ownership of the incident, without signing in to the system. The notice goes out only after the retest confirms the failure, and the first layer of the ladder receives it before the others.

What happens when someone acknowledges?

Escalation stops. One acknowledgment, one tap on the email link, is enough: the next tiers are not woken up, and the incident shows who took ownership and when. The time to that first acknowledgment also goes into the monthly report as the team's average response.

What if nobody acknowledges?

The alert moves up tier by tier, on the timing you set for each one. If no tier acknowledges, your company's administrators are notified. That way an overnight incident does not sit with a single person who never saw the message.

How do I avoid alerts during scheduled maintenance?

With a maintenance window for the affected circuits, at a scheduled time. During the window, failures on those circuits do not open an incident and nobody is woken up. Outside the window, monitoring returns to normal on its own, and a failure that continues is handled like any other.

SLA REPORTS

Carrier SLA report for telephony, with the 12-month figure

Every month, Monitor by 2N5 delivers a carrier SLA report for telephony, with each voice circuit's availability for the period and the last 12 months, the incidents, the downtime, and the time until someone took ownership. It is the evidence to enforce the carrier's SLA with proof and to bring to the next contract negotiation.

By [2N5 team](#) Updated October 20, 2026

How do I produce a telephony SLA report by month and by year?

By measuring every circuit all the time and adding it up. Because Monitor calls each voice circuit at the agreed pace, it knows when each one was down and for how long. Every month it builds the service report with the availability for the period and the last 12 months, and emails it to the people you chose. Reports and the SLA are for Enterprise customers.

What is in the report?

Section	What it shows
Voice availability	The report period and the last 12 months, against the contract target
Incidents	How many, at what severity, and how long each one lasted
Downtime	The total for the period, per circuit
Average response	The time to first acknowledgment of each incident
The period, day by day	Which days had a confirmed failure on some circuit
By circuit and by carrier	The availability of each circuit and each carrier

Example with demo data: 99.89% availability for the month and 99.91% over the last 12 months, 99.90% target met, 4 incidents, 3 h 55 min of downtime, and an average of 6 minutes to first acknowledgment.

How much downtime fits in 99.9%?

In a 365-day year, 99.9% availability allows about 8 hours and 46 minutes of downtime; 99.95%, about 4 hours and 23 minutes; 99.99%, about 53 minutes. In a 30-day month, 99.9% allows about 43 minutes. The math is always the same: the length of the period multiplied by (100% minus the target). The [SLA calculator](#) does the math for any target and period, and the guide [Telephony SLA: yearly and monthly calculation](#) explains what counts as an outage. The [SLA glossary](#) entry sums up the terms.

How do I use the report to enforce the carrier's SLA?

Each incident in the report carries the start, end and duration, and each test call carries the SIP code and the Q.850 cause the route returned. With that, the ticket to the carrier stops being "the line went down" and becomes something like "the circuit failed from 02:14 to 02:51, with SIP 503 and cause 34" (example with demo data). The guide on [how to enforce a carrier SLA with proof](#) walks through it, and the [SAC decree guide](#) explains how to prove availability in regulated industries.

What is left out of the calculation?

The periods when the monitoring itself was blind because of a problem on 2N5's side. They are taken out of the calculation and shown separately, because monitoring blindness is not an outage of your circuit, and an honest SLA figure cannot mix the two.

How does Monitor compare with SLA tools?

See the comparison with [Sipfront](#), which also has an SLA feature, with the source for every claim.

Frequently asked questions

How do I calculate the availability SLA of my 0800 number?

Divide the time the 0800 completed calls by the total time in the period and multiply by 100. In 30 days (43,200 minutes), 90 minutes of downtime gives 99.79%. The hard part is measuring downtime: Monitor by 2N5 measures it with real calls, circuit by circuit, and adds up the duration of every confirmed incident.

How do I prove a carrier breached its SLA?

Take the carrier the report and, for each incident, the start and end time, the number called, the SIP code and Q.850 cause of every test call, and the last healthy call. It is the record that is usually missing when the carrier says it found no fault. Check the contract for which measurement counts and whether scheduled maintenance is left out.

How do I calculate 99.9% uptime per month?

In a 30-day month, 99.9% allows about 43 minutes of downtime; in a 365-day year, about 8 hours and 46 minutes. The math is always the length of the period multiplied by (100% minus the target): 43,200 minutes times 0.1% is 43.2 minutes. The site's SLA calculator does it for any target and period.

Why does the report show the last 12 months?

Because in Brazil the telephony SLA is usually measured over the year. The 12-month figure is what decides whether the target was met; the month and the last 30 days are shown for reference. On a contract with a monthly SLA, the month's figure counts, and it is in the report too.

Who receives the report?

The people your company chooses, by email, each in their own language, English or Portuguese. A branch manager can receive only the report for their part of the organization, so each area sees its own circuits without exposing the other areas' figures.

Does a monitoring failure count against me?

No. A period when 2N5 could not test the circuit because of a problem on 2N5's side is left out of the calculation and shown separately. Monitoring blindness never becomes downtime for your circuit, and an honest SLA figure does not mix the two.

DATA CIRCUITS

Internet link monitoring for every branch and site

Monitor by 2N5 does internet link monitoring for every branch with a small agent that 2N5 installs on site. It answers the first question of every outage, was it the voice line or the connection, by measuring latency, packet loss, jitter and speed, and it alerts your team when the link goes down.

By [2N5 team](#) Updated October 14, 2026

How do I tell whether the outage was the internet link or the voice line?

By measuring both. In Monitor, the **data circuit** at each site, the dedicated internet link or broadband, is tracked by an agent installed right there, alongside the test calls to the voice circuits. It is internet link monitoring for every branch, tied to what the voice feels. The guide to [failing calls](#) shows how to tell one cause from the other.

Each voice circuit can be linked to its site's internet link. With the wait-for-the-link option, alerts for that circuit hold while the link has an open outage: your team gets one alert about the link, not one for every voice circuit that runs through it.

What does the agent measure?

Measurement	What it is for
Latency	Shows delay on the route to the internet
Packet loss	Explains choppy audio and dropped calls
Jitter	Variation in delay, which ruins voice before it drops the call
DNS time	Shows when names do not resolve, even with the network up
Speed	The actual speed, to compare with what you pay for

The measurements appear in charts covering the last 100 hours on each data circuit's page, and a confirmed link outage opens an incident and climbs the same [escalation ladder](#) as the voice circuits. The terms are in the glossary: [jitter](#) and [packet loss](#).

How is the agent installed?

At your company's written request, 2N5 installs the agent on a computer at the site, over remote access. It opens no inbound port: it only sends measurements to 2N5. When the site stops being monitored, 2N5 removes the agent itself.

What about a site where the agent cannot be installed?

2N5 tests the link from the outside. At every interval it attempts a connection to the site's public address, and a series of unanswered attempts is treated as a link outage. It is a simpler check than the agent's: it tells you whether the site responds, without measuring latency, loss, or speed.

What does the agent never collect?

Files, browsing, traffic content, usernames, programs, Wi-Fi network names, and other devices on the network. The details are in the [Privacy Policy](#) and in [security and LGPD](#).

How does Monitor compare with network monitoring tools?

See the comparison with [Obkio](#), an agent-based synthetic network monitoring tool, with the source for every claim.

Frequently asked questions

How do I tell whether a call dropped because of the carrier or my own network?

Compare both sides at the same time: the test call, placed from the outside, and the measurements of the site's link. If the test call fails and the site's link shows loss or is down, the problem is in the site's connection. If the link is healthy and the call fails, the cause is on the voice carrier's route. Monitor by 2N5 shows both sides together.

How do I monitor the internet link at several branches?

By installing an agent at each site, which measures the link's latency, packet loss, jitter, DNS and speed and sends the measurements to the same dashboard. With Monitor by 2N5, 2N5 installs the agent over remote access, with your written authorization, and each manager sees only their part of the company. Where installing is not possible, 2N5 tests the link from the outside.

What does the agent measure?

Latency, packet loss, jitter, DNS response time, and the actual connection speed, plus the CPU and memory usage of the computer it runs on. The measurements appear in charts covering the last 100 hours, and a confirmed link outage opens an incident and climbs the same escalation ladder as the voice circuits.

Does the agent read my traffic or my files?

No. It never collects files, browsing, traffic content, usernames, installed programs, Wi-Fi network names, or other devices on the network. It only measures the connection, opens no inbound port, and sends nothing but the measurements to 2N5, so your traffic and your files stay where they are.

Who installs the agent?

2N5 does, over remote access, on a computer at the site, after written authorization from your company. You do not download or configure anything. When the site stops being monitored, 2N5 removes the agent itself. If installing is not possible, 2N5 uses the simpler outside check.

How long are the measurements kept?

For 13 months, so the yearly reports are complete, with month-to-month comparison. After that period the measurements are no longer kept. The details of what is kept, for how long, and how to exercise your rights are on the security and LGPD page and in the Privacy Policy.

WEBSITES

Website monitoring: downtime and content changes

Monitor by 2N5 visits your website every five minutes and alerts your team when it goes down. With one extra switch, it also compares the content with a reference and shows the website as Content changed when the page changes without authorization, as in a hacked site, with swapped text or a new script.

By [2N5 team](#) Updated October 21, 2026

How do I know whether my website is down or has been altered?

By checking two things at every interval: whether the website answers, and whether the page is still the one you approved. In Monitor, a **website** is added with just its address. 2N5's servers visit it from the outside, as a visitor would, and a confirmed failure opens an incident and climbs the same [escalation ladder](#) as the voice circuits.

How does Monitor know the website is down?

Every five minutes by default, Monitor visits the address and waits for the page to answer. Each result has a state.

Result	What Monitor does
The website answers normally	Records the response time and moves on
The website answers slowly, or the certificate expires in a few days	Marks it degraded: a warning that does not count as downtime
The address does not exist, the connection is refused, the secure connection fails, the server does not answer in time, or it returns an error	Repeats the check; consecutive failures open a downtime incident
The certificate has expired	Counts as an outage

The interval, the number of retries, and the slowness limit can be adjusted per website. Two weeks before a certificate expires, your company's contacts already get a notice.

If the problem is on 2N5's own side, for example the checker having no internet access, the result is recorded as not judged: it does not count against your website, and 2N5 is the one notified.

What if the website stays up but has been altered?

That is the hacked or tampered website: the page opens normally, but it is no longer yours. An availability check does not notice, because the server answers. That is why Monitor offers, for each website, the **content check**: one switch, **Also check this website's content**, when you add or edit the website.

When you switch it on, Monitor reads the live page three times and keeps the result as the **reference**. There is nothing to upload and no rules to write: the reference is always the page that is live at that moment, and the check follows the same schedule as the website's own test. At every check, it reads the website again and compares:

- **Defacement or error phrases**, from a default list kept by 2N5.

- **Text and structure:** how much of the text changed and how much of the page structure stays the same, within fixed limits.
- **Unexpected scripts:** a script from an address that was not in the reference, a new file, code written into the page that changed, or a new event handler. The alert says “possible code injection”.
- **Redirects:** the website now ends up at another address.

A small change within the limit raises no alert: the reference follows the page. Parts that change by themselves, such as a clock or a visitor counter, are already ignored. A page built entirely by JavaScript cannot be content-checked: it reads as unreadable.

What does your team receive when the content changes?

The website shows as **Content changed**, a state worse than Up and better than Down, and an incident of its own, which Monitor handles like any other: it alerts the first tier by email, with a link to acknowledge in one tap, moves up a tier if nobody takes ownership, and tells your company's contacts when the page is back to normal. [Maintenance windows](#) silence the alerts during planned work.

If the website goes down at the same time, there is a single incident, the outage, and the content change stays as a line in the observation log. A slow website plus a content change are two incidents. The website's page has two content sections, under Answers and Certificate: the state and what changed, and the checks. The incident page shows where on the site the text changed, what was removed and what was added, and which scripts are new. There is a single question: *was this change made on purpose?* If it was, **This change is expected** replaces the reference and closes the incident.

How does the website appear in the report?

The [monthly report](#) has an online services page with each website's availability, calculated like the data circuits': an outage is downtime, degraded is shown but does not count, and time when 2N5 could not check is left out of the calculation. It also shows the average response time.

Content changes have a **section of their own**. They never enter availability, because an altered website is up, and they appear with the number of incidents, how long the page differed from the reference, how many changes were accepted, and the date of the last one. The availability calculation follows the [SLA](#) concept.

How does this fit with the rest of Monitor?

The website is one more item on the same account, beside the voice circuits and the [internet links](#): the same contacts, the same ladder, the same maintenance windows, the same report. For the company's mailbox, see [email monitoring](#). How 2N5 runs the service for you is described in [How it works](#).

Frequently asked questions

What happens if my website is altered?

The website shows as Content changed and Monitor opens a content incident, which alerts the first tier of your ladder by email. If nobody acknowledges in time, the alert moves up tier by tier. The incident shows what changed on the page, in which part of it, and which scripts appeared.

How does Monitor know the website was hacked?

It does not guess: it compares. The page it reads now is checked against the stored reference, and Monitor alerts you when the text or structure changes beyond the limit, a defacement or error phrase appears, the website ends up at another address, or a script appears that was not in the reference. The limits are fixed: there is nothing to configure.

Can I accept an expected change?

Yes. On the incident, the “This change is expected” button makes today's page the new reference and closes the incident, with your name recorded. There is no approval queue: whoever has access decides on the spot, and small changes within the limit are followed automatically.

Does a content change count as the website being down?

No. An altered website is still up, so a content incident never enters the downtime calculation. It appears in its own section of the monthly report. If the website goes down while its content is changed, there is a single incident, the outage, and the change stays noted in the observation log.

Does a slowdown or a brief error trigger an alert?

No. A check that fails is repeated, and only consecutive failures open an incident. A slow website, or one whose certificate is close to expiring, shows as degraded, a warning that does not count as downtime. If the failure is on 2N5's own side, the result is recorded as not judged and does not count against your website.

How do I detect website defacement?

By comparing the live page with a saved reference on every check. Monitor by 2N5 alerts when the text or structure changes beyond the limit, a defacement or error phrase appears, the website ends up at another address or a new script shows up, and it opens a content incident. Defacement is caught even while the site stays online.

How do I monitor whether my website is down and get an alert?

With automatic checks run from outside your company at a fixed interval. Monitor by 2N5 loads the site, repeats a failed check and opens an incident only after consecutive failures, which avoids false alarms. The alert reaches the first layer of your contact ladder by email and moves up a layer if nobody acknowledges in time.

EMAIL

Email monitoring: are your messages arriving?

Monitor by 2N5 sends one test email an hour to a mailbox of your company, checks where it landed, and deletes it. If the message does not arrive, lands in spam, or the mailbox is filling up, your team is alerted before a customer complains.

By [2N5 team](#) Updated October 20, 2026

How do I know whether my company's emails are arriving?

By sending a real email and checking that it arrived. In Monitor, an **email check** uses a test address of your company, a mailbox nobody reads. By default, every hour, 2N5's sender emails it a message, Monitor looks for that message in the mailbox, and records how long it took to arrive and which folder it landed in. That answers the question of anyone who depends on email: *can the people writing to us reach us?*

How does the round-trip test work?

1. **Send.** Monitor sends the test address a plain message, with no links, images, or trackers. The sender is named "Monitor by 2N5", and the subject is in Portuguese when the address ends in `.br` and in English otherwise. The message carries an identifier unique to that round, in a header and on the last line of the text.
2. **Look.** About once a minute, Monitor searches the mailbox for that identifier, the spam folder included.
3. **Verdict.** It reached the inbox: passed. It landed in spam or arrived after the check's allowed time: degraded. It did not appear in fifteen minutes: Monitor sends another test and, if that one does not arrive either, opens an incident.
4. **Clean-up.** The message that was found goes to trash at once, along with any earlier test that arrived late. Nobody at the company sees a test.

Each round keeps the delivery time and the result of the sender checks made by the receiving server (SPF, DKIM, and DMARC), so that a spam result can say why.

What kinds of mailbox can Monitor follow?

Account type	What Monitor does
Google account or Google Workspace	Whoever administers the mailbox signs in with Google once, through a connect link valid for 7 days, and Monitor then looks in it for the test message. No server, port or password is asked for
Sending only (SMTP)	Reads no mailbox: each round, the company's own outgoing server sends one test from that address to a 2N5 mailbox, where Monitor looks for it, inbox or spam, and deletes it. It reads down when the server cannot be reached, will not encrypt, refuses the login, refuses to send, or the test does not arrive, and degraded when it lands in spam
Mailbox from another provider, over IMAP	Reads the mailbox with a login and password, the spam folder included, and can also check that the company can send, as above, with the same login

In practice, 2N5 does the setup with you, as with the rest of the [managed service](#). The mail provider is always the one your company already uses.

What happens when the mailbox fills up?

When the mailbox reports its space used and its limit, or you enter its size, Monitor shows the percentage used and compares it with two marks on each check:

- **Warning mark:** the responsible people are told once each time the mark is crossed, and the check stays up.
- **Alert mark:** the check starts failing and climbs the escalation ladder. The same goes for a test that comes back bounced because the mailbox is full.

Without a known limit, Monitor shows only the size, never an invented percentage. At 100%, anyone writing to the address gets a bounce, and the warning screen reminds you to free space or raise the quota with the provider.

Which alerts does your team receive?

An email that does not arrive opens an incident and triggers the [escalating alerts](#): email to the first tier, with one-tap acknowledgment, and a climb tier by tier if nobody takes ownership. When messages arrive again, your company's contacts are told it is resolved.

Spam, delay, and a nearly full mailbox are warnings. They inform the contacts without waking the ladder. A [maintenance window](#) silences the check during planned work.

How does email appear in the report?

The [monthly report](#) includes the email check on its online services page: availability, calculated like the data circuits', the share of tests that reached the inbox, and the delivery time. The same [SLA](#) logic applies: an outage is downtime, a warning is shown but does not count, and time when 2N5 could not check is left out of the calculation.

What does email monitoring not do?

The test confirms that email sent **to** the company arrives. It does not measure whether the messages the company sends land in other people's inboxes. For the company's website, see [website monitoring](#).

Frequently asked questions

Does the test email land in the inbox?

That is what is expected, and Monitor checks it. If the test lands in spam or arrives later than the check allows, the result reads degraded, a warning that does not climb the ladder. The test is written like an ordinary email, with no links or images, and is deleted as soon as it is found.

What happens when the mailbox fills up?

The check reads the space used and compares it with two marks, a warning and an alert. At the warning mark, people are told once and the check stays up. At the alert mark it starts failing and climbs the ladder, because with a full mailbox anyone writing to the company gets a bounce.

Does anyone at the company have to read the test emails?

No. The test address is a mailbox nobody reads, and each message is moved to trash as soon as it is found, along with any earlier test that arrived late. The test text says it is automatic and needs no reply.

Does Monitor check the emails my company sends?

The round-trip test confirms that email sent to the company arrives. Optionally, Monitor also checks sending: it logs in to the company's outgoing server and asks it to accept a message from that address, but it stops before anything is sent.

Does a problem on 2N5's side count as email being down?

No. If 2N5's sender cannot send the test, or the mailbox can no longer be read, the result is recorded as not judged: it does not count against your company, and 2N5 is the one notified. That way, your team is only woken by problems that really affect the company's email.

How do I monitor email delivery with round-trip tests?

Send a test email to a company mailbox on a schedule and check where it landed. Monitor by 2N5 sends one test email an hour, confirms whether it reached the inbox or the spam folder, and deletes it right after. If it does not arrive, arrives late or the mailbox is filling up, your team is alerted before customers complain.

SECURITY AND LGPD

Security and LGPD in VoIP monitoring at Monitor by 2N5

VoIP monitoring from Monitor by 2N5 processes the minimum personal data needed to watch voice and data circuits in Brazil. This page summarizes what is kept, for how long, and how it is protected, and the main rule: a person's voice is never kept. The full, binding text is in the Privacy Policy.

By 2N5 team Updated October 14, 2026

Who is responsible for the data?

It depends on the data. 2N5 is the **controller** of the data it decides to process itself: demo requests, user accounts, and access and security logs. For the monitoring data a customer provides or that the service generates for them, such as alert contacts, numbers called, recordings, and reports, the customer company is the controller and 2N5 is the **processor**, handling the data only on the customer's instructions.

What is kept, and for how long?

Data	How long
Demo request	12 months after the last contact
Technical record of each test call	The period 2N5 sets
Stored audio (recording or IVR only)	The period 2N5 sets, encrypted
Alert messages sent and data circuit measurements	13 months
Access logs (Marco Civil da Internet, Brazil's Internet Civil Framework, art. 15)	6 months
Verdicts, quality scores, and incidents	For the duration of the contract

What is the call audio rule?

A person's voice is never kept. The audio of a test call is kept only when what answers is the recording or IVR the customer specified, and only to prove the line works and to measure quality. If anything else answers, including a person, the audio is destroyed at the end of the call and only the result is kept. A call that may be answered by a person first plays a message saying it is an automated test.

How is the data protected?

- All traffic to and from the service is encrypted.
- Audio, recordings, passwords for connected services, and signing keys are each encrypted individually and at rest, with keys kept separate from the data.
- Each company's data is kept separate, and each manager sees only the part of the company assigned to them.
- 2N5 administrators sign in with two-step verification; any manager can turn it on, and a company can require it for everyone.
- Repeated failed sign-in attempts pause access, and changing a password ends all other sessions.
- Administrative access to the servers is restricted and audited, and daily backups are kept for 14 days.

What about the agent installed at sites?

It is installed only with the customer's written authorization, opens no inbound port, and measures only the connection: latency, loss, jitter, DNS, and speed. It never collects files, browsing, traffic content, usernames, or other devices on the network.

How do I exercise my rights?

Write to 2N5's data protection officer at privacidade@2n5.com.br or call **+55 (11) 2050-1096**. The [LGPD](#) page explains each right and the response time, and the [Privacy Policy](#) has the full text.

Where can I read more?

The [guide to SIP trunk, 0800 and IVR monitoring](#) shows how the test calls work, the [synthetic monitoring](#) entry defines the term, and the [comparisons](#) show how Monitor compares with other tools.

Frequently asked questions

Does Monitor record the conversations on my phone system?

No. It only places its own test calls and does not listen to your customers' calls or to the conversations on your phone system. The audio of a test call is kept only when what answers is the recording or IVR you specified, and a person's voice is never kept.

Does VoIP monitoring need access to my company's voice network?

No. Monitor by 2N5 does not enter the PBX, the SBC or the voice network: it only calls the numbers you specify, from the outside, like a customer. The only software installed, and only with your written authorization, is the small internet link agent, which opens no inbound port and measures only the connection.

Where is the data stored?

With cloud infrastructure providers hired by 2N5, some of them outside Brazil, including in the United States. Transfers follow Article 33 of the LGPD (Lei Geral de Proteção de Dados, Brazil's data protection law), using the standard contractual clauses approved by the ANPD, Brazil's data protection authority. Each company's data is kept separate from the others', and each manager sees only the part assigned to them.

Does the Monitor website use tracking cookies?

No. This website does not set cookies, and it does not load analytics or advertising tools. It keeps only the theme you choose (light, dark or your computer's setting) in this browser, so the website and the product open in the same theme; that choice is never sent to 2N5. The product stores only what it needs to work, such as keeping you signed in.

How do I request my data or its deletion?

Write to 2N5's data protection officer at privacidade@2n5.com.br or call **+55 (11) 2050-1096**. The [LGPD](#) page explains each right, such as access, correction and deletion, and how quickly we respond, and the [Privacy Policy](#) has the full text, with the retention period for each kind of data.

PLANS

Monitor by 2N5 plans for 0800, IVR and SIP trunk monitoring, and how to buy

Monitor by 2N5 has published monthly plans for 0800, IVR and SIP trunk monitoring, in US dollars for the USA and Canadian dollars for Canada. Each plan includes a number of monitored channels and daily tests per channel, and 2N5 handles onboarding as a managed service. Choose a plan below and book a demo on your own numbers.

By [2N5 team](#) Updated October 20, 2026

Plans

On the website, this page's plan cards show each plan's monthly price and what it includes. Use the USA | Canada toggle to pick your country: USA prices are in US dollars with tax included, and Canada prices are in Canadian dollars plus applicable taxes. Each plan includes a number of **monitored channels** and **daily tests per channel**. A channel is a voice circuit, a data circuit, a website, or an email check.

If a plan is not enough, there are two add-ons, billed monthly:

- **Extra channel:** one more monitored channel.
- **Extra daily test:** one more test per day, billed per channel.

The monthly price is: plan price + extra channels × extra-channel price + extra daily tests × extra-daily-test price × number of channels (included plus extra).

For example, a plan with 1 channel, plus 2 extra channels and 1 extra daily test: plan price + 2 × extra channel + 1 × extra daily test × 3 channels.

The **Enterprise** plan is on request, priced per organization.

What each plan includes, such as its alert channels, the escalation list and reports with the SLA, is on its card, which comes from 2N5's own plan table. If a card does not list something, the plan does not include it.

What is included?

Included	How
Onboarding	2N5 adds each circuit, proves it with a real call, and records the reference
Test calls	From 2N5's platform, at your plan's daily test pace
Measured voice quality	MOS-LQO for every call, against the expected recording
Alerts	Through the channels the plan's card lists, with a retest before alerting
Evidence	SIP code and Q.850 cause for every failure, ready for the carrier
Reports and SLA	On the plans whose card lists them, every month, in English or Portuguese
Dashboard access	For your company's managers, each in their own part

How do I buy?

1. Choose a plan and click the button on its card.

2. Book the demo: the request reaches 2N5 with the plan you chose. It is three questions, about a minute.
3. 2N5 shows you Monitor with the numbers you chose.
4. You send the form with the circuits and contacts, and 2N5 handles onboarding.

Want to compare before choosing a plan?

The [guide to VoIP monitoring tools in Brazil](#) and the [comparisons](#) show, with the source for every claim, where Monitor differs from other tools. The [SIP trunk, 0800 and IVR monitoring guide](#) explains what is monitored, and the [SLA](#) entry defines the term used in the reports.

Frequently asked questions

How much does it cost to monitor 0800 numbers and SIP trunks?

It depends on the plan: the monthly plans are published on this page, in US dollars with tax included for the USA and in Canadian dollars plus applicable taxes for Canada. Each plan includes a number of monitored channels and daily tests per channel; a channel is a voice circuit, a data circuit, a website, or an email check. If you need more, you add extra channels and extra daily tests. The Enterprise plan is on request, priced per organization.

Is there a free trial?

There is the demo: 2N5 sets up Monitor with the numbers you choose and shows how it would handle a failure on them, with the SIP code and the alert your team would receive. Because it is a managed service, there is no self-service sign-up to try it on your own.

Is onboarding charged separately?

Onboarding is done by 2N5 as part of the managed service. You send the form with the circuits and contacts; 2N5 adds them, verifies each circuit with a real call, records the reference greeting, and sets up the alert ladder and the reports.

Who signs the contract?

Your company, with 2N5 Tecnologia da Informação e Telecomunicações Ltda., CNPJ (Brazilian company registration) 26.593.241/0001-85. The usage rules are in the Terms of Use, and data processing is covered in the Privacy Policy, and the security and LGPD page explains what is kept and for how long.

Reference

REFERENCE

SIP response codes: what each one means

A SIP response code is the three-digit number a server sends back to a call: 1xx while it is in progress, 2xx when it connects, and 3xx to 6xx when it does not. This is the full list from the IANA registry, with how Monitor reads each one on a test call.

How do you read a SIP code?

The first digit gives the class: 4xx is a refusal of the request (number does not exist, busy, credentials), 5xx is a failure of the server that tried to connect the call (overload, route, maintenance) and 6xx is a final refusal at any destination. The Q.850 cause that usually comes with it, in the Reason header, gives the reason on the telephone network's side.

In the "Reading on a test call" column, Down means the call counts as a circuit failure; Monitor repeats it every 2 minutes before it opens an incident. "Depends" means the reading changes with what should answer the number, and the code's page explains each case. 1xx and 2xx codes are not failures.

1xx: provisional (the call is in progress)

Code	Official name	Defined in	Reading on a test call
SIP 100	Trying	RFC 3261	Not a failure
SIP 180	Ringing	RFC 3261	Not a failure
SIP 181	Call Is Being Forwarded	RFC 3261	Not a failure
SIP 182	Queued	RFC 3261	Not a failure
SIP 183	Session Progress	RFC 3261	Not a failure
SIP 199	Early Dialog Terminated	RFC 6228	Not a failure

2xx: success

Code	Official name	Defined in	Reading on a test call
SIP 200	OK	RFC 3261	Not a failure
SIP 202	Accepted (Deprecated)	RFC 6665	Not a failure
SIP 204	No Notification	RFC 5839	Not a failure

3xx: redirection

Code	Official name	Defined in	Reading on a test call
SIP 300	Multiple Choices	RFC 3261	Down
SIP 301	Moved Permanently	RFC 3261	Down
SIP 302	Moved Temporarily	RFC 3261	Down

Code	Official name	Defined in	Reading on a test call
SIP 305	Use Proxy	RFC 3261	Down
SIP 380	Alternative Service	RFC 3261	Down

4xx: request failure

Code	Official name	Defined in	Reading on a test call
SIP 400	Bad Request	RFC 3261	Down
SIP 401	Unauthorized	RFC 3261	Monitoring blind
SIP 402	Payment Required	RFC 3261	Down
SIP 403	Forbidden	RFC 3261	Down
SIP 404	Not Found	RFC 3261	Down
SIP 405	Method Not Allowed	RFC 3261	Down
SIP 406	Not Acceptable	RFC 3261	Down
SIP 407	Proxy Authentication Required	RFC 3261	Monitoring blind
SIP 408	Request Timeout	RFC 3261	Down
SIP 410	Gone	RFC 3261	Down
SIP 412	Conditional Request Failed	RFC 3903	Down
SIP 413	Request Entity Too Large	RFC 3261	Down
SIP 414	Request-URI Too Long	RFC 3261	Down
SIP 415	Unsupported Media Type	RFC 3261	Down
SIP 416	Unsupported URI Scheme	RFC 3261	Down
SIP 417	Unknown Resource-Priority	RFC 4412	Down
SIP 420	Bad Extension	RFC 3261	Down
SIP 421	Extension Required	RFC 3261	Down
SIP 422	Session Interval Too Small	RFC 4028	Down
SIP 423	Interval Too Brief	RFC 3261	Down
SIP 424	Bad Location Information	RFC 6442	Down
SIP 425	Bad Alert Message	RFC 8876	Down
SIP 428	Use Identity Header	RFC 8224	Down
SIP 429	Provide Referrer Identity	RFC 3892	Down
SIP 430	Flow Failed	RFC 5626	Down
SIP 433	Anonymity Disallowed	RFC 5079	Down
SIP 436	Bad Identity Info	RFC 8224	Down
SIP 437	Unsupported Credential	RFC 8224	Down
SIP 438	Invalid Identity Header	RFC 8224	Down
SIP 439	First Hop Lacks Outbound Support	RFC 5626	Down
SIP 440	Max-Breadth Exceeded	RFC 5393	Down
SIP 469	Bad Info Package	RFC 6086	Down

Code	Official name	Defined in	Reading on a test call
SIP 470	Consent Needed	RFC 5360	Down
SIP 480	Temporarily Unavailable	RFC 3261	Depends on the circuit
SIP 481	Call/Transaction Does Not Exist	RFC 3261	Down
SIP 482	Loop Detected	RFC 3261	Down
SIP 483	Too Many Hops	RFC 3261	Down
SIP 484	Address Incomplete	RFC 3261	Down
SIP 485	Ambiguous	RFC 3261	Down
SIP 486	Busy Here	RFC 3261	Depends on the circuit
SIP 487	Request Terminated	RFC 3261	Depends on the circuit
SIP 488	Not Acceptable Here	RFC 3261	Down
SIP 489	Bad Event	RFC 6665	Down
SIP 491	Request Pending	RFC 3261	Down
SIP 493	Undecipherable	RFC 3261	Down
SIP 494	Security Agreement Required	RFC 3329	Down

5xx: server failure

Code	Official name	Defined in	Reading on a test call
SIP 500	Server Internal Error	RFC 3261	Down
SIP 501	Not Implemented	RFC 3261	Down
SIP 502	Bad Gateway	RFC 3261	Down
SIP 503	Service Unavailable	RFC 3261	Down
SIP 504	Server Time-out	RFC 3261	Down
SIP 505	Version Not Supported	RFC 3261	Down
SIP 513	Message Too Large	RFC 3261	Down
SIP 555	Push Notification Service Not Supported	RFC 8599	Down
SIP 580	Precondition Failure	RFC 3312	Down

6xx: global failure

Code	Official name	Defined in	Reading on a test call
SIP 600	Busy Everywhere	RFC 3261	Depends on the circuit
SIP 603	Decline	RFC 3261	Depends on the circuit
SIP 604	Does Not Exist Anywhere	RFC 3261	Down
SIP 606	Not Acceptable	RFC 3261	Down
SIP 607	Unwanted	RFC 8197	Down
SIP 608	Rejected	RFC 8688	Down

Sources: [IANA SIP response code registry](#) and [RFC 3261](#), checked on October 7, 2026. See also the [list of Q.850 causes](#) and the [SIP code lookup](#).

SIP CODES

SIP 180 Ringing: what it means and how to fix it

By 2N5 team Updated October 7, 2026

Short answer: SIP 180 Ringing is a provisional response: the destination received the INVITE and is alerting the user, which means the phone is ringing. It is not a failure. The trouble starts when the 180 repeats and nothing follows it. In that case, check whether the call ends in 200 OK, 480 or 487.

SIP 180 in the reference

Code	180
Official name	Ringling
Class	1xx: provisional (the call is in progress)
Defined in	RFC 3261

How Monitor treats SIP 180

SIP 180 on its own is not a failure: the call is still in progress. What counts is how it ends. The table shows how a test call that rings until the test stops waiting is read for each kind of circuit, depending on what should answer the number.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Up

For "Other": With "Count ringback as reachable" unticked: Down.

What Monitor records for an IVR: "The line rang and nothing answered, but this circuit is supposed to answer every call by itself."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

What does SIP 180 mean in practice?

SIP 180 Ringing is the provisional response that says "the destination is ringing." Under RFC 3261, the agent that received the INVITE is trying to alert the user, and the calling side can use this signal to generate the ringback tone locally. The 180 creates a provisional dialog (early dialog), but it does not settle anything: the call is still waiting for a final response, which can be 200 OK (answered), a 4xx, 5xx or 6xx error, or a CANCEL from the caller.

That is why a 180 on its own is not a diagnosis. It tells you the signaling reached something capable of ringing. What matters to anyone monitoring a trunk or an 0800 number is what happens after the 180: whether the destination answers, how long it takes, and whether the call keeps ringing forever without anyone, not even the IVR, picking up.

Common causes in Brazil

- An extension or IP phone ringing normally on an Asterisk, Issabel, FreePBX, 3CX or Intelbras PBX: expected behavior.
- An 0800 number or IVR that should answer right away but keeps ringing because the IVR server froze and the PBX keeps alerting a destination that does not respond.
- A call-center queue with no agents logged in, set up to ring until someone joins instead of rejecting the call or sending it to a message.
- A wrong time-of-day rule on the PBX that sends after-hours calls to a group of extensions nobody answers.
- An interconnection with the phone network where the carrier passes the destination's alerting (mobile or landline) through as a 180.
- No ringback tone for the caller, because the 180 arrived without SDP and the originating equipment does not generate the tone locally.

How to diagnose it

1. Capture the signaling with `sngrep` (for example, `sngrep port 5060`) and check the sequence: INVITE, 100 Trying, 180 Ringing, and which final response comes after.
2. On Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and repeat the call. Check whether the 180 carries an SDP body.
3. Use `core show channels` during the call to see which extension, queue or application the channel is sitting on while it rings.
4. On 3CX, enable "Verbose logging" and read the call's activity log to see which extension or queue it was delivered to.
5. Measure the time between the 180 and the final response. If the call ends in 487 Request Terminated, the calling side gave up. If it ends in 480, the destination ended it.
6. If the 180 repeats many times, check whether the destination uses reliable provisional responses (`Require: 100rel` and PRACK, RFC 3262) and whether the PRACKs are arriving.

What to send the carrier

- Date and time of each attempt, with the time zone (for example, Brasília time, UTC-3).
- Calling number (A) and called number (B), in the exact format sent in the INVITE.
- The Call-ID of at least one affected call.
- The full sequence of responses: how many 180s, whether a 183 came, and the final response line (for example, `SIP/2.0 480 Temporarily Unavailable`).
- The Reason header, if present, with the Q.850 cause (`Reason: Q.850;cause=19`).
- The pcap trace produced by `sngrep` or `tcpdump`, when available.

Frequently asked questions

Why do I get 180 Ringing but hear no ringback tone?

When the 180 arrives without SDP, your side is expected to generate the ringback tone locally. If the PBX or the phone does not generate it, the call stays silent until it is answered. If the destination sends its own audio, the usual approach is 183 Session Progress with SDP. Check the ringback settings on the trunk and the extension.

Is a call that rings and nobody answers a problem?

It depends on the destination. A personal extension may simply go unanswered. An 0800 (toll-free) number or an IVR, on the other hand, should always answer. If the 180 repeats and the call ends in 480 Temporarily Unavailable or is canceled by the calling side (487), the automated answering has stopped working, or the queue has no free positions.

Does a real test call notice when the 0800 just rings?

Yes. The test call is a real call, so it sees the sequence of SIP responses all the way to the final response, not just the first 180. A 180 on its own does not decide the result. When the test fails, it is repeated before any alert goes out, to tell a real problem apart from a passing blip.

SIP CODES

SIP 183 Session Progress: what it means and how to fix it

By 2N5 team Updated October 7, 2026

Short answer: SIP 183 Session Progress is a provisional response that reports call progress before the call is answered. In practice it almost always carries SDP and opens early media: ringback, music or a recorded announcement. It is not a failure. Listen to what that audio says and check which final response comes after it.

SIP 183 in the reference

Code	183
Official name	Session Progress
Class	1xx: provisional (the call is in progress)
Defined in	RFC 3261

How Monitor treats SIP 183

SIP 183 on its own is not a failure: the call is still in progress. What counts is how it ends. The table shows how a test call that hears a message before it is answered is read for each kind of circuit, depending on what should answer the number.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: “The line never rang and nothing answered, but this circuit is supposed to answer every call by itself.”

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

What does SIP 183 mean in practice?

Under RFC 3261, SIP 183 Session Progress reports call progress when no other provisional response applies. In practice it almost always carries an SDP and opens the audio path before the call is answered, which is known as early media (RFC 3960). That way the caller hears whatever the destination wants to play before the 200 OK: ringback generated by the network, music, a recording notice or a carrier message.

This is where the risk lies for anyone doing monitoring. An 0800 (toll-free) number that plays a message before answering is normal, but the network uses the same technique to announce that the number does not exist, that the subscriber is out of the coverage area or that the route is congested. The call “talks” without having been answered, and only the final response and the content of the audio tell you whether it worked.

Common causes in Brazil

- An 0800 number or IVR that plays a notice before answering, such as “this call may be recorded,” sent as early media.
- The carrier’s network (Vivo, Claro, Algar, Embratel or others) generating ringback as audio instead of letting your side generate it locally.
- A recorded network message announcing a nonexistent number, a phone that is switched off or out of the coverage area, played before a final error response.
- An interconnection gateway to the traditional phone network that translates call progress into a 183 with audio.
- An Asterisk, Issabel or FreePBX PBX using `Progress()` or `Playback(...,noanswer)` in the dial plan to play audio without answering.
- 3CX or Intelbras passing the carrier’s early media through to the extension that placed the call.

How to diagnose it

1. Capture the call with `sngrep port 5060` and open the 183: check whether it carries an `application/sdp` body and which IP and port the audio is going to.
2. On Asterisk, Issabel or FreePBX, use `pjsip set logger on` (or `sip set debug on` with `chan_sip`) to see the full 183 and the final response that follows.
3. Record or listen to the early media. In `sngrep`, the RTP flow screen helps confirm whether audio packets are arriving before the 200 OK.
4. Check whether the 183 requires acknowledgment (`Require: 100rel`). If it does, your side must answer with PRACK (RFC 3262), or the destination may resend the 183 and end the attempt.
5. If the audio is an error message, look for the `Reason: Q.850;cause=...` header (RFC 6432) in the final response to see the exact cause.
6. If there is no audio, check NAT and the firewall: the SDP in the 183 may point to a private IP or to blocked RTP ports.

What to send the carrier

- Date and time of the calls, with the time zone.
- Calling number (A) and called number (B), as they appear in the INVITE.
- The Call-ID of an affected call and the SDP received in the 183.
- What was heard in the early media (the text of the announcement or a recorded excerpt).

- The final response line and the Reason header with the Q.850 cause, if present.
- The pcap of the call, including the RTP, when available.

Frequently asked questions

Why does the 0800 number play a message before it answers?

Because the destination or the carrier sent 183 Session Progress with SDP and started sending audio before the 200 OK. That is how announcements like "this call may be recorded" or hold music before the call is picked up work. Until the 200 OK arrives, the call has not been answered as far as SIP is concerned, even if you can already hear something.

Does hearing audio on the 183 mean the call went through?

No. Early media may be precisely the announcement that the call will not complete, such as "the number you dialed does not exist" or "this phone is switched off." A check that only listens for sound can mark an 0800 (toll-free) number as working when it is actually playing a network error message. You have to look at the final response and at what the audio says.

Does a real test call see what happens after the 183?

Yes. Because it is a real call, it follows the call all the way to the final SIP response instead of stopping at the 183. The 183 on its own does not decide the result. The call also checks audio in both directions and the measured voice quality (MOS-LQO), and a failure is repeated before any alert goes out.

SIP CODES

SIP 400 Bad Request: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 400 Bad Request means the server received the SIP message but could not interpret it because of a syntax error. The most common cause is a malformed header, often altered by a SIP ALG on the router. Read the response's reason phrase, which should point to the problem, and compare the INVITE with what the carrier expects.

SIP 400 in the reference

Code	400
Official name	Bad Request
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	41 (Temporary failure)

How Monitor treats SIP 400

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider refused the call (SIP 400)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 400 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 400 mean in practice?

Under RFC 3261, SIP 400 Bad Request indicates that the request could not be understood because of malformed syntax. The RFC itself recommends that the reason phrase describe the problem in more detail, for example "Missing Call-ID header field." In practice, many servers reply with just "Bad Request," and the diagnosis comes down to comparing the message you sent with what the server expects.

It is a final failure: the call does not happen, and sending the same message again gives the same result. On a trunk that used to work, a 400 that suddenly appears usually comes from a change along the path, such as a new router, a PBX update or a new rule on the carrier's SBC. Seeing the SIP 400 error on every call points to configuration; on only some calls, to specific numbers or headers.

Common causes in Brazil

- SIP ALG enabled on an internet provider's router, rewriting headers and the SDP in an invalid way.
- A Request-URI with characters the carrier does not accept, such as spaces, hyphens or a `+` where it expects digits only.
- A display name (Caller ID) with badly encoded accents or quotes, generated from the extension's settings in Asterisk, Issabel or FreePBX.
- A required header missing or duplicated after a PBX customization or a manipulation rule on the SBC.
- An inconsistent SDP, for example with an empty IP or lines out of the expected order, after a NAT change.
- A 3CX, Intelbras or Asterisk version upgrade that changed the format of a header sent to the trunk.

How to diagnose it

1. Capture with `sngrep port 5060` and read the entire response: the reason phrase and the `Warning` header, if present, usually say what is wrong.
2. On Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and compare the rejected INVITE line by line with an INVITE that works.
3. Take the same capture on the PBX and at the network edge. If the message changes between the two points, a device is altering the SIP. Disable the router's SIP ALG.

- 4. Test with a simple destination number, digits only, and with a display name without accents, to isolate the problem field.
- 5. On 3CX, enable "Verbose logging" and look for the 400 response in the activity log, together with the INVITE that triggered it.
- 6. If the 400 shows up on the REGISTER, before any call, review the username, domain and port configured on the trunk.

What to send the carrier

- Date and time, with the time zone, of a few rejected attempts.
- Calling number (A) and called number (B), in the format sent.
- The Call-ID and the full message that received the 400 (INVITE or REGISTER).
- The full response line, with the reason phrase and the Warning header.
- How many attempts failed and whether all calls are affected or only some.
- The pcap of the capture, when available.

Frequently asked questions

Is SIP 400 my problem or the carrier's?

Almost always, the message that left your side arrived out of spec, either because of the PBX configuration or because a device along the path altered it. But there are cases where the carrier's server is stricter than RFC 3261 and rejects something valid. Compare the message captured as it leaves the PBX with the one that reaches the server.

Which Q.850 cause shows up when the 400 crosses into the phone network?

Under RFC 3398, a gateway that converts SIP to ISUP translates 400 Bad Request into Q.850 cause 41 (Temporary failure). So on the traditional phone network side, a 400 can look like a temporary failure even though it started as a malformed SIP message. No Q.850 cause is converted into a 400 in the other direction.

SIP CODES

SIP 401 Unauthorized: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 401 Unauthorized is the authentication challenge from the registrar or server: it asks for the request to be resent with digest credentials. Once is normal on REGISTER and INVITE. If the 401 comes back after the resend, the username, password or realm is wrong. Check the trunk credentials with the carrier.

SIP 401 in the reference

Code	401
Official name	Unauthorized
Class	4xx: request failure
Defined in	RFC 3261

Q.850 causes that usually come with it	21 (Call rejected)
--	--------------------

How Monitor treats SIP 401

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Monitoring blind
A test tone or echo service	Monitoring blind
Other: it may ring or be answered	Monitoring blind

Monitoring blind means the failure is on 2N5's side, not in your circuit: the test does not count against the circuit, stays out of the SLA and does not alert the customer. 2N5 is told if it lasts.

What Monitor records for an IVR: "The provider did not accept 2N5's calling account, usually a wrong username or password, so no test call reached the line. This is a configuration problem on 2N5's side, not an outage on this line (SIP 401)."

If the carrier returns SIP 401 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 401 mean in practice?

Under RFC 3261, SIP 401 Unauthorized indicates that the request requires user authentication. It is sent by registrars and by the agents receiving the call, and it carries the `WWW-Authenticate` header with the digest challenge: realm, nonce and algorithm. The client computes the response from the username and password and resends the request with the `Authorization` header.

That is why a single 401 is routine: many trunks challenge every REGISTER and every INVITE. The warning sign is the cycle REGISTER, 401, REGISTER with credentials and another 401. If the repeated challenge carries `stale=true`, the nonce just expired and the client should try again with the new nonce. Without `stale`, the credentials were rejected, and that repeated SIP 401 error is what keeps the trunk from registering. For proxies, the equivalent is [SIP 407](#).

Common causes in Brazil

- The trunk password was changed in the carrier's portal and not updated in Asterisk, Issabel, FreePBX or 3CX.
- An authentication username that differs from the trunk number: many carriers use an authentication login separate from the user that goes in the From header.
- A wrong realm or domain, so the PBX answers the challenge with another server's details.
- A trunk contracted for IP-based authentication, but the PBX started going out through a different public IP, so the carrier asks for credentials again.
- Two PBXs registering with the same credentials, each knocking out the other's registration.
- Special characters in the password that the PBX's admin panel saves or escapes differently.

How to diagnose it

1. On Asterisk or FreePBX with PJSIP, run `pjsip show registrations` and check whether the trunk is `Registered` or `Rejected`. With `chan_sip` (common on Issabel), use `sip show registry`.
2. Turn on `pjsip set logger on` (or `sip set debug on`) and confirm whether a second 401 arrives after the attempt carrying `Authorization`.
3. Compare the realm in the `WWW-Authenticate` header with the one configured on the trunk, and check the user sent in the `username` field of `Authorization`.
4. On FreePBX, review the authentication fields of the PJSIP trunk (username, password and server). On 3CX, check the "Authentication ID" and password, and enable "Verbose logging."
5. Use `sngrep port 5060` to see the full ladder and make sure the resend goes out from the same IP and port as the first attempt.
6. If the trunk is IP-based, confirm with `curl ifconfig.me` which public IP the server is using.

What to send the carrier

- Date and time, with the time zone, of the registration or call attempts.
- The trunk number or identifier and the authentication username used.
- The Call-ID of a rejected REGISTER or INVITE.
- Both 401 responses, with the full `WWW-Authenticate` header showing the realm.
- The originating public IP and how many attempts failed.
- The pcap of the capture, when available, without exposing the password in clear text.

Frequently asked questions

Does seeing 401 in the Asterisk log mean the trunk has a problem?

Not necessarily. The normal authentication flow is: REGISTER or INVITE without credentials, 401 with the `WWW-Authenticate` header, a resend with the `Authorization` header and then 200 OK. The problem is when the second attempt, already carrying credentials, also gets a 401. That means the credentials were rejected, not just challenged.

What is the difference between 401 and 403 on a SIP trunk?

The 401 says "I need to know who you are" and offers a challenge to authenticate. The 403 Forbidden says "I know what you asked for and I will not do it," and authenticating again does not help. If you see a 401 followed by a 403, authentication passed, but the carrier rejected the call for another reason, such as the calling number.

SIP CODES

SIP 403 Forbidden: what it means and how to fix it

By 2N5 team Updated October 7, 2026

Short answer: SIP 403 Forbidden means the server understood the request and refused to carry it out, and authenticating again does not help. On trunks in Brazil, the most common cause is the carrier not recognizing the calling number sent. Check the number in the From and P-Asserted-Identity headers and confirm with the carrier which numbers the trunk may use.

SIP 403 in the reference

Code	403
Official name	Forbidden
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	21 (Call rejected), 55 (Incoming calls barred within CUG), 57 (Bearer capability not authorized), 87 (User not member of CUG)

How Monitor treats SIP 403

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider refused this call. The usual cause is that the provider does not recognize the number the call comes from; it can also mean this number is barred or suspended, or the account is out of credit (SIP 403)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 403 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 403 mean in practice?

Under RFC 3261, SIP 403 Forbidden indicates that the server understood the request but refuses to fulfill it. Authorization will not help, and the request should not be repeated. It differs from [SIP 401](#): with a 401 the server does not yet know who you are; with a 403 it knows and has decided to say no.

On SIP trunks in Brazil, the 403 is one of the most frequent codes when "the trunk registers but cannot call out." The carrier accepted the registration or the IP, but a business rule blocked that particular call: the calling number, the destination, the account status or the address the call came from. The reason phrase and the `Warning` header sometimes say which rule was applied, and it is worth reading both before opening a ticket about the SIP 403 error.

Common causes in Brazil

- A calling number (caller ID) that does not belong to the trunk, sent in the From or P-Asserted-Identity header by an extension or a PBX rule.
- A calling number in a format other than the one the carrier expects, for example with or without the area code (DDD), with a leading 0 or in E.164 (+55...).
- A blocked or suspended line or trunk, for example for non-payment or an administrative block.
- A prepaid or allowance-based plan with no credit left for that type of call.
- A destination blocked under the contract: mobile, long distance, international or premium-rate numbers.
- An unauthorized source IP on an IP-based trunk, after switching internet links or migrating Asterisk, Issabel, FreePBX or 3CX.

How to diagnose it

1. Capture with `sngrep port 5060` and note the 403's reason phrase, the `Warning` header and `Reason: Q.850;cause=21`, if present (`Q.850 cause 21, call rejected`).
2. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and check the From, P-Asserted-Identity and Remote-Party-ID headers in the INVITE.
3. Call the same destination using the trunk's main number as the calling number. If that works, the problem is the calling number of the extension or route.
4. Call different kinds of destinations (local landline, mobile, long distance) to see whether the block is by destination type.
5. On 3CX, enable "Verbose logging" and review the outbound caller ID configured on the trunk and in the outbound rule.
6. Confirm in the carrier's portal or with its billing team that the account is active and has credit.

What to send the carrier

- Date and time, with the time zone, of a few rejected attempts.
- Calling number (A) and called number (B), exactly as they went out in the INVITE.
- The Call-ID of at least one rejected call.
- The full response line, with the reason phrase, the `Warning` header and the `Reason` header with the Q.850 cause.
- The public IP the calls go out from and how many attempts failed.
- The pcap of the capture, when available.

Frequently asked questions

Why do I get 403 on only some calls?

When the 403 shows up only for certain destinations, the carrier is usually blocking that type of call under your contract, such as mobile, long distance, international or premium-rate numbers. When it shows up only for certain extensions, the problem is usually the calling number those extensions send, which does not belong to the trunk.

Which Q.850 cause comes with a 403?

Under RFC 3398, Q.850 cause 21 (Call rejected) is the most common, and a 403 is translated back as cause 21. Causes 55 and 87, which relate to closed user groups, and cause 57, bearer capability not authorized, are also converted into 403. In practice, cause 21 is the one you will see most.

SIP CODES

SIP 404 Not Found: what it means and how to fix it

By 2N5 team Updated October 7, 2026

Short answer: SIP 404 Not Found means the server has definitive information that the user or number does not exist in the called domain. The most common cause on trunks is the number format: a missing area code (DDD), a wrong prefix, or E.164 where the carrier expects another format. Compare the number sent in the INVITE with the format the carrier accepts.

SIP 404 in the reference

Code	404
Official name	Not Found
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	1 (Unallocated (unassigned) number), 2 (No route to specified transit network (national use)), 3 (No route to destination), 26 (Non-selected user clearing)

How Monitor treats SIP 404

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The number does not exist or was dialed wrongly, so the provider did not connect the call (SIP 404)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 404 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 404 mean in practice?

Under RFC 3261, SIP 404 Not Found indicates that the server has definitive information that the user does not exist in the domain given in the Request-URI. The same code is used when the domain in the Request-URI does not match any domain served by whoever received the request.

On a trunk, the SIP 404 error almost always boils down to “that number does not exist here.” Sometimes the number really is not allocated, and the network usually responds with [Q.850 cause 1](#). Often, though, the number exists and was sent in the wrong format, or it reached the destination PBX, which has no route for it. When a call comes into your own PBX and gets a 404, the DID is not configured in the inbound routes. If the number existed and was deactivated, the more precise code would be [SIP 410](#); if digits are missing, [SIP 484](#).

Common causes in Brazil

- A number sent in a format the carrier does not accept: without the area code (DDD), with an extra `0`, with a carrier selection code (CSP) where it does not belong, or in E.164 (`+55`) where the carrier expects digits only.
- A mobile number dialed without the ninth digit, or a landline dialed with one digit too many.
- An 0800 number sent incomplete or with an area code in front.
- A destination number that is not allocated or has been deactivated.
- A ported number routed to the old carrier (RFC 3398 maps cause 26 to 404 in this case).
- An inbound DID with no route in Asterisk, Issabel, FreePBX or 3CX: the PBX answers 404 because it cannot find the extension in the context.

How to diagnose it

1. Capture with `sngrep port 5060` and look at the INVITE's Request-URI: that is where the number appears exactly as it went out.
2. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`). For inbound calls, the log shows the extension that was not found in the context. Check it with `dialplan show <number>@<context>`.
3. Review the outbound rules (dial patterns) in FreePBX or 3CX: a prefix stripped or added by mistake changes the number sent.
4. Test the same destination in different formats, for example `11987654321` and `5511987654321`, to find out what the carrier accepts.
5. Look for the `Reason: Q.850;cause=1` header (RFC 6432) or another cause in the response: it tells an unallocated number apart from a missing route (cause 3).
6. Call the same number from another line. If that does not go through either, the number really does not exist.

What to send the carrier

- Date and time of the attempts, with the time zone.
- Calling number (A) and called number (B), exactly as they appeared in the Request-URI.
- The Call-ID of at least one call.
- The full response line and the Reason header with the Q.850 cause, if present.
- The dialing format you are using and how many attempts failed.
- The pcap of the capture, when available.

Frequently asked questions

What format should I send the number to the carrier in?

It depends on the carrier and the contract. Some accept the area code (DDD) plus the number (10 or 11 digits), others want E.164 with or without the +55, and others accept traditional dialing, with `0`, the carrier selection code and the area code. An 0800 (toll-free) number has 11 digits, as in 0800 123 4567. Ask for the official format in writing and adjust the PBX's outbound rule.

What is the difference between 404, 410 and 484?

A 404 says the number does not exist at that destination. A 410 Gone says it used to exist and is permanently out of service. A 484 Address Incomplete says digits are missing. In practice, many carriers use 404 for all three cases, so it is worth checking the Q.850 cause that comes with it too.

SIP CODES

SIP 407 Proxy Authentication Required: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 407 Proxy Authentication Required is the authentication challenge from a proxy, usually the carrier's, asking for the INVITE to be resent with digest credentials. Getting a 407 on the first attempt is normal. If the 407 repeats after the resend, the credentials or the realm are wrong. Review the trunk authentication on the PBX.

SIP 407 in the reference

Code	407
Official name	Proxy Authentication Required
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	21 (Call rejected)

How Monitor treats SIP 407

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Monitoring blind
A test tone or echo service	Monitoring blind
Other: it may ring or be answered	Monitoring blind

Monitoring blind means the failure is on 2N5's side, not in your circuit: the test does not count against the circuit, stays out of the SLA and does not alert the customer. 2N5 is told if it lasts.

What Monitor records for an IVR: "The provider did not accept 2N5's calling account, usually a wrong username or password, so no test call reached the line. This is a configuration problem on 2N5's side, not an outage on this line (SIP 407)."

If the carrier returns SIP 407 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 407 mean in practice?

Under RFC 3261, SIP 407 Proxy Authentication Required is similar to [SIP 401](#), but it indicates that the client must first authenticate with the proxy. The proxy sends the `Proxy-Authenticate` header with the digest challenge (realm and nonce), and the client resends the request with the `Proxy-Authorization` header.

On SIP trunks, the normal flow of a call is INVITE, 407, ACK, a new INVITE with credentials and then 100 Trying and 180 or 183. That is why the 407 shows up in almost every outbound call log and does not signal a problem on its own. The sign of failure, the real SIP 407 error, is a second 407 right after the INVITE that already carried credentials, or a PBX that never resends the INVITE. If the new challenge carries `stale=true`, the nonce has just expired and simply trying again is enough.

Common causes in Brazil

- An outdated trunk password or authentication username in Asterisk, Issabel, FreePBX or 3CX.
- A trunk on the PBX with no outbound authentication configured, so the INVITE is never resent with credentials.
- A different realm for the REGISTER and the INVITE, with the PBX using the wrong credential for the proxy's challenge.
- A user in the INVITE's From header that differs from the authentication login the carrier tied to the trunk.
- An IP-based trunk that started going out through a different public IP, so the carrier's proxy demands credentials.
- An SBC along the path that strips or rewrites the `Proxy-Authorization` header.

How to diagnose it

1. Capture with `sngrep port 5060` and check whether the INVITE is resent after the 407 and whether the second attempt also gets a 407.
2. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and compare the realm in `Proxy-Authenticate` with the one in `Proxy-Authorization`.
3. Run `pjsip show endpoint <trunk>` and check that the endpoint has `outbound_auth` pointing to the right authentication. With `chan_sip`, review `secret`, `defaultuser` and `fromuser`.
4. On 3CX, enable "Verbose logging" and review the trunk's "Authentication ID" and password.
5. Place a test call right after restarting the registration with `pjsip send register <trunk>`, to rule out an expired nonce.
6. If there is an SBC or a firewall doing SIP inspection, capture before and after it to see whether the authentication header arrives intact.

What to send the carrier

- Date and time, with the time zone, of the rejected calls.
- Calling number (A), called number (B) and the trunk's authentication username.
- The Call-ID of an affected call.

- Both 407 responses received, with the full `Proxy-Authenticate` header (realm and nonce).
- The public IP the calls go out from and how many attempts failed.
- The pcap of the capture, when available, without exposing the password.

Frequently asked questions

What is the difference between 401 and 407?

Both are digest authentication challenges and work the same way. The 401 comes from a registrar or the final destination and uses the WWW-Authenticate and Authorization headers. The 407 comes from a proxy along the path and uses Proxy-Authenticate and Proxy-Authorization. On trunks, it is common for the REGISTER to get a 401 and the INVITE to get a 407.

The trunk registers fine, but calls get 407. Why?

Registration and calls can use different credentials or realms, and some PBXs only answer the INVITE challenge if the trunk has outbound authentication configured. It also happens when the INVITE goes out with a user in the From header other than the one the carrier expects. Compare the INVITE's Proxy-Authorization with the REGISTER's Authorization.

SIP CODES

SIP 408 Request Timeout: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 408 Request Timeout means the request got no response in time. Either the server could not locate the destination in time, or the originating device itself gave up after waiting. The most common cause is the network: firewall, NAT or the route to the server. Capture the traffic and check whether the INVITE goes out and whether any response comes back.

SIP 408 in the reference

Code	408
Official name	Request Timeout
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	18 (No user responding), 102 (Recovery on timer expiry)

How Monitor treats SIP 408

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider did not respond to the call in time (SIP 408)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 408 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 408 mean in practice?

Under RFC 3261, SIP 408 Request Timeout indicates that the server could not produce a response within a suitable amount of time, for example because it could not locate the user in time. The client may repeat the request later without changes.

In practice, the SIP 408 error has two origins. It can come from a server along the path, such as the carrier's proxy or the destination PBX, which forwarded the call and got no response. Or it can be generated locally by your own PBX, when the INVITE got no response at all before timer B expired. On calls coming from the traditional phone network, it usually corresponds to [Q.850 cause 18](#); in the other direction, to [Q.850 cause 102](#). In every case, something stopped responding.

Common causes in Brazil

- A firewall or NAT rule blocking the SIP responses coming back from the carrier, after a router or link change.
- A destination extension that shows as registered in Asterisk, Issabel, FreePBX or 3CX, but whose NAT mapping has expired so it no longer receives the INVITE.
- An INVITE too large for UDP, with many codecs or headers, fragmented along the path and dropped by some device.
- The carrier's server unavailable, or an unstable internet route to it.
- The trunk's DNS pointing to a decommissioned server.
- An IP phone switched off or off the network, with its registration still valid on the PBX.

How to diagnose it

1. Capture with `sngrep port 5060` and confirm whether the INVITE goes out, how many retransmissions happen and whether any 100 Trying comes back.
2. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`). If the 408 does not appear on the wire, it was generated locally.
3. For extensions, run `pjsip show contacts` (or `sip show peers`) and check whether the contact is `Avail` or `Unavail`, along with the qualify response time.
4. Check the size of the INVITE in the capture. If it goes over about 1,300 bytes, trim the codec list or test the trunk over TCP.
5. Also run `tcpdump -i any -w trace.pcap port 5060` on the router or firewall, if possible, to see where the message gets lost.
6. On 3CX, enable "Verbose logging" and check in the activity log whether the trunk or the extension stopped responding.

What to send the carrier

- Date and time, with the time zone, of the unanswered calls.
- Calling number (A) and called number (B).
- The Call-ID of an affected call and the time of each INVITE retransmission.
- The full 408 line, if it came from the network, and the Reason header with the Q.850 cause, if present.
- The destination IP and port used and how many attempts failed.
- The pcap captured at the edge of your network.

Frequently asked questions

Does the 408 always come from the carrier?

No. Under RFC 3261, when an INVITE gets no response at all before timer B expires (32 seconds with the default values), the originating device itself treats the failure as a 408. Often no 408 ever crossed the network. If the capture shows the INVITE going out and nothing coming back, the problem is along the path or at the destination server.

Which Q.850 cause comes with a 408?

Under RFC 3398, Q.850 cause 18 (No user responding) is converted into 408 when the call leaves the phone network for SIP. In the other direction, a 408 entering the phone network becomes Q.850 cause 102 (Recovery on timer expiry). Both describe a timeout with no response from the destination.

SIP CODES

SIP 410 Gone: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 410 Gone means the destination is no longer available on that server and no forwarding address is known, in a condition considered permanent. In telephony, it is almost always a number that was deactivated, canceled or changed. Confirm with the number's owner whether it is still active and update the record or route to the new number.

SIP 410 in the reference

Code	410
Official name	Gone
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	22 (Number changed), 23 (Redirection to new destination)

How Monitor treats SIP 410

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider refused the call: the number is no longer in service (SIP 410)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 410 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 410 mean in practice?

Under RFC 3261, SIP 410 Gone indicates that the requested resource is no longer available at the server and no forwarding address is known. The condition is expected to be permanent. If the server does not know whether the condition is permanent, the RFC says to use [SIP 404](#) instead.

On a phone trunk, the SIP 410 error usually means "this number used to exist and no longer does." It is more specific than a 404 and therefore less common: many networks answer 404 or play a recorded announcement even when the number was deactivated. When a 410 does appear, treat it as a strong sign of a change in the number's registration, not as a passing network failure. Calling again will not change the result.

Common causes in Brazil

- A landline or mobile line canceled by the subscriber and not yet reassigned to another customer.
- An 0800 or service number deactivated during a contract or carrier migration.
- A number the company replaced, with the old one still saved in phone books, IVRs or PBX rules.
- An inbound DID removed from the trunk but still advertised or still present in other systems' forwarding.
- An extension deleted on an Asterisk, Issabel, FreePBX, 3CX or Intelbras PBX configured to answer 410 for removed users.

How to diagnose it

1. Capture with `sngrep port 5060` and check that the response really is `410 Gone` and whether there is a `Reason: Q.850;cause=22` header (RFC 6432).
2. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and see which server the 410 came from: the carrier or a PBX further along.
3. Call the same number from another line, ideally on another carrier, to confirm the result repeats.
4. Check that the number is not mistyped on the PBX or in the outbound rule: one swapped digit can lead to an old, deactivated number.
5. On 3CX, enable "Verbose logging" and check the exact number sent to the trunk.
6. If the number is yours, ask the carrier that holds it about the line's status and the date it changed.

What to send the carrier

- Date and time of the attempts, with the time zone.
- Calling number (A) and called number (B), as sent in the INVITE.
- The Call-ID of at least one call.
- The full 410 response line and the Reason header with the Q.850 cause, if present.
- Since when the number has stopped going through and how many attempts failed.
- The pcap of the capture, when available.

Frequently asked questions

Can a number that gets 410 start working again?

It can, but not on its own. The 410 signals a condition the server considers permanent, such as the line being canceled. It only goes through again if the number is reactivated or reassigned. If your company's 0800 (toll-free) number or main number gets a 410, talk to the carrier that holds the number: it may have been canceled by mistake or during a migration.

Which Q.850 cause comes with a 410?

Under RFC 3398, Q.850 cause 22 (Number changed) with no new number given is converted into 410, as is cause 23 (Redirection to new destination). When the network does provide the new number, cause 22 becomes 301 Moved Permanently, with the new number in the Contact header. In the other direction, a 410 becomes cause 22.

SIP CODES

SIP 415 Unsupported Media Type: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 415 Unsupported Media Type means the server rejected the request because the message body is in a format it does not accept for that method. It is not a codec mismatch, which produces a 488. The most common cause is an unexpected Content-Type. Read the response's Accept header, which lists the accepted formats.

SIP 415 in the reference

Code	415
Official name	Unsupported Media Type
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	79 (Service or option not implemented, unspecified)

How Monitor treats SIP 415

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider refused the call (SIP 415)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 415 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 415 mean in practice?

Under RFC 3261, SIP 415 Unsupported Media Type indicates that the server refuses to service the request because the message body is in a format it does not support for the requested method. The server should return the list of accepted formats in the `Accept`, `Accept-Encoding` or `Accept-Language` headers, depending on the problem.

On trunks, the SIP 415 error is rare on an ordinary call, because almost every INVITE carries only `application/sdp`. It shows up when some device adds other content, such as a multipart body with signaling data from the phone network or a compressed body, or when in-call messages, such as INFO for DTMF, use a type the other side does not accept. For codec incompatibility, the correct code is [SIP 488](#).

Common causes in Brazil

- A PBX or SBC sending an INVITE with a `multipart/mixed` body (SDP plus phone-network signaling data) to a carrier that only accepts SDP.
- DTMF sent by SIP INFO with `application/dtmf-relay` or `application/dtmf`, rejected by the other side while navigating an IVR.
- A compressed body (`Content-Encoding: gzip`) that the destination server cannot open.
- A `Content-Type` header missing or misspelled after a customization in Asterisk, Issabel, FreePBX or 3CX.
- An interconnection gateway configured for SIP-I or SIP-T talking to an ordinary SIP trunk.

How to diagnose it

1. Capture with `sngrep port 5060` and check which method got the 415 (INVITE, INFO, MESSAGE) and the `Content-Type` that was sent.
2. Read the `Accept`, `Accept-Encoding` and `Accept-Language` headers in the response to learn what the server accepts.
3. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and check the trunk's DTMF mode. In PJSIP, prefer `dtmf_mode=rfc4733`.
4. If the INVITE carries a multipart body, review the SBC or gateway that adds it and turn off the inclusion of phone-network signaling data for that trunk.
5. On 3CX, enable "Verbose logging" and check the trunk's DTMF mode and advanced parameters.

6. Place a test call to an IVR and press options to confirm the digits arrive without a 415.

What to send the carrier

- Date and time of the occurrences, with the time zone.
- Calling number (A) and called number (B).
- The Call-ID of the call and the method that got the 415.
- The full response line and the `Accept` headers received.
- The `Content-Type` sent and how many attempts failed.
- The pcap of the capture, when available.

Frequently asked questions

Is a 415 the same thing as an incompatible codec?

No. Codec incompatibility is handled inside the SDP and usually produces SIP 488 Not Acceptable Here. The 415 is about the type of the whole body, given in the Content-Type header, such as application/sdp, multipart/mixed or application/dtmf-relay. If the server does not accept that type for the method used, it answers 415 without ever getting to codec negotiation.

Why do I get a 415 mid-call when pressing keys in an IVR?

When the PBX sends DTMF digits as SIP INFO messages, the body goes out with a type such as application/dtmf-relay. If the other side does not accept that type, it answers the INFO with a 415 and the digit is lost, although the call continues. Switch DTMF to RFC 4733 (events in the RTP), which is the standard carriers accept most widely.

SIP CODES

SIP 480 Temporarily Unavailable: what it means and how to fix it

By 2N5 team Updated October 7, 2026

Short answer: SIP 480 Temporarily Unavailable means the destination was located but is not available to take the call right now: it is not registered, it is on Do Not Disturb or it did not answer. The most common cause on a PBX is an extension or IP phone that is not registered. Check the destination's registration and look for the Retry-After header, which may suggest when to try again.

SIP 480 in the reference

Code	480
Official name	Temporarily Unavailable
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	19 (No answer from user (user alerted)), 20 (Subscriber absent), 31 (Normal, unspecified), 18 (No user responding)

How Monitor treats SIP 480

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	After ringing	Without ringing
A greeting or recorded message (the IVR)	Down	Down
A test tone or echo service	Down	Down
Other: it may ring or be answered	Up	Down

For "Other": With "Count ringback as reachable" unticked: Down.

What Monitor records for an IVR: "The line rang and nobody answered, but this circuit is supposed to answer every call by itself (SIP 480)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 480 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 480 mean in practice?

Under RFC 3261, SIP 480 Temporarily Unavailable indicates that the destination's system was contacted, but the user is not available at the moment: not logged in, in a state that prevents communication, or with Do Not Disturb turned on. The response may carry the `Retry-After` header with a better time to call. Proxies and redirect servers also use 480 when they recognize the user but have no valid address to deliver the call to.

In telephony, the 480 also carries network situations: **Q.850 cause 20**, such as a mobile phone switched off or out of coverage, cause 19, when the destination rang and nobody answered, and **Q.850 cause 31**. It is a "temporary" code: the same number may go through minutes later. A SIP 480 error that keeps repeating on the same destination, however, deserves a closer look.

Common causes in Brazil

- An extension or IP phone not registered on Asterisk, Issabel, FreePBX, 3CX or Intelbras, because of a network outage, a restart or a changed password.
- An extension with Do Not Disturb turned on, set to reject the call instead of sending it to voicemail.
- A destination mobile phone that is switched off, has no signal or is out of the coverage area.
- A call that rang until the PBX's or the network's time limit with no answer.
- An IVR or auto-attendant server that lost its registration with the PBX, leaving the 0800 number with no destination.
- A queue with no available agents, set to end the call instead of keeping it on hold.

How to diagnose it

1. Capture with `sngrep port 5060` and see who generated the 480: the carrier or your PBX. Check the `Reason: Q.850;cause=...` and `Retry-After` headers.
2. On Asterisk or FreePBX, run `pjsip show contacts` (or `sip show peers` with `chan_sip`) and confirm whether the destination extension is registered and responding.
3. Turn on `pjsip set logger on` (or `sip set debug on`) and repeat the call to see whether it rang (180 or 183) before the 480.
4. Run `core show channels` during the call to see which queue, group or extension it stopped at.
5. Review the PBX's Do Not Disturb, time-of-day and queue rules for the affected destination.
6. On 3CX, enable "Verbose logging" and check the status of the extension or queue at the time of the call.

What to send the carrier

- Date and time of the calls, with the time zone.
- Calling number (A) and called number (B).
- The Call-ID of at least one call.
- The full 480 line, with the Reason (Q.850 cause) and Retry-After headers, if present.
- Whether the call rang and how many attempts failed.
- The pcap of the capture, when available.

Frequently asked questions

Which Q.850 cause comes with a 480?

Under RFC 3398, three causes are converted into 480: 19 (No answer from user, the destination rang and was not answered), 20 (Subscriber absent, such as a mobile phone switched off or out of coverage) and 31 (Normal, unspecified). In the other direction, a 480 entering the phone network becomes Q.850 cause 18 (No user responding).

Does a 480 mean my 0800 number is down?

For an 0800 (toll-free) number or an IVR, which should always answer, a 480 is a sign of trouble: the final destination is not available, either because the IVR server dropped its registration or because the queue had nobody to answer. For a personal extension, it may just be someone on Do Not Disturb or away from their desk.

SIP CODES

SIP 481 Call/Transaction Does Not Exist: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 481 Call/Transaction Does Not Exist means the server received a request that does not match any dialog or transaction it knows. The most common cause is one side having already ended the call while the other keeps sending messages. If calls always drop at the same minute, review the session timers and NAT.

SIP 481 in the reference

Code	481
Official name	Call/Transaction Does Not Exist
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	41 (Temporary failure)

How Monitor treats SIP 481

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider refused the call (SIP 481)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 481 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 481 mean in practice?

Under RFC 3261, SIP 481 Call/Transaction Does Not Exist indicates that the server received a request that does not match any existing dialog or transaction. The dialog is identified by the Call-ID and by the From and To tags. If any of these values changes along the path, or if one side has already forgotten the call, the next message gets a 481.

The RFC also says that when a request inside the dialog gets a 481, the client must end the dialog. That is why the SIP 481 error rarely keeps a call from starting. The typical effect is a call that connects, carries audio normally and drops on its own after a few minutes, or a CANCEL that arrives too late. A gateway that carries the 481 into the phone network translates it as Q.850 cause 41, temporary failure.

Common causes in Brazil

- A session timer refresh (re-INVITE or UPDATE) reaching an SBC or PBX that has already discarded the call, dropping it at the same minute every time.
- A NAT mapping on the router that expires during the call, so later messages arrive at a different destination.
- The router's SIP ALG rewriting tags or the Call-ID.
- A restart or node switch on the carrier's SBC or on the PBX (Asterisk, Issabel, FreePBX, 3CX) in the middle of active calls.
- A CANCEL sent after the call had already been answered or ended on the other side.

- A BYE sent by both sides at almost the same moment.

How to diagnose it

1. Capture with `sngrep port 5060` and find the message that got the 481: BYE, CANCEL, re-INVITE, UPDATE or INFO.
2. Compare that message's Call-ID and From and To tags with those of the original INVITE. If they changed, some device along the path is rewriting the SIP.
3. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and check the negotiated `Session-Expires` and `Min-SE` headers.
4. Measure how long the dropped calls last. If it is always close to half of `Session-Expires`, the problem lies in the session refresh.
5. Review the UDP timeout on the firewall and disable the router's SIP ALG.
6. On 3CX, enable "Verbose logging" and check whether the trunk received the 481 from the carrier or whether 3CX sent it.

What to send the carrier

- Date and time, with the time zone, of when the calls started and when they dropped.
- Calling number (A) and called number (B).
- The Call-ID of the call and the message that got the 481.
- The full response line and the negotiated `Session-Expires` value.
- How many calls dropped and whether the duration repeats.
- The pcap of the whole call, from the INVITE to the drop.

Frequently asked questions

Why do my calls always drop after the same amount of time?

A common pattern is the session timer (RFC 4028). Each side refreshes the session with a re-INVITE or UPDATE at the agreed interval. If that refresh gets a 481 or no response, RFC 3261 says the dialog must be ended, and the call drops. When this always happens at the same minute, check the negotiated `Session-Expires`.

Is getting a 481 on a BYE serious?

Usually not. If the other side had already ended the call and your BYE arrives afterward, the 481 just confirms there is nothing left to end. The problem is a 481 on messages inside a call that should still be active, such as re-INVITE, UPDATE or INFO, because that drops the call midway.

SIP CODES

SIP 484 Address Incomplete: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 484 Address Incomplete means the number sent in the Request-URI is incomplete and the destination cannot complete the call without more digits. The most common cause is a PBX dialing rule that cuts digits, such as the area code (DDD), the ninth digit of a mobile number or part of an 0800 number. Compare the number sent with the full expected format.

SIP 484 in the reference

Code	484
Official name	Address Incomplete
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	28 (Invalid number format (address incomplete))

How Monitor treats SIP 484

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider refused the call: the number dialed is incomplete (SIP 484)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 484 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 484 mean in practice?

Under RFC 3261, SIP 484 Address Incomplete indicates that the server received a request with an incomplete Request-URI. The RFC recommends that the reason phrase carry more information. The code is also used in overlap dialing (RFC 3578), where the number is sent a few digits at a time and the 484 asks for more digits.

On SIP trunks in Brazil, overlap dialing is rare, and a 484 almost always means the PBX sent fewer digits than it should have. Formats vary: a local call with 8 digits (landline) or 9 (mobile), long distance with 0 + carrier selection code (CSP) + area code (DDD) + number, 0800 numbers with 11 digits, or E.164 with +55. A rule that strips the wrong prefix produces the SIP 484 error. If the number is complete but does not exist, the expected code is SIP 404.

Common causes in Brazil

- An outbound rule in FreePBX, Issabel or 3CX that strips too many digits, for example removing the area code along with the 0.
- A mobile number dialed without the ninth digit, still present in old phone books and records.
- An 0800 number saved with one digit missing in the IVR, the CRM or the PBX.
- An IP phone with a dial plan that sends the number too early, before the user finishes keying it in.
- A long-distance number sent without an area code to a carrier that requires the area code on every call.
- A short service number that the carrier does not recognize in that form.

How to diagnose it

1. Capture with `sngrep port 5060` and count the digits in the Request-URI of the INVITE that got the 484.
2. On Asterisk or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and compare the number dialed on the extension with the number that went out to the trunk.
3. Review the outbound routes (dial patterns, prefixes to strip and to add) in FreePBX or Issabel. On 3CX, review the outbound rules and enable "Verbose logging."
4. On Asterisk, use `dialplan show <number>@<context>` to see which rule matched the number.
5. Test the same destination in the full format, for example area code plus number, and in E.164, to find out what the carrier accepts.
6. Check for the `Reason: Q.850;cause=28` header (RFC 6432) in the response, if the carrier sends it.

What to send the carrier

- Date and time of the attempts, with the time zone.
- Calling number (A) and called number (B), exactly as they appeared in the Request-URI.
- The Call-ID of at least one call.
- The full 484 line and the Reason header with the Q.850 cause, if present.
- The dialing format you are using and how many attempts failed.
- The pcap of the capture, when available.

Frequently asked questions

Why does the PBX send the number before I finish dialing?

Some phones and PBXs send the number as soon as it matches a short dialing rule. If the rule accepts, say, eight digits, a nine-digit mobile number is sent incomplete and comes back as a 484. Adjust the phone's dial plan and the PBX's outbound rules to the current formats.

Which Q.850 cause comes with a 484?

Under RFC 3398, Q.850 cause 28 (Invalid number format, address incomplete) is converted into 484, and a 484 is translated back as cause 28 when the call crosses into the phone network. So in an interconnection log, the same problem can appear as a 484 on one side and cause 28 on the other.

SIP CODES

SIP 486 Busy Here: what it means and how to fix it

By 2N5 team Updated October 20, 2026

Short answer: SIP 486 Busy Here means the destination was reached but cannot accept another call at that point. On an extension or mobile phone, the person is already on a call. On a SIP trunk, the most common cause is the concurrent-channel limit being used up. Check how many calls were active at the time and compare that with the contracted limit.

SIP 486 in the reference

Code	486
Official name	Busy Here
Class	4xx: request failure
Defined in	RFC 3261
Q.850 causes that usually come with it	17 (User busy)

How Monitor treats SIP 486

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Up

For "Other": With "Count busy as reachable" unticked: Down.

What Monitor records for an IVR: "The line was busy: the phone is on another call, but this circuit is supposed to answer every call by itself (SIP 486)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 486 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 486 mean in practice?

SIP 486 Busy Here is a final 4xx-class response: the destination's system was contacted, but it will not or cannot accept another call at that point. RFC 3261 lets the response carry the Retry-After header, suggesting when to try again. When converting to the phone network, RFC 3398 maps the 486 to **Q.850 cause 17 (User busy)**, and the reverse mapping also leads to 486.

On an ordinary line, “busy” means someone is already talking. On a SIP trunk, the meaning is usually different: all the concurrent channels contracted or configured are in use, and the next call is refused as busy. That is why a SIP 486 error at peak hours deserves attention even when nobody is on the phone. Unlike [SIP 600](#), the 486 applies only to that destination; another point may still answer.

Common causes in Brazil

- Concurrent channels used up on the SIP trunk contracted with the carrier (Vivo, Claro, Algar, Embratel or another), usually at peak hours or during an outbound calling campaign.
- A channel limit configured on the PBX itself (Asterisk, Issabel, FreePBX, 3CX) lower than the contracted one, refusing calls before the trunk is full.
- A destination extension or mobile phone already on a call, with no call waiting enabled.
- An IVR’s queue or answering group with no free agents, set to return busy instead of keeping the call in the queue.
- “Do Not Disturb” turned on for an Intelbras extension, IP phone or softphone, which many devices signal as 486.
- Stuck calls: channels that were not released after a failed BYE still count as in use.

How to diagnose it

1. Capture the signaling with `sngrep` on the server and see who generated the 486: the PBX, the SBC or the carrier. The response’s Via, Server and User-Agent headers help identify it.
2. On Asterisk, Issabel or FreePBX, run `core show channels` at the moment of the failure to count the active calls and compare them with the trunk limit.
3. Turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and look for `Reason: Q.850;cause=17` and any `Retry-After` in the response.
4. Look for stuck channels: calls with a very long duration in `core show channels` are usually sessions that never got a BYE.
5. On 3CX, enable detailed (Verbose) logging and check in the Activity Log whether the refusal came from the trunk or from an internal rule.
6. Call the same number through another trunk or from a mobile phone. If it goes through, the problem is the route’s capacity, not the destination.

What to send the carrier

- Exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- Calling number (A) and destination number (B), in the format sent in the INVITE.
- The Call-ID of the call and the full response line, such as `SIP/2.0 486 Busy Here`.
- The Reason header with the Q.850 cause, if it came, and the Retry-After value.
- How many attempts failed, over what interval, and how many calls were active on the trunk at that moment.
- The pcap file of the capture, if you have one.

Frequently asked questions

What is the difference between SIP 486 and SIP 600?

486 Busy Here says the destination is busy at that point, and another device or another route may still answer. 600 Busy Everywhere is a global failure: the destination knows no other point will accept the call, and the proxy should not try other routes. On trunks and PBXs in Brazil, the 486 shows up far more often.

Why does the trunk return busy when nobody is on the phone?

Because on a SIP trunk 'busy' usually means there is no free channel. If the contract allows, for example, 30 concurrent calls and all of them are in use, the next one may get a 486. Some SBCs and carriers use SIP 503 with Q.850 cause 34 in this same situation. Compare your peak call count with the contracted limit.

Does a real test call get this busy signal?

Yes. A real test call goes through the trunk like any other call, so it gets the final SIP response the network returns, including the 486. When the result counts as a failure, the call is repeated before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 487 Request Terminated: what it means and how to fix it

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP 487 Request Terminated means the INVITE was ended by a CANCEL before the call was answered. Most often the caller hung up while it was ringing, or the PBX gave up after a timeout. It is not a destination failure. Only investigate when it shows up in excess: see who sent the CANCEL and after how many seconds.

SIP 487 in the reference

Code	487
Official name	Request Terminated
Class	4xx: request failure
Defined in	RFC 3261

How Monitor treats SIP 487

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	After ringing	Without ringing
A greeting or recorded message (the IVR)	Down	Down
A test tone or echo service	Down	Down
Other: it may ring or be answered	Up	Down

For "Other": With "Count ringback as reachable" unticked: Down.

What Monitor records for an IVR: "The line rang and nobody answered, but this circuit is supposed to answer every call by itself (SIP 487)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 487 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 487 mean in practice?

SIP 487 Request Terminated is the response a request gets when it is ended by a BYE or a CANCEL, according to RFC 3261. In the typical case, the flow goes like this: the INVITE goes out, the destination responds 180 or 183 and starts ringing, the caller gives up and your side sends a CANCEL. The destination confirms the CANCEL with 200 OK and answers the original INVITE with 487.

So the 487 comes from the destination, but the decision to end the call came from the caller. It shows up all the time in PBX logs and is almost always normal. The SIP 487 error only deserves investigation when the volume breaks the usual pattern, when the CANCEL always goes out after the same number of seconds (a sign of a configured timeout), or when the CANCEL comes from a device you do not control.

Common causes in Brazil

- The customer gave up while it was ringing or while listening to a message before the call was answered.
- Ring time ran out on the PBX: Asterisk's `Dial()` timeout, or the ring group or queue timeout in FreePBX and Issabel, expired and the PBX sent a CANCEL.
- A simultaneous ring group: one extension answered and the others received a CANCEL.
- An IVR queue that overflows to another destination after a while, canceling the previous attempt.
- A campaign dialer with a short ring time that gives up before the person can answer.
- An SBC between the PBX and the carrier's trunk with its own timer, which cancels the INVITE before the PBX's timeout.

How to diagnose it

1. Open the call in `sngrep` and find the CANCEL: see who sent it and how many seconds passed since the INVITE and since the 180 or 183.
2. If the CANCEL leaves your PBX after the same interval every time, such as 20, 30 or 60 seconds, review the `Dial()` timeout in the dial plan or the ring time of the group and the queue.
3. Turn on `pjsip set logger on` (or `sip set debug on` with `chan_sip`) and read the CANCEL's Reason header, such as `Reason: SIP;cause=200` or `Reason: Q.850;cause=16`.
4. On Asterisk, check the `DIALSTATUS` variable: `CANCEL` means the caller hung up before the call was answered, and `NOANSWER` means the ring time ran out.
5. On 3CX, enable detailed (Verbose) logging and see in the Activity Log who ended the call before it was answered.

6. Compare the share of calls ending in 487 with the day's total. A sudden rise usually comes from a configuration change or an overloaded queue.

What to send the carrier

- Only open a ticket if the CANCEL came from the carrier's side without anyone on your side hanging up.
- Exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- Calling number (A), destination number (B) and the Call-ID.
- The full sequence: INVITE, provisional responses, CANCEL and `SIP/2.0 487 Request Terminated`, with the time of each message.
- The CANCEL's Reason header and how many calls were affected.
- The pcap file of the capture, if you have one.

Frequently asked questions

Is SIP 487 an error?

Most of the time, no. It is the standard response to the INVITE after the calling side sends a CANCEL, and it appears in any PBX log with unanswered calls. It becomes a problem when it shows up in unusual volume, with a ring time that is too short, or with a CANCEL that nobody on your side asked for.

Which Q.850 cause corresponds to SIP 487?

RFC 3398 defines no Q.850 cause for the 487. On the way from the phone network to SIP, Q.850 cause 16 (Normal call clearing) usually becomes a BYE or a CANCEL, not a response code. That is why, when a call comes from a telephone exchange and the caller gives up, you see the CANCEL followed by the 487.

Why do I see 487 on extensions that did not answer a ring group?

In a simultaneous ring group, the PBX calls several extensions at once. When one answers, it sends a CANCEL to the others, and each one responds 487. The CANCEL may carry the Reason header indicating that the call was completed elsewhere. This is expected behavior.

SIP CODES

SIP 488 Not Acceptable Here: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 488 Not Acceptable Here means the destination received the INVITE but did not accept the media offer in the SDP. The most common cause is having no codec in common, such as offering only G.729 to a side that accepts G.711 A-law (PCMA). Compare the codecs in the offer with those the trunk accepts and read the Warning header.

SIP 488 in the reference

Code	488
Official name	Not Acceptable Here
Class	4xx: request failure

Defined in	RFC 3261
Q.850 causes that usually come with it	65 (Bearer capability not implemented), 70 (Only restricted digital information bearer capability is available (national use))

How Monitor treats SIP 488

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The phone at the other end cannot use any of the audio formats (codecs) the call offered, so the call could not connect (SIP 488)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 488 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 488 mean in practice?

SIP 488 Not Acceptable Here has the same meaning as [SIP 606](#), but it applies only to the resource addressed in the Request-URI, so the request may work elsewhere. Put simply: the destination received the INVITE, read the media offer in the SDP and did not accept the codec, the transport profile or some other session detail. The response may carry an SDP with the destination's capabilities and a Warning header, such as `305 Incompatible media format` or `304 Media type not available`.

It also shows up mid-call, as a response to a re-INVITE, for example when switching to T.38 fax. Under RFC 3398, a gateway converts the SIP 488 error into a Q.850 cause based on the Warning: cause 31 by default, or [cause 65 \(Bearer capability not implemented\)](#) when the Warning refers to bearer capability. In the other direction, causes 65 and 70 become 488.

Common causes in Brazil

- The PBX offers only G.729 or Opus, and the carrier's trunk (Vivo, Claro, Algar, Embratel or another) accepts G.711 A-law (PCMA).
- An imported phone or softphone with G.711 μ -law (PCMU) as its only codec, the default in the United States, talking to a trunk that only accepts PCMA.
- An inbound call where the carrier offers only G.729 and Asterisk does not have the codec module, so the PBX itself answers 488.
- Encrypted media: an offer with `RTP/SAVP` (SRTP) to a trunk that only accepts `RTP/AVP`, or the other way around.
- A re-INVITE for T.38 fax sent to a destination that does not support T.38.

- A codec change during a transfer or when taking a call off hold, rejected by the SBC mid-call.

How to diagnose it

1. Open the call in `sngrep`, go into the INVITE and read the `m=audio` line and the `a=rtpmap` lines. Payload 8 is PCMA, 0 is PCMU and 18 is G.729.
2. Read the response's Warning, for example `Warning: 305 sbc.carrier "Incompatible media format"`, and the SDP that came with it, if any.
3. On Asterisk, Issabel or FreePBX, run `pjsip show endpoint <trunk>` and check the `allow` line (with `chan_sip`, `sip show peer <trunk>`). `core show translation` shows which codec conversions are available.
4. Compare the profile on the `m=audio` line: RTP/AVP is unencrypted audio, RTP/SAVP is SRTP. Check the endpoint's `media_encryption` option.
5. Run a test with the trunk accepting only `alaw`. If the call goes through, adjust the codec order.
6. If the 488 comes on a re-INVITE, see what triggered it (T.38, hold, transfer) with `pjsip set logger on` or, on 3CX, with detailed (Verbose) logging.

What to send the carrier

- Exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- Calling number (A), destination number (B) and the Call-ID.
- The full `SIP/2.0 488 Not Acceptable Here` line and the Warning header.
- The SDP of the INVITE sent, with the `m=audio` and `a=rtpmap` lines.
- Whether it happened on the initial INVITE or on a re-INVITE, and how many calls failed.
- The pcap file of the capture, if you have one.

Frequently asked questions

Which codec should I use on a SIP trunk in Brazil?

G.711 A-law (PCMA, payload 8) is the codec of the Brazilian phone network and the one carriers accept most widely. Many trunks also accept G.729 to save bandwidth. Keeping PCMA in the trunk's codec list avoids most 488s. Confirm the exact list with the carrier, because each contract may differ.

What is the difference between SIP 488 and SIP 606?

Both say the media offer was not accepted. The 488 applies only to that destination, and the request may work elsewhere. The 606 is a global failure: no other point should accept it, and the proxy stops trying. In practice, the 488 shows up far more often on trunks and PBXs.

Does a real test call catch this error?

Yes. A real test call negotiates media with the destination like any other call, so it gets the 488 when the offer is rejected. When the result counts as a failure, the call is repeated before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 500 Server Internal Error: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 500 Server Internal Error means a server on the route hit an unexpected condition and could not process the request. Common causes are a software fault in the PBX, the SBC or the carrier, or a 503 from a downstream server that a proxy turned into a 500. Find out who generated the response before you open a ticket.

SIP 500 in the reference

Code	500
Official name	Server Internal Error
Class	5xx: server failure
Defined in	RFC 3261
Q.850 causes that usually come with it	111 (Protocol error, unspecified), 127 (Interworking, unspecified), 41 (Temporary failure)

How Monitor treats SIP 500

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider could not complete the call because of a fault on its side: an internal fault at the provider (SIP 500)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 500 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 500 mean in practice?

SIP 500 Server Internal Error is the generic 5xx response: the server ran into an unexpected condition that kept it from fulfilling the request. RFC 3261 allows it to carry Retry-After, and the client may try again after a few seconds. The code does not say what failed, only that the failure happened on the side of the server that responded.

One detail matters for troubleshooting: under RFC 3261, a proxy that receives only a SIP 503 from a downstream server should not pass it along, and should generate a 500 instead. So a SIP 500 error can hide an outage further down the route. Where SIP meets the phone network, RFC 3398 converts Q.850 cause 127 (Interworking, unspecified), cause 111 and any unmapped cause into a 500.

Common causes in Brazil

- A software fault in the PBX (Asterisk, Issabel, FreePBX, 3CX) after a version or module upgrade.

- An internal dependency down in the PBX or the carrier's platform, such as a database or a routing service.
- A proxy or SBC that received a 503 from the next server and turned it into a 500.
- A SIP-to-E1 gateway that received a Q.850 cause with no mapping from the exchange, such as 111 or 127, and returned a 500.
- The carrier's SBC (Vivo, Claro, Algar, Embratel or another) failing intermittently during maintenance or a restart.

How to troubleshoot

1. Capture the call with `sngrep` and identify who generated the 500 from the Via, Server and User-Agent headers.
2. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and look for the Reason header, such as `Reason: Q.850;cause=127`, and any Retry-After.
3. If the 500 came from your PBX, read the full Asterisk log for the same time window, looking for module or database errors.
4. Run `pjsip show registrations` and `core show channels` to rule out a downed trunk or an overloaded PBX.
5. In 3CX, turn on verbose logging and check the Activity Log for the step that failed.
6. Place the call a few times. A 500 on every call points to configuration or software; an intermittent 500 points to overload or a failure further along.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full `SIP/2.0 500 Server Internal Error` line and the Server, Reason and Retry-After headers, if present.
- How many attempts failed and whether the failure is constant or intermittent.
- The pcap file of the capture, if you have one.

Frequently asked questions

Does a SIP 500 always come from the carrier?

No. Any server on the route can generate a 500: your PBX, an SBC on your network, the carrier's SBC or equipment further along. Before you open a ticket, capture the response and check the Via, Server and User-Agent headers. If the 500 originated in your PBX, the problem is local.

Which Q.850 cause goes with a SIP 500?

Under RFC 3398, a gateway converts a SIP 500 into Q.850 cause 41 (Temporary failure). In the other direction, cause 111 (Protocol error, unspecified), cause 127 (Interworking, unspecified) and any cause without its own mapping become a 500. A `Reason: Q.850;cause=127` header points to a failure in the conversion between networks.

Does a real test call catch this error?

Yes. A real test call receives the final SIP response the network returns, including a 500. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 501 Not Implemented: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 501 Not Implemented means the server does not support the functionality needed to fulfill the request, typically a SIP method it does not recognize. It shows up more often on REFER, UPDATE, INFO or MESSAGE than on INVITE. Check the trace for the method that got the 501 and set the PBX to stop sending it to that trunk.

SIP 501 in the reference

Code	501
Official name	Not Implemented
Class	5xx: server failure
Defined in	RFC 3261
Q.850 causes that usually come with it	29 (Facility rejected), 79 (Service or option not implemented, unspecified)

How Monitor treats SIP 501

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider could not complete the call because of a fault on its side (SIP 501)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 501 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 501 mean in practice?

SIP 501 Not Implemented says the server does not support the functionality needed to fulfill the request. Under RFC 3261, it is the right response when the recipient does not recognize the request method and cannot support it for any user. Proxies forward requests of any method, so a 501 usually comes from the

end device, such as the carrier's SBC or the destination PBX.

In practice, a SIP 501 error rarely blocks the whole call. It tends to break one feature: a transfer that uses REFER, a session refresh with UPDATE, DTMF over INFO or a MESSAGE. The CSeq header in the response shows the refused method, such as `CSeq: 102 REFER`. Where SIP meets the phone network, RFC 3398 converts a 501 into Q.850 cause 79 and maps causes 29 and 79 back to 501 in the other direction.

Common causes in Brazil

- A call transfer from the phone or the PBX sending REFER to a carrier trunk (Vivo, Claro, Algar, Embratel or another) that does not implement that method.
- UPDATE sent to refresh the session or update call details to equipment that does not support it.
- INFO for DTMF or MESSAGE for text sent by IP phones or softphones to a trunk that does not handle those methods.
- SUBSCRIBE or PUBLISH for presence and busy lamp field (BLF) going to the trunk by mistake instead of staying between extensions.
- A SIP-to-E1 gateway that received Q.850 cause 79 or 29 from the exchange and returned a 501.

How to troubleshoot

1. Open the trace in `sngrep` and filter for the 501 response. The CSeq header shows which method was refused.
2. Send or capture an OPTIONS to the trunk and read the Allow header in the response: it lists the methods the equipment accepts.
3. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and repeat the action that failed, such as the transfer, to see which request went out.
4. Run `pjsip show endpoint <trunk>` and review the options that generate extra methods, such as the DTMF mode and session timers.
5. In 3CX, turn on verbose logging and check the Activity Log for the request that got the 501.
6. If the 501 came in response to the INVITE, look in the SDP or the Require and Supported headers for an extension the destination does not implement.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full `SIP/2.0 501 Not Implemented` line and the CSeq header with the refused method.
- Which feature failed for the user (transfer, DTMF, the whole call).
- The pcap file of the capture, if you have one.

Frequently asked questions

What is the difference between SIP 501 and SIP 405?

501 Not Implemented says the server does not recognize the method and does not support it for any user. 405 Method Not Allowed says the server recognizes the method, but it is not allowed for that resource. In both cases, the Allow header in the response, or in a reply to OPTIONS, shows which methods the equipment accepts.

Does a 501 in reply to OPTIONS take the trunk down?

Usually not. OPTIONS is used as a keepalive for the trunk, and any response shows the equipment on the other side is up. What matters for calls is the response to INVITE. Still, check whether your PBX or SBC treats a 501 to OPTIONS as a failure, because that depends on the configuration.

Does a real test call catch this error?

Yes, when it comes in response to the INVITE. A real test call receives the final SIP response the network returns. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 502 Bad Gateway: what it means and how to fix it

By 2N5 team Updated October 20, 2026

Short answer: SIP 502 Bad Gateway means a server acting as a gateway or proxy received an invalid response from the next server on the route. On trunks, it usually points to a failure in the carrier's interconnection with the destination network, or a destination out of service. Use the headers to find who generated the 502 and give the carrier the Call-ID.

SIP 502 in the reference

Code	502
Official name	Bad Gateway
Class	5xx: server failure
Defined in	RFC 3261
Q.850 causes that usually come with it	27 (Destination out of order), 38 (Network out of order)

How Monitor treats SIP 502

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider could not complete the call because of a fault on its side: a server on the route got a bad reply from the next one (SIP 502)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 502 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 502 mean in practice?

SIP 502 Bad Gateway is a 5xx response. Under RFC 3261, the responding server was acting as a gateway or proxy and received an invalid response from the downstream server it contacted to fulfill the request. In other words, whoever sends the 502 is not the direct culprit: it is reporting that the next hop failed.

Where SIP meets the phone network, RFC 3398 converts [Q.850 cause 27 \(Destination out of order\)](#) into a 502, and a 502 into [cause 38 \(Network out of order\)](#). That is why a SIP 502 error on a call to a landline, a mobile or an 0800 (toll-free) number usually means the destination exchange could not deliver the call. In [SIP trunk monitoring](#), what matters is telling a one-off 502 apart from a 502 that keeps repeating for the same destination.

Common causes in Brazil

- A destination out of service on another carrier's network, signaled by the exchange with Q.850 cause 27 and converted into a 502 by the gateway.
- A failure in the interconnection between your trunk's carrier (Vivo, Claro, Algar, Embratel or another) and the network that hosts the number.
- A destination PBX with its own IVR, or an E1 link, that is down, so the carrier's gateway receives a failure.
- An SBC on your network, in front of Asterisk, Issabel, FreePBX or 3CX, receiving a response from the PBX that it does not accept.
- Upgraded equipment on the route that started sending non-standard messages to the previous hop.

How to troubleshoot

1. Capture the call with `sngrep` and see who generated the 502 from the Via, Server and User-Agent headers.
2. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and look for `Reason: Q.850;cause=27` or another cause in the response.
3. Call the same number through another trunk or from a mobile phone. If that also fails, the problem is at the destination.
4. Try other numbers on the same destination carrier. If they all fail, suspect the interconnection.
5. If the 502 comes from your own SBC, capture on the PBX side as well and compare the response the PBX sent with the one the SBC passed on.
6. In 3CX, turn on verbose logging and check the Activity Log for the trunk that returned the 502.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full `SIP/2.0 502 Bad Gateway` line and the Server and Reason headers.

- Whether other numbers on the same destination network also fail, and how many attempts were made.
- The pcap file of the capture, if you have one.

Frequently asked questions

Is a SIP 502 a problem with my PBX?

Rarely. A 502 is generated by equipment in the middle of the route, such as a proxy, an SBC or a gateway to the phone network, when the next hop returns something invalid. If the 502 comes from the carrier, the failure is usually beyond it. If it comes from your own SBC, check the response the PBX sent to it.

Which Q.850 cause goes with a SIP 502?

Under RFC 3398, a gateway converts a SIP 502 into Q.850 cause 38 (Network out of order). In the other direction, Q.850 cause 27 (Destination out of order) becomes a 502. A `Reason: Q.850;cause=27` header means the destination exchange considered the destination out of service.

Does a real test call catch this error?

Yes. A real test call receives the final SIP response the network returns, including a 502. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 503 Service Unavailable: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 503 Service Unavailable means the server is temporarily unable to process the call because of overload or maintenance. On trunks, the most common reasons are the carrier running out of channels or routes, or an SBC being down. Read the Retry-After and the Q.850 cause in the Reason header and, if it persists, open a ticket with the Call-ID.

SIP 503 in the reference

Code	503
Official name	Service Unavailable
Class	5xx: server failure
Defined in	RFC 3261
Q.850 causes that usually come with it	34 (No circuit/channel available), 38 (Network out of order), 41 (Temporary failure), 42 (Switching equipment congestion), 47 (Resource unavailable, unspecified), 58 (Bearer capability not presently available), 88 (Incompatible destination)

How Monitor treats SIP 503

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: “The provider could not complete the call because of a fault on its side: the provider's service is unavailable (SIP 503).”

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 503 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 503 mean in practice?

SIP 503 Service Unavailable says the server is temporarily unable to process the request because of overload or maintenance. Under RFC 3261, it may carry Retry-After to say when to try again, and the recipient should look for an alternate server. Without Retry-After, the client treats the response as a 500.

On trunks, the SIP 503 error is the most common way of saying “there is no way to complete this right now.” The Q.850 cause in the Reason header helps explain why. Under RFC 3398, the following become a 503: [cause 34 \(No circuit/channel available\)](#), [cause 38 \(Network out of order\)](#), [cause 41 \(Temporary failure\)](#), [cause 42 \(Switching equipment congestion\)](#), [cause 47 \(Resource unavailable, unspecified\)](#), plus causes 58 and 88. In the other direction, a 503 becomes cause 41.

Common causes in Brazil

- Concurrent channels used up on the carrier's trunk (Vivo, Claro, Algar, Embratel or another), with Reason: Q.850;cause=34 .
- The calls-per-second limit reached by dialers and campaigns that place many calls at once.
- Scheduled maintenance or a restart of the carrier's SBC, sometimes with Retry-After.
- Congestion or a temporary failure on the route to the destination network, with cause 41 or 42.
- A PBX (Asterisk, Issabel, FreePBX, 3CX) with no outbound trunk available, generating the 503 for the extension.
- An SBC license or capacity exhausted on your own network.

How to troubleshoot

1. Capture the call with `sngrep` and identify who generated the 503 from the Via, Server and User-Agent headers.
2. Turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and read the Reason, such as Reason: Q.850;cause=34 , and the Retry-After.
3. Run `pjsip show registrations` to confirm the trunk is registered and `core show channels` to count the active calls.
4. Compare the number of active calls with the channel limit in your contract. If you are at the ceiling, the 503 is about capacity.

- 5. In 3CX, turn on verbose logging and check the Activity Log for which trunk responded and whether there was an alternate route.
- 6. Place the call again after the Retry-After time and through another trunk. If the 503 keeps happening on only one trunk, the problem is in that trunk.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full SIP/2.0 503 Service Unavailable line and the Reason and Retry-After headers.
- How many attempts failed, over what time window, and how many calls were active on the trunk.
- Whether the failure affects all destinations or only some.
- The pcap file of the capture, if you have one.

Frequently asked questions

What do I do with the Retry-After header in a SIP 503?

It says how many seconds to wait before sending new requests to that server. Under RFC 3261, whoever receives the 503 should try an alternate server and not send other requests to the same server during that time. Without Retry-After, the 503 should be treated as a 500. If your PBX has a second trunk, set it up as an alternate route.

Does a 503 always come from the carrier?

No. The PBX itself or an SBC on your network also generates 503s, for example when no outbound trunk is available. Before you open a ticket, capture the response and check the Via, Server and User-Agent headers. If the 503 originated inside your network, the carrier will not find anything.

Does a real test call catch this error?

Yes. A real test call goes through the trunk like any other call and receives the final SIP response the network returns, including a 503. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 504 Server Time-out: what it means and how to fix it

By 2N5 team Updated October 20, 2026

Short answer: SIP 504 Server Time-out means a server on the route did not get a timely response from another server it contacted to process the call. On trunks, it usually points to a slow or downed next hop, such as the carrier's interconnection or a routing lookup. Check who generated the 504 and how long it took.

SIP 504 in the reference

Code	504
Official name	Server Time-out
Class	5xx: server failure

Defined in	RFC 3261
Q.850 causes that usually come with it	102 (Recovery on timer expiry)

How Monitor treats SIP 504

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The provider could not complete the call because of a fault on its side: a server on the route got no reply from the next one in time (SIP 504)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 504 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 504 mean in practice?

SIP 504 Server Time-out is a 5xx response. Under RFC 3261, the responding server did not receive a timely response from an external server it contacted to process the request. That external server can be the next proxy on the route, a gateway to the phone network or a lookup service, such as a DNS server or a routing database.

Where SIP meets the phone network, RFC 3398 ties the 504 to [Q.850 cause 102 \(Recovery on timer expiry\)](#) in both directions. A SIP 504 error shows the first hop is up, but something after it took too long. If your PBX gets no response at all, the code you see is usually a [SIP 408](#), generated locally.

Common causes in Brazil

- A slow or downed interconnection between your trunk's carrier (Vivo, Claro, Algar, Embratel or another) and the destination network.
- A slow external lookup in call routing, such as DNS or a routing database queried by the SBC.
- A SIP-to-E1 gateway that got no response from the exchange before its timer ran out, signaled with Q.850 cause 102.
- An SBC on your network, in front of Asterisk, Issabel, FreePBX or 3CX, waiting on an overloaded or hung PBX.
- Packet loss or a firewall block between two servers on the route, so responses never arrive.

How to troubleshoot

1. Capture the call with `sngrep` and see who generated the 504 from the Via, Server and User-Agent headers, and how much time passed between the INVITE and the response.

2. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and look for `Reason: Q.850;cause=102`.
3. If the 504 comes from your own SBC, capture on the PBX side as well and check whether the PBX responded and how late.
4. Run `core show channels` and check the server load at the time of the failure.
5. Call the same number through another trunk and call other numbers on the same destination network, to tell the route apart from the destination.
6. In 3CX, turn on verbose logging and check the Activity Log for how long the trunk took to respond.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full `SIP/2.0 504 Server Time-out` line, the Server and Reason headers and the time between the INVITE and the response.
- Whether other destinations also fail and how many attempts were made.
- The pcap file of the capture, if you have one.

Frequently asked questions

What is the difference between SIP 504 and SIP 408?

408 Request Timeout says the request got no response within the expected time, and it is often generated by the PBX itself when nobody answered. A 504 comes from a server in the middle of the route that received your request but did not get a timely response from the server it needed to contact. With a 504, you know at least the first hop is up.

Which Q.850 cause goes with a SIP 504?

Under RFC 3398, a SIP 504 becomes Q.850 cause 102 (Recovery on timer expiry), and cause 102 becomes a 504 in the other direction. A `Reason: Q.850;cause=102` header means a protocol timer expired in the phone network before the call could be completed.

Does a real test call catch this error?

Yes. A real test call receives the final SIP response the network returns, including a 504. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 600 Busy Everywhere: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 600 Busy Everywhere means the destination was reached, is busy, and knows no other endpoint would accept the call right now. It is a 6xx global failure, so the proxy stops trying other routes. On a trunk, it can mean the concurrent channels are used up. Check who generated the response and the channel limit on the route.

SIP 600 in the reference

Code	600
Official name	Busy Everywhere
Class	6xx: global failure
Defined in	RFC 3261
Q.850 causes that usually come with it	17 (User busy)

How Monitor treats SIP 600

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Up

For "Other": With "Count busy as reachable" unticked: Down.

What Monitor records for an IVR: "The line was busy: the phone is on another call, but this circuit is supposed to answer every call by itself (SIP 600)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 600 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 600 mean in practice?

SIP 600 Busy Everywhere is a 6xx response, the global failure class. According to RFC 3261, the destination's system was contacted successfully, but the destination is busy and does not wish to take the call right now. The destination should only send this response when it knows no other endpoint could accept the call; it may carry Retry-After with a better time to try.

The practical consequence is that a proxy that receives a 600 cancels its pending attempts and does not look for other routes. That is what sets it apart from [SIP 486](#), which says "busy here" and leaves other routes open. On trunks, a 600 can also mean the concurrent channels are used up, when the carrier's or the company's SBC is configured to respond that way. A SIP 600 error is less common than a 486, which is why it is worth identifying the equipment that generated it.

Common causes in Brazil

- Concurrent channels used up on an SBC configured to respond 600 instead of 486, to keep the PBX from trying other routes.
- The concurrent-call limit per SIP account reached on a carrier trunk (Vivo, Claro, Algar, Embratel or another).

- A hosted contact-center platform or IVR that responds 600 when every seat is in use.
- A ring group with every extension busy, when the destination system decides none of them can answer.
- Intermediate equipment that turns a busy signal from another network into a 600 instead of passing the 486 along.

How to troubleshoot

1. Capture the call with `sngrep` and see who generated the 600 from the Via, Server and User-Agent headers in the response.
2. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and look for `Reason: Q.850;cause=17` and any `Retry-After`.
3. Run `core show channels` at the time of the failure to count the active calls and compare them with the route's channel limit.
4. Check the log for whether the PBX tried another route after the response. If the overflow route was never used, the 600 may be why.
5. In 3CX, turn on verbose logging and check the Activity Log for the trunk that returned the 600.
6. Place the call through another trunk or from a mobile phone. If it completes, the problem is the route's capacity, not the destination.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A) and the called number (B), in the format sent in the INVITE.
- The Call-ID and the full response line, such as `SIP/2.0 600 Busy Everywhere`.
- The Reason and Retry-After headers, if present.
- How many attempts failed, over what time window, and how many calls were active on the trunk.
- The pcap file of the capture, if you have one.

Frequently asked questions

Is a SIP 600 more serious than a 486?

Not necessarily. Both say the destination is busy. The difference is scope: a 486 applies only to that endpoint, while a 600 warns that no other endpoint will accept, so the proxy should not try other routes. If your PBX has an overflow route, a 600 can keep it from being used.

Which Q.850 cause goes with a SIP 600?

Under RFC 3398, a gateway that converts a SIP 600 for the phone network uses Q.850 cause 17 (User busy). In the other direction, cause 17 becomes a 486, not a 600. So a 600 that reaches your PBX was almost always generated by SIP equipment, not translated from a telephone exchange.

Does a real test call catch this code?

Yes. A real test call receives the final SIP response the network returns, including a 600. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP 603 Decline: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 603 Decline means the destination device was reached, but the user or the system explicitly refused the call, and no other endpoint should be tried. The most common causes are the person rejecting the call or call blocking. Check whether the refusal came from the destination, the PBX or the carrier.

SIP 603 in the reference

Code	603
Official name	Decline
Class	6xx: global failure
Defined in	RFC 3261
Q.850 causes that usually come with it	21 (Call rejected)

How Monitor treats SIP 603

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Up

For "Other": With "Count busy as reachable" unticked: Down.

What Monitor records for an IVR: "The phone at the other end declined the call, but this circuit is supposed to answer every call by itself (SIP 603)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 603 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 603 mean in practice?

SIP 603 Decline is a 6xx global failure. Under RFC 3261, the destination device was contacted successfully, but the user does not wish to or cannot take part in the call. The response may carry Retry-After with a better time to call, and should only be sent when the destination knows no other endpoint will answer. A proxy that receives a 603 cancels its pending attempts.

What separates it from a busy signal is intent: with a 603, someone or some rule decided to refuse. Where SIP meets the phone network, RFC 3398 converts a 603 into [Q.850 cause 21 \(Call rejected\)](#). A SIP 603 error that shows up on every call to a number is usually a blocking rule, not a person declining.

There are more specific codes for unwanted calls, such as 607 Unwanted and 608 Rejected, but a lot of equipment still uses 603.

Common causes in Brazil

- A person who declined the call on a mobile, an IP phone or a softphone.
- A block list or time-of-day rule on the destination PBX (Asterisk, Issabel, FreePBX, 3CX, Intelbras) configured to refuse.
- A destination that refuses calls with no caller ID (anonymous) or from an unrecognized calling number.
- A spam-call blocking app or service on the destination side flagging the calling number.
- An IVR or contact-center platform set to refuse after hours instead of playing a message.

How to troubleshoot

1. Capture the call with `sngrep` and see who generated the 603 from the Via, Server and User-Agent headers, and whether there was ringing (180 or 183) before the refusal.
2. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and look for `Reason: Q.850;cause=21` and any Retry-After.
3. Check the calling number (A) and the privacy header in the INVITE. An anonymous caller or an odd number format is a common reason for refusal.
4. Call the same number from another trunk and from a mobile phone. If only one calling number is refused, suspect that number is blocked.
5. In 3CX, turn on verbose logging and check the Activity Log for whether the refusal came from the trunk or from an internal rule.
6. If the 603 comes from your own PBX, review block lists, time-of-day rules and inbound routes.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full `SIP/2.0 603 Decline` line and the Reason header, if present.
- Whether other calling numbers can complete calls to the same destination.
- How many attempts were refused and over what time window.
- The pcap file of the capture, if you have one.

Frequently asked questions

What is the difference between SIP 603 and SIP 486?

486 Busy Here says the destination is busy and cannot answer right now at that endpoint. 603 Decline says the destination actively refused the call, and that no other endpoint will answer. Depending on the network and the phone, pressing 'decline' on a mobile can arrive as a 603 or as a 486.

Which Q.850 cause goes with a SIP 603?

Under RFC 3398, a gateway converts a SIP 603 into Q.850 cause 21 (Call rejected). In the other direction, cause 21 normally becomes 403 Forbidden, and a 603 may be used when the refusal came from the user. A `Reason: Q.850;cause=21` header confirms a rejection.

Does a real test call catch this code?

Yes. A real test call receives the final SIP response the network returns, including a 603. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 604 Does Not Exist Anywhere: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 604 Does Not Exist Anywhere means the server has authoritative information that the user or number in the Request-URI does not exist anywhere. It is the global version of 404 Not Found. The most common causes are a disconnected number, a typo or an invalid format. Check the number sent in the INVITE before you open a ticket.

SIP 604 in the reference

Code	604
Official name	Does Not Exist Anywhere
Class	6xx: global failure
Defined in	RFC 3261
Q.850 causes that usually come with it	1 (Unallocated (unassigned) number)

How Monitor treats SIP 604

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The call was refused, and the provider said it would fail everywhere (SIP 604)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 604 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 604 mean in practice?

SIP 604 Does Not Exist Anywhere is a 6xx global failure. Under RFC 3261, the server has authoritative information that the user named in the Request-URI does not exist anywhere. That is a stronger statement than [SIP 404](#), which only says the user was not found in that domain. A proxy that receives a

604 cancels its pending attempts and does not look for other routes.

Where SIP meets the phone network, RFC 3398 converts a 604 into [Q.850 cause 1 \(Unallocated number\)](#). The SIP 604 error is rare: most carriers signal a nonexistent number with a 404. When it does show up, check the number that went out in the INVITE first, because a wrong format produces the same result as a disconnected number.

Common causes in Brazil

- A number that was disconnected or never assigned by the carrier.
- The wrong format in the INVITE: a missing area code (DDD), a missing ninth digit on a mobile number or an incomplete 0800 (toll-free) number, or a country code when the trunk expects the national format.
- Digit manipulation on the PBX's outbound route (Asterisk, Issabel, FreePBX, 3CX) cutting digits or adding extra ones.
- An extension or user that does not exist on a destination PBX configured to respond 604 instead of 404.
- A recently ported number whose routing has not been updated yet on the route used.

How to troubleshoot

1. Open the call in `sngrep` and read the Request-URI and the To header of the INVITE. Check the number digit by digit.
2. In Asterisk, Issabel or FreePBX, turn on `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and compare the number dialed by the extension with the one sent to the trunk.
3. Review the digit manipulation rules on the outbound route, such as prefixes being added or removed.
4. Call the same number from a mobile phone. If it completes, the problem is the format or the route, not the number.
5. In 3CX, turn on verbose logging and check the Activity Log for the number sent to the trunk.
6. Look for `Reason: Q.850;cause=1` in the response, which confirms an unallocated number.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A) and the called number (B), exactly as they went out in the Request-URI.
- The Call-ID and the full `SIP/2.0 604 Does Not Exist Anywhere` line, with the Reason header, if present.
- Whether the number completes from a mobile phone or another trunk.
- The pcap file of the capture, if you have one.

Frequently asked questions

What is the difference between SIP 604 and SIP 404?

404 Not Found says the user was not found in that domain or on that server. 604 Does Not Exist Anywhere goes further: the server has authoritative information that the user does not exist anywhere, so no other route should be tried. In practice, the 404 is far more common on Brazilian trunks.

Which Q.850 cause goes with a SIP 604?

Under RFC 3398, a gateway converts a SIP 604 into Q.850 cause 1 (Unallocated number). In the other direction, cause 1 becomes a 404, not a 604. So a 604 that reaches your PBX was almost always generated by SIP equipment, not translated from a telephone exchange.

Does a real test call catch this error?

Yes. A real test call receives the final SIP response the network returns, including a 604. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

SIP CODES

SIP 606 Not Acceptable: what it means and how to fix it

By [2N5 team](#) Updated October 20, 2026

Short answer: SIP 606 Not Acceptable means the destination was reached but did not accept parts of the proposed session, such as media, bandwidth or addressing, and indicates that no other endpoint would. It is the global version of 488. The most common cause is a codec or SRTP mismatch. Read the Warning header and compare the SDP offer.

SIP 606 in the reference

Code	606
Official name	Not Acceptable
Class	6xx: global failure
Defined in	RFC 3261

How Monitor treats SIP 606

The reading depends on what should answer the number, which is the question 2N5 answers when it adds each circuit. These are the rules the product applies today, generated from the same function that judges every test call.

What should answer the number	Reading
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

What Monitor records for an IVR: "The call was refused, and the provider said it would fail everywhere (SIP 606)."

A call read as **Down** alerts nobody on its own: it is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident, with each call's SIP code and Q.850 cause as evidence for the carrier. At the default pace, every circuit is tested every 15 minutes.

If the carrier returns SIP 606 when 2N5's calling account registers, before any call, the test is **monitoring blind**: it is a problem on 2N5's side, it does not count against the circuit and it does not alert the customer.

What does SIP 606 mean in practice?

SIP 606 Not Acceptable is a 6xx global failure. Under RFC 3261, the destination's system was contacted successfully, but some aspects of the session description, such as the requested media, the bandwidth or the addressing style, are not acceptable. It indicates the destination wants to communicate but cannot support the session as described. The response may carry one or more Warning headers explaining why.

What separates it from [SIP 488](#) is scope: a 488 applies only to that destination, while a 606 warns that no other endpoint should accept, and the proxy cancels its pending attempts. In practice, the SIP 606 error is rare on Brazilian trunks. When it appears, it almost always comes from an SBC or a platform configured to refuse the media offer outright.

Common causes in Brazil

- An offer without G.711 A-law (PCMA), the codec of the Brazilian phone network, sent to a carrier SBC (Vivo, Claro, Algar, Embratel or another) that only accepts that codec.
- Encrypted media (RTP/SAVP) offered to a destination that only accepts unencrypted audio, or the other way around.
- An offer with only video, or with a mandatory video stream, to a destination that only handles audio.
- A softphone or phone with Opus or G.722 as its only codecs, talking to an Asterisk, Issabel, FreePBX or 3CX PBX that does not have those codecs on the trunk.
- Bandwidth declared in the SDP above what the destination accepts.

How to troubleshoot

1. Open the call in `sngrep` and read the Warning in the response, for example `Warning: 305 sbc.destino "Incompatible media format"`.
2. In the INVITE, check the `m=audio` line and the `a=rtpmap` lines. Payload 8 is PCMA, 0 is PCMU and 18 is G.729.
3. Check the transport profile on the `m=audio` line: `RTP/AVP` is unencrypted audio and `RTP/SAVP` is SRTP.
4. In Asterisk, Issabel or FreePBX, run `pjsip show endpoint <trunk>` and check `allow` and `media_encryption` (in `chan_sip`, `sip show peer <trunk>`). With `pjsip set logger on`, you can see the full SDP.
5. In 3CX, turn on verbose logging and check the Activity Log for which codecs were offered.
6. Place the call again with only `alaw` on the trunk. If it completes, adjust the codec list on the phone or the PBX.

What to send your carrier

- The exact date and time with the time zone, for example 2026-10-07 14:32:10 -03:00.
- The calling number (A), the called number (B) and the Call-ID.
- The full `SIP/2.0 606 Not Acceptable` line and every Warning header.
- The SDP of the INVITE you sent.
- How many calls failed and whether the problem started after a change.
- The pcap file of the capture, if you have one.

Frequently asked questions

What does the Warning header say in a SIP 606?

It explains why the session was not accepted. The three-digit code gives the reason, for example 304 for media type not available, 305 for incompatible media format and 370 for insufficient bandwidth. The quoted text after the server name often adds detail. Without a Warning, compare the SDP you sent with what the destination accepts.

Which Q.850 cause goes with a SIP 606?

Under RFC 3398, it depends on the Warning: the gateway uses Q.850 cause 31 (Normal, unspecified) by default, or cause 65 (Bearer capability not implemented) when the Warning is about bearer capability. No Q.850 cause is converted into a 606 in the other direction; cause 65 becomes a 488.

Does a real test call catch this error?

Yes. A real test call negotiates media with the destination like any other call, so it receives the 606 when the session is refused. When the result counts as a failure, the call is retried before any alert goes out, and the alert is sent by email.

REFERENCE

Q.850 causes (hangup cause): the full list and SIP map

The Q.850 cause is the number the telephone network gives for the end of a call: 16 when it ends normally, 17 when the number is busy, 34 when no channel is free. It travels in the release message on E1 trunks and, on SIP trunks, in the `Reason: Q.850;cause=N` header.

Where does the Q.850 cause show up?

On E1 trunks (SS7/ISUP or ISDN), the cause comes in the call's release message. On SIP trunks, the carrier or the SBC repeats it in the Reason header of the response (RFC 6432). In Asterisk, Issabel and FreePBX it is the `HANGUPCAUSE` variable and the CDR's `hangupcause` field.

The 72 ITU-T Q.850 causes

Cause	Official name	Class
Cause 1	Unallocated (unassigned) number	Normal event (1 to 31)
Cause 2	No route to specified transit network (national use)	Normal event (1 to 31)
Cause 3	No route to destination	Normal event (1 to 31)
Cause 4	Send special information tone	Normal event (1 to 31)
Cause 5	Misdialled trunk prefix (national use)	Normal event (1 to 31)
Cause 6	Channel unacceptable	Normal event (1 to 31)
Cause 7	Call awarded and being delivered in an established channel	Normal event (1 to 31)
Cause 8	Pre-emption	Normal event (1 to 31)
Cause 9	Pre-emption – circuit reserved for reuse	Normal event (1 to 31)
Cause 13	Call completed elsewhere	Normal event (1 to 31)
Cause 14	QoR: ported number	Normal event (1 to 31)
Cause 16	Normal call clearing	Normal event (1 to 31)
Cause 17	User busy	Normal event (1 to 31)
Cause 18	No user responding	Normal event (1 to 31)
Cause 19	No answer from user (user alerted)	Normal event (1 to 31)
Cause 20	Subscriber absent	Normal event (1 to 31)
Cause 21	Call rejected	Normal event (1 to 31)
Cause 22	Number changed	Normal event (1 to 31)
Cause 23	Redirection to new destination	Normal event (1 to 31)
Cause 24	Call rejected due to feature at the destination	Normal event (1 to 31)
Cause 25	Exchange – routing error	Normal event (1 to 31)
Cause 26	Non-selected user clearing	Normal event (1 to 31)
Cause 27	Destination out of order	Normal event (1 to 31)
Cause 28	Invalid number format (address incomplete)	Normal event (1 to 31)
Cause 29	Facility rejected	Normal event (1 to 31)

Cause	Official name	Class
Cause 30	Response to STATUS ENQUIRY	Normal event (1 to 31)
Cause 31	Normal, unspecified	Normal event (1 to 31)
Cause 34	No circuit/channel available	Resource unavailable (34 to 47)
Cause 38	Network out of order	Resource unavailable (34 to 47)
Cause 39	Permanent frame mode connection out of service	Resource unavailable (34 to 47)
Cause 40	Permanent frame mode connection operational	Resource unavailable (34 to 47)
Cause 41	Temporary failure	Resource unavailable (34 to 47)
Cause 42	Switching equipment congestion	Resource unavailable (34 to 47)
Cause 43	Access information discarded	Resource unavailable (34 to 47)
Cause 44	Requested circuit/channel not available	Resource unavailable (34 to 47)
Cause 46	Precedence call blocked	Resource unavailable (34 to 47)
Cause 47	Resource unavailable, unspecified	Resource unavailable (34 to 47)
Cause 49	Quality of service not available	Service or option not available (49 to 63)
Cause 50	Requested facility not subscribed	Service or option not available (49 to 63)
Cause 53	Outgoing calls barred within CUG	Service or option not available (49 to 63)
Cause 55	Incoming calls barred within CUG	Service or option not available (49 to 63)
Cause 57	Bearer capability not authorized	Service or option not available (49 to 63)
Cause 58	Bearer capability not presently available	Service or option not available (49 to 63)
Cause 62	Inconsistency in designated outgoing access information and subscriber class	Service or option not available (49 to 63)
Cause 63	Service or option not available, unspecified	Service or option not available (49 to 63)
Cause 65	Bearer capability not implemented	Service or option not implemented (65 to 79)
Cause 66	Channel type not implemented	Service or option not implemented (65 to 79)
Cause 69	Requested facility not implemented	Service or option not implemented (65 to 79)
Cause 70	Only restricted digital information bearer capability is available (national use)	Service or option not implemented (65 to 79)
Cause 79	Service or option not implemented, unspecified	Service or option not implemented (65 to 79)
Cause 81	Invalid call reference value	Invalid message (81 to 95)
Cause 82	Identified channel does not exist	Invalid message (81 to 95)
Cause 83	A suspended call exists, but this call identity does not	Invalid message (81 to 95)
Cause 84	Call identity in use	Invalid message (81 to 95)

Cause	Official name	Class
Cause 85	No call suspended	Invalid message (81 to 95)
Cause 86	Call with the requested call identity has been cleared	Invalid message (81 to 95)
Cause 87	User not member of CUG	Invalid message (81 to 95)
Cause 88	Incompatible destination	Invalid message (81 to 95)
Cause 90	Non-existent CUG	Invalid message (81 to 95)
Cause 91	Invalid transit network selection (national use)	Invalid message (81 to 95)
Cause 95	Invalid message, unspecified	Invalid message (81 to 95)
Cause 96	Mandatory information element is missing	Protocol error (96 to 111)
Cause 97	Message type non-existent or not implemented	Protocol error (96 to 111)
Cause 98	Message not compatible with call state or message type non-existent or not implemented	Protocol error (96 to 111)
Cause 99	Information element/parameter non-existent or not implemented	Protocol error (96 to 111)
Cause 100	Invalid information element contents	Protocol error (96 to 111)
Cause 101	Message not compatible with call state	Protocol error (96 to 111)
Cause 102	Recovery on timer expiry	Protocol error (96 to 111)
Cause 103	Parameter non-existent or not implemented – passed on (national use)	Protocol error (96 to 111)
Cause 110	Message with unrecognized parameter discarded	Protocol error (96 to 111)
Cause 111	Protocol error, unspecified	Protocol error (96 to 111)
Cause 127	Interworking, unspecified	Interworking (127)

From Q.850 cause to SIP code (RFC 3398, §7.2.4.1)

Q.850 cause	SIP response	Note
1 (unallocated number)	SIP 404 Not Found	
2 (no route to network)	SIP 404 Not Found	
3 (no route to destination)	SIP 404 Not Found	
16 (normal call clearing)	None	(*) usually results in a BYE or CANCEL, no status code
17 (user busy)	SIP 486 Busy Here	
18 (no user responding)	SIP 408 Request Timeout	
19 (no answer from the user)	SIP 480 Temporarily Unavailable	
20 (subscriber absent)	SIP 480 Temporarily Unavailable	
21 (call rejected)	SIP 403 Forbidden	(+) if cause location is user, 603 could be given instead
22 (number changed (w/o diagnostic))	SIP 410 Gone	
22 (number changed (w/ diagnostic))	301 Moved Permanently	new number from diagnostic SHOULD go in Contact

Q.850 cause	SIP response	Note
23 (redirection to new destination)	SIP 410 Gone	
26 (non-selected user clearing)	SIP 404 Not Found	(=) ANSI: misrouted ported number
27 (destination out of order)	SIP 502 Bad Gateway	
28 (address incomplete)	SIP 484 Address Incomplete	
29 (facility rejected)	SIP 501 Not Implemented	
31 (normal unspecified)	SIP 480 Temporarily Unavailable	
34 (no circuit available)	SIP 503 Service Unavailable	
38 (network out of order)	SIP 503 Service Unavailable	
41 (temporary failure)	SIP 503 Service Unavailable	
42 (switching equipment congestion)	SIP 503 Service Unavailable	
47 (resource unavailable)	SIP 503 Service Unavailable	
55 (incoming calls barred within CUG)	SIP 403 Forbidden	
57 (bearer capability not authorized)	SIP 403 Forbidden	
58 (bearer capability not presently available)	SIP 503 Service Unavailable	
65 (bearer capability not implemented)	SIP 488 Not Acceptable Here	
70 (only restricted digital avail)	SIP 488 Not Acceptable Here	
79 (service or option not implemented)	SIP 501 Not Implemented	
87 (user not member of CUG)	SIP 403 Forbidden	
88 (incompatible destination)	SIP 503 Service Unavailable	
102 (recovery of timer expiry)	SIP 504 Server Time-out	
111 (protocol error)	SIP 500 Server Internal Error	
127 (interworking unspecified)	SIP 500 Server Internal Error	
Any other ((not listed))	SIP 500 Server Internal Error	default for unlisted causes (SHOULD)
44 (requested circuit or channel not available)	None	untranslatable: should not result in any SIP status; originating switch retries another CIC

From SIP code to Q.850 cause (RFC 3398, §8.2.6.1)

SIP response	Q.850 cause	Note
SIP 400 Bad Request	41 (Temporary Failure)	
SIP 401 Unauthorized	21 (Call rejected)	(*) try to authenticate first
SIP 402 Payment required	21 (Call rejected)	
SIP 403 Forbidden	21 (Call rejected)	
SIP 404 Not found	1 (Unallocated number)	
SIP 405 Method not allowed	63 (Service or option unavailable)	
SIP 406 Not acceptable	79 (Service/option not implemented)	(+)
SIP 407 Proxy authentication required	21 (Call rejected)	(*) try to authenticate first
SIP 408 Request timeout	102 (Recovery on timer expiry)	
SIP 410 Gone	22 (Number changed (w/o diagnostic))	
SIP 413 Request Entity too long	127 (Interworking)	(+)
SIP 414 Request-URI too long	127 (Interworking)	(+)
SIP 415 Unsupported media type	79 (Service/option not implemented)	(+)
SIP 416 Unsupported URI Scheme	127 (Interworking)	(+)
SIP 420 Bad extension	127 (Interworking)	(+)
SIP 421 Extension Required	127 (Interworking)	(+)
SIP 423 Interval Too Brief	127 (Interworking)	(+)
SIP 480 Temporarily unavailable	18 (No user responding)	
SIP 481 Call/Transaction Does not Exist	41 (Temporary Failure)	
SIP 482 Loop Detected	25 (Exchange - routing error)	
SIP 483 Too many hops	25 (Exchange - routing error)	
SIP 484 Address incomplete	28 (Invalid Number Format)	(+)
SIP 485 Ambiguous	1 (Unallocated number)	
SIP 486 Busy here	17 (User busy)	
SIP 487 Request Terminated	From the Warning header	
SIP 488 Not Acceptable here	From the Warning header	default 31; 65 if Warning is about bearer capability
SIP 500 Server internal error	41 (Temporary failure)	
SIP 501 Not implemented	79 (Not implemented, unspecified)	

SIP response	Q.850 cause	Note
SIP 502 Bad gateway	38 (Network out of order)	
SIP 503 Service unavailable	41 (Temporary failure)	
SIP 504 Server time-out	102 (Recovery on timer expiry)	
SIP 505 Version Not Supported	127 (Interworking)	(+) RFC text prints 504; corrected to 505 by Errata ID 2161 (Held for Document Update)
SIP 513 Message Too Large	127 (Interworking)	(+)
SIP 600 Busy everywhere	17 (User busy)	
SIP 603 Decline	21 (Call rejected)	
SIP 604 Does not exist anywhere	1 (Unallocated number)	
SIP 606 Not acceptable	From the Warning header	default 31; 65 if Warning is about bearer capability
Any other (not listed)	31 (Normal, unspecified)	default for unlisted status codes

Sources: [ITU-T Q.850](#) (10/18, with amendment 1 of 04/19), [RFC 3398](#) (with its [errata](#)) and [RFC 6432](#), checked on October 7, 2026.

Q.850 CAUSES

Q.850 cause 1 (Unallocated (unassigned) number): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 1 (Unallocated (unassigned) number) means the network that cleared the call has no subscriber assigned to that number. It is not congestion or a trunk failure: the number does not exist there. It usually comes from a deactivated number or a wrong dialing format. Check the number exactly as it went out in the INVITE.

Cause 1 in the reference

Cause	1
Official name (ITU-T Q.850)	Unallocated (unassigned) number
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 404 Not Found

How Monitor treats cause 1

When the carrier gives the Q.850 cause in the **Reason** header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 1 reads: "the number is not allocated".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 1 mean in practice?

Q.850 cause 1, Unallocated (unassigned) number, reports that the destination number reached a switch that does not have it assigned. The format may be correct and the route to that switch may exist; what is missing is an active subscriber with that number. The cause belongs to the normal event class, so it does not indicate an equipment fault. On E1 trunks with SS7, it arrives in the cause parameter of the ISUP REL message; on an ISDN E1 (PRI), in the Q.931 DISCONNECT or RELEASE. In SIP, RFC 3398 (section 7.2.4.1) maps cause 1 to **SIP 404 Not Found**, and many carriers repeat the original value in the `Reason: Q.850;cause=1` header (RFC 6432). In the other direction, 404, 485 and 604 map back to cause 1. In Asterisk it shows up in `${HANGUPCAUSE}` and, where the CDR records the field, as `hangupcause 1`.

When it appears

- A deactivated or canceled number, including an 0800 (toll-free) number taken down when its contract ended.
- A mobile number dialed without the ninth digit, or a landline dialed with one digit too many or too few.
- A PBX (Asterisk, Issabel, FreePBX, 3CX) that normalizes the number incorrectly: it strips the area code, keeps the outbound 0, or sends the carrier selection code (CSP) on a trunk that expects E.164 format.
- A newly contracted or ported DID range that has not yet been activated on the carrier's switch.
- A ported number whose routing database still points to the old carrier.

How to diagnose it

1. Call the same number from a mobile phone. If that does not complete either, the problem is the number, not your trunk.
2. Capture the signaling with `pjsip set logger on` (or `sip set debug on` in `chan_sip`) and copy the Request-URI and the To header exactly as they went out.
3. Compare the format sent with what the carrier's contract asks for: `+5511...`, `11...` or `0xx11...` with CSP.
4. In `sngrep`, open the call and look at the final response: a 404 with `Reason: Q.850;cause=1` confirms the rejection came from the network.
5. In 3CX, turn on Verbose logging and check in the Activity Log which number went out on the trunk after the dialing rules.
6. If the number is yours (a DID), ask the carrier to confirm that the range is active and routed to your trunk.

What to send the carrier

- Date and time of each attempt, with the time zone (for example, 2026-10-07 14:32, UTC-3).
- The A number (caller) and B number (called) in the format they went out in the INVITE.
- The Call-ID and the full response line, such as `SIP/2.0 404 Not Found`, with the Reason header.
- If the B number is one of your DID ranges, the contract number or the circuit ID.
- The pcap of the call, if you have it.

Frequently asked questions

Does cause 1 always mean the number does not exist?

It means the network that cleared the call does not recognize the number as assigned. In practice it can be a deactivated number, but also a valid number sent in the wrong format (missing the ninth digit Brazilian mobile numbers carry, missing the area code, with an extra prefix) or a DID range the carrier has not yet activated on its switch.

What is the difference between cause 1 and cause 3?

Cause 1 says the number is not assigned. Cause 3 (No route to destination) says the network has no route to the destination, even though the number may exist on another network. Both become SIP 404 Not Found under RFC 3398, so only the Reason header or the trace on the TDM side shows which of the two happened.

Q.850 CAUSES

Q.850 cause 3 (No route to destination): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 3 (No route to destination) means the network that cleared the call has no path to the called number's network, although the number may exist. It usually appears with prefixes not enabled on the contracted plan or with missing interconnection routes. Test the same destination through another carrier and send the trace to yours.

Cause 3 in the reference

Cause	3
Official name (ITU-T Q.850)	No route to destination
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 404 Not Found

How Monitor treats cause 3

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier.

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 3 mean in practice?

Q.850 cause 3, No route to destination, says the switch that received the call does not know how to reach the network where the destination lives. Unlike [cause 1](#), it does not claim the number is deactivated: the number may work normally when dialed through another network. The cause is in the normal class and travels in the ISUP REL message on SS7 links, in the ISDN DISCONNECT on an E1 PRI, or in a `Reason: Q.850;cause=3` header (RFC 6432) on SIP trunks. RFC 3398 (section 7.2.4.1) maps cause 3 to [SIP 404 Not Found](#), the same response as cause 1, so a 404 on its own is not enough to tell

the two apart. In Asterisk, the `$_HANGUPCAUSE` variable is set to 3. Note: on mobile phones and LTE modems, the "cause code 3" in EMM messages comes from a different table (Illegal UE) and has nothing to do with this cause.

When it appears

- Destinations outside the contracted plan: international numbers, 0300, 0500 or 0900 numbers not enabled on the trunk.
- Calls to a carrier your carrier has no route configured to for that prefix.
- A new or recently ported prefix that the carrier's routing table does not yet recognize.
- An intermediate SBC or gateway with no outbound route for the dialed pattern, which clears with cause 3 instead of passing the call on.
- A PBX with the wrong outbound route, sending the number through a trunk that only accepts local calls.

How to diagnose it

1. Dial the same number through another carrier or from a mobile phone. If it completes, the number exists and the problem is the route.
2. Test other numbers in the same prefix to find out whether the missing route affects the whole prefix or a single number.
3. Turn on `pjsip set logger on` and check whether the 404 carries `Reason: Q.850;cause=3`; without the Reason, ask the carrier for the original cause.
4. In FreePBX or Issabel, check under Outbound Routes which route matched the dialed pattern and which trunk the call went out on.
5. In 3CX, turn on Verbose logging and see in the Activity Log which trunk was used and the response received.
6. On an ISDN E1, use `pri set debug on span 1` and read the cause's location field: it shows whether the rejection came from the local network, the transit network or the remote network.

What to send the carrier

- Date and time of each attempt, with the time zone.
- The A and B numbers as they went out, and the affected prefix.
- The Call-ID, the `SIP/2.0 404 Not Found` line and the Reason header received.
- The result of the test through the other path (the number completes through another carrier or from a mobile phone).
- The explicit request: confirm whether the prefix is enabled on the plan and routed on the interconnection.

Frequently asked questions

Why does the carrier return 404 if the number exists?

Because RFC 3398 maps both cause 1 (unassigned number) and cause 3 (no route) to SIP 404 Not Found. The 404 only says the network did not complete the call to that number. A Reason header with `cause=3`, or the trace on the TDM side, shows that the number may exist and that what is missing is the route.

Does enabling international or 0300 calling on the trunk fix cause 3?

If the destination was outside the contracted plan, yes: the carrier enables the prefix and the call gets a route. If the prefix was already enabled, the missing route is further along, in an interconnection or an outdated table, and only the carrier can fix it. That is why it pays to send example numbers and times.

Q.850 CAUSES

Q.850 cause 16 (Normal call clearing): what it means

By [2N5 team](#) Updated October 20, 2026

Short answer: Q.850 cause 16 (Normal call clearing) is the normal ending: one of the parties hung up after the call was set up. It does not indicate a failure. It becomes a problem when many calls end with 16 within a few seconds, with no audio, or right after the IVR answers. In that case, check who sent the BYE and whether RTP flowed in both directions.

Cause 16 in the reference

Cause	16
Official name (ITU-T Q.850)	Normal call clearing
Class	Normal event (1 to 31)
SIP response per RFC 3398	None: it usually ends the call with BYE or CANCEL

How Monitor treats cause 16

When the carrier gives the Q.850 cause in the **Reason** header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 16 reads: "the call was cleared normally".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 16 mean in practice?

Q.850 cause 16, Normal call clearing, indicates the call was cleared because one of the users involved asked for it, in other words, someone hung up. It is the most common cause in any call log and does not represent a failure. On an SS7 link it travels in the ISUP REL message; on an ISDN E1, in the Q.931 DISCONNECT. In SIP there is no equivalent error response: RFC 3398 (section 7.2.4.1) says cause 16 normally results in a BYE or a CANCEL, often with **Reason: Q.850;cause=16**. In Asterisk it appears as hangup cause 16 in `_${HANGUPCAUSE}` and, in the CDR, as hangupcause 16 when the field is recorded. In 3CX logs and those of other PBXs, the same cause usually comes alongside the side that hung up. On LTE networks, the "cause code 16" in EMM messages is something else (MSC temporarily not reachable) and should not be confused with this cause.

When it appears

- On every answered call that ends when someone hangs up: this is the normal case.
- On calls lasting 1 to 3 seconds after the 200 OK, when the IVR answers and drops. The cause is normal, but the answer is not.

- When audio flows in only one direction and the person hangs up because they hear nothing: the log shows 16 and the real defect is in the media (NAT, firewall, SIP ALG).
- When a device clears the call on its own, because of a maximum call duration or missing RTP, and reports 16 in the BYE, because as far as signaling goes the clearing was orderly.
- On dialers that hang up when they detect voicemail or silence.

How to diagnose it

1. Split the cause 16 calls by duration. A high volume under 5 seconds deserves attention; the rest is usually normal.
2. In `sngrep`, see who sent the BYE: your side or the carrier. That tells you who hung up.
3. Check whether RTP flowed in both directions: in Asterisk, `rtp set debug on` shows the packets; in `sngrep`, the flow shows the RTP when it passes through the machine.
4. Open the BYE and read the Reason header. A `cause=16` coming from the carrier right after the answer means the destination hung up.
5. If the drops happen at a fixed time (for example, always near 32 seconds), look for a lost ACK or a session timer: the hangup is automatic, even if it arrives with cause 16.

What to send the carrier

- Usually nothing: an isolated cause 16 is no reason to open a ticket.
- If there is a pattern, a list of short calls with date, time and time zone, the A and B numbers, and the duration of each.
- The Call-ID of each example and which side sent the BYE, with the Reason header.
- If audio was missing, the pcap with SIP and RTP showing the direction with no packets.
- When the pattern started and whether it repeats with the same destination.

Frequently asked questions

Is hangup cause 16 in Asterisk an error?

No. In Asterisk, ``${HANGUPCAUSE}`` equal to 16 means someone hung up in an orderly way, whether your user or the destination. It is the most common value in any call log. It is only worth investigating when it comes with very short durations, missing audio or complaints of dropped calls, because then the normal hangup may be the result of another problem.

Why does cause 16 have no SIP response code?

Because it is not a rejection. Under RFC 3398, cause 16 normally results in a BYE, when the call was already answered, or a CANCEL, when it was still ringing, with no status code. The value usually travels inside those messages in the ``Reason: Q.850;cause=16`` header, defined in RFC 3326.

Does a real test call help reveal when cause 16 hides a problem?

Yes, because it is a real call: besides the final SIP response, it checks the audio in both directions and measures voice quality (MOS-LQO). So a call that ends with cause 16 but has no audio does not go unnoticed. Failures are retried before they become an alert, and the alert goes out by email.

Q.850 cause 17 (User busy): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 17 (User busy) means the destination did not accept the call because it was busy. It is a normal cause: the network worked and the destination's equipment is compatible. It becomes a problem when an 0800 number, an IVR or a trunk returns busy often, which usually means a channel limit has been reached. Compare concurrent calls with what you contracted.

Cause 17 in the reference

Cause	17
Official name (ITU-T Q.850)	User busy
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 486 Busy Here

How Monitor treats cause 17

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 17 reads: "the line was busy".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 17 mean in practice?

Q.850 cause 17, User busy, is used when the destination cannot take another call because it is busy. The called equipment is compatible and the network worked; it was the destination that refused. Under RFC 3398 (section 7.2.4.1), cause 17 becomes [SIP 486 Busy Here](#), and in the other direction both 486 and [SIP 600 Busy Everywhere](#) map back to 17. On SS7 links it comes in the ISUP REL; on ISDN E1, in the DISCONNECT; in SIP, sometimes in a `Reason: Q.850; cause=17`. In Asterisk, `${HANGUPCAUSE}` is set to 17 and `DIALSTATUS` is usually BUSY. Many PBXs also respond 486 when they hit a concurrent call limit, so "busy" does not always mean a person on the phone: it may just be capacity running out.

When it appears

- The person is already on another call, with no call waiting.
- A mobile user who rejects the call: many networks return busy in that case.
- An extension with "do not disturb" set to respond busy.
- A concurrent call limit reached on the extension, the contracted SIP trunk or the PBX license.
- An 0800 number or IVR with every channel in use at peak time.
- A call queue configured to return busy when it fills up.

How to diagnose it

1. Call from another line at the same moment. If it is always busy, suspect a limit, not a busy person.
2. In Asterisk, `core show channels` shows how many calls are active; compare it with the number of channels in the contract.
3. In `sngrep`, see where the 486 came from: your PBX, the SBC or the carrier. The Server or User-Agent headers in the response help.

4. Check the call limits configured on the extension and the trunk (FreePBX, Issabel) and on the 3CX license, when it limits concurrent calls.
5. See whether the busy responses cluster at peak hours, which points to a lack of capacity.

What to send the carrier

- Date, time and time zone of the attempts that returned busy.
- The A and B numbers and how many attempts in a row failed.
- The Call-ID, the `SIP/2.0 486 Busy Here` line and the Reason received.
- How many concurrent calls there were on your side at the time (the output of `core show channels`).
- If the destination is one of your 0800 numbers, the number of channels contracted for it.

Frequently asked questions

Is cause code 17 on a mobile phone the same as Q.850 cause 17?

No. On LTE mobile networks, the 'cause code 17' that shows up in modem or phone logs comes from the 3GPP EMM causes and means Network failure. Q.850 cause 17 is the busy cause in telephony, used in ISUP, ISDN and the SIP Reason header. The numbers match; the tables do not.

Is a busy signal on my 0800 number always the carrier's fault?

Not necessarily. If the 486 comes from your PBX or your IVR, the limit is yours: channels, license or queue configuration. If it comes from the carrier before reaching your equipment, it may be the channel limit in the 0800 contract. The trace shows who responded, and that decides who to open the ticket with.

Q.850 CAUSES

Q.850 cause 18 (No user responding): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 18 (No user responding) means the destination did not respond to the call request, either by ringing or by answering, within the allowed time. It is not busy or a rejection: the device simply did not react. In SIP it is usually a registered extension that stopped responding, because of NAT or a firewall. Check the registration and the traffic.

Cause 18 in the reference

Cause	18
Official name (ITU-T Q.850)	No user responding
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 408 Request Timeout

How Monitor treats cause 18

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 18 reads: "no device responded".

Cause 18 also decides how a **SIP 480** is read, since on its own a 480 does not say whether the destination rang or nothing answered:

What should answer the number	SIP 480 with cause 18
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 18 mean in practice?

Q.850 cause 18, No user responding, is used when the destination does not respond to the call setup request with an alerting (ringing) or answer indication within the allowed time. The difference from cause 19 is that with 19, the phone rang and nobody answered; with 18, not even the ringing was confirmed. Under RFC 3398 (section 7.2.4.1), cause 18 becomes **SIP 408 Request Timeout**. In the other direction, the RFC maps **SIP 480 Temporarily Unavailable** to 18, which is why a 480 from the carrier can show up on your PBX as cause 18. In SS7 it comes in the ISUP REL; in SIP, in a Reason: Q.850;cause=18 (RFC 6432). In Asterisk, `${HANGUPCAUSE}` is set to 18 and, where the CDR records the field, it appears as hangupcause 18.

When it appears

- A SIP extension registered behind NAT whose mapping has expired: the INVITE does not arrive and nobody responds.
- A mobile softphone in the background that does not wake up in time.
- An IP phone unplugged before its registration expired.
- A router with SIP ALG that alters or drops the INVITE.
- A mobile phone with no signal, on networks that return 18 instead of 20.
- A customer's PBX, behind an 0800 number, that receives the INVITE and does not even return 100 Trying.

How to diagnose it

1. In `sngrep`, see whether the destination responded at all: with no 100 Trying and no 180 Ringing, the INVITE probably did not arrive.
2. In Asterisk, `pjsip show contacts` shows whether the contact is reachable on qualify; `pjsip show registrations` shows your trunk's registration with the carrier.
3. Turn on `qualify_frequency` on the AOR (PJSIP) or `qualify=yes` (chan_sip) to detect dead contacts before the call.
4. Disable SIP ALG on the router and adjust the phone's NAT keepalive.
5. If the 408 comes from the carrier, check whether the INVITE left your firewall with `tcpdump -i any port 5060`.

What to send the carrier

- Date, time and time zone of each attempt.
- The A and B numbers.

- The Call-ID and the final response received (SIP/2.0 408 Request Timeout or 480), with the Reason.
- Whether there was a 100 Trying or 180 Ringing before the failure.
- The pcap showing that the INVITE left your network.

Frequently asked questions

What is the difference between cause 18 and cause 19?

With cause 18 (No user responding), the destination did not even confirm ringing in time. With cause 19 (No answer from user, user alerted), the phone rang and nobody answered. The first points to equipment that is unreachable or silent; the second, to nobody being available. Under RFC 3398, 18 becomes SIP 408 and 19 becomes SIP 480.

Why did a SIP 480 become cause 18 on my PBX?

Because RFC 3398 maps SIP 480 Temporarily Unavailable to Q.850 cause 18 when the call passes from SIP to the telephone network. If the carrier sent the Reason header, it carries the original cause, and it is better to trust it than the table's translation, which can hide a cause 20 or 31.

Q.850 CAUSES

Q.850 cause 20 (Subscriber absent): what it means

By 2N5 team Updated October 7, 2026

Short answer: Q.850 cause 20 (Subscriber absent) means the subscriber cannot be reached: a mobile phone that is off or has no signal, or a user not registered on any interface. It is a normal cause and does not point to a trunk failure. It becomes a problem when it happens on an extension or destination that should always be registered, or on many numbers at once.

Cause 20 in the reference

Cause	20
Official name (ITU-T Q.850)	Subscriber absent
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 480 Temporarily Unavailable

How Monitor treats cause 20

When the carrier gives the Q.850 cause in the Reason header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier.

Cause 20 also decides how a SIP 480 is read, since on its own a 480 does not say whether the destination rang or nothing answered:

What should answer the number	SIP 480 with cause 20
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down

What should answer the number	SIP 480 with cause 20
Other: it may ring or be answered	Down

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 20 mean in practice?

Q.850 cause 20, Subscriber absent, is used when the mobile station has left the network, when there is no radio contact with it, or when a user is not addressable on any user-network interface. In IP telephony, the direct equivalent is an extension with no active registration. The cause is in the normal class: the network worked and found the subscriber absent. Under RFC 3398 (section 7.2.4.1), cause 20 becomes **SIP 480 Temporarily Unavailable**. It arrives in the ISUP REL on SS7 links and, in SIP, can come in a `Reason: Q.850;cause=20` (RFC 6432). In Asterisk, `$_HANGUPCAUSE` is set to 20, and the CDR shows hangupcause 20 when the installation records that field. Because the same 480 is also used for other situations, only the Reason confirms that the subscriber is absent.

When it appears

- A mobile phone that is switched off, out of battery or outside the coverage area.
- A mobile phone in airplane mode or with no active SIM.
- A SIP extension with no registration: phone off, softphone closed or registration expired.
- A remote extension that lost its internet connection.
- A call to a mobile number the destination carrier cannot locate at that moment.

How to diagnose it

1. Confirm whether the destination is a mobile phone or an extension. With a mobile phone that is off, the cause is expected and there is nothing to fix on the trunk.
2. For extensions, use `pjsip show contacts` (or `sip show peers` in `chan_sip`) and see whether the extension has a registered contact.
3. Check the registration expiry time and whether the phone renews before it expires.
4. In `sngrep`, check whether the 480 came with `Reason: Q.850;cause=20` and from which device.
5. In 3CX, the extensions panel shows which ones are registered; cross-check it with the time of the failure.
6. If the cause appears for many different numbers at the same time, it is not an absent subscriber: investigate the route.

What to send the carrier

- Usually nothing: an absent subscriber is a condition at the destination.
- If many different numbers return 20 at the same time: date, time and time zone, and the list of B numbers.
- The Call-ID, the `SIP/2.0 480 Temporarily Unavailable` line and the Reason.
- The test through another line showing that the destinations were switched on and completing calls.

Frequently asked questions

Does a mobile phone that is switched off always return cause 20?

Not always. It depends on how the destination carrier handles the case: some return cause 20, others play a switched-off or out-of-area announcement before clearing, and others use cause 18 or 31. On calls to mobile phones, what reaches your PBX depends on the interconnection and on how each network translates the cause to SIP.

Is cause code 20 on a 4G modem the same thing?

No. The 'cause codes' in LTE EMM and ESM messages come from a different table, defined by 3GPP, with different meanings. Q.850 cause 20 belongs to telephony: it appears in ISUP, ISDN and the SIP Reason header. If the number shows up in a modem or phone log, look it up in the 3GPP table, not Q.850.

Q.850 CAUSES

Q.850 cause 21 (Call rejected): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 21 (Call rejected) means the equipment that cleared the call refused it, although it was neither busy nor incompatible. The rejection can come from the user or from the network. On SIP trunks it is common after an authentication failure, an IP that is not allowed, or a caller number that is not accepted. Check the credentials and the A number sent.

Cause 21 in the reference

Cause	21
Official name (ITU-T Q.850)	Call rejected
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 403 Forbidden

How Monitor treats cause 21

When the carrier gives the Q.850 cause in the **Reason** header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 21 reads: "the call was rejected".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 21 mean in practice?

Q.850 cause 21, Call rejected, means the equipment sending the cause does not want to accept the call, even though it could, because it is neither busy nor incompatible. The rejection can come from the user or from the network, through blocking rules. Under RFC 3398 (section 7.2.4.1), cause 21 becomes [SIP 403 Forbidden](#); if the cause's location is the user, the gateway may send [SIP 603 Decline](#). In the other direction, 401, 402, 403, 407 and 603 map back to 21, which explains why authentication errors show

up as rejected calls on the PBX. The cause arrives in the ISUP REL, the ISDN DISCONNECT or a Reason: Q.850;cause=21 (RFC 6432). In Asterisk, `${HANGUPCAUSE}` is set to 21, and the CDR shows hangupcause 21 when the field is recorded.

When it appears

- A wrong trunk password or username, or an expired registration: a 401 or 407 without re-authentication becomes 21.
- Your PBX's IP address not allowed by the carrier, which responds 403.
- A caller number (From or P-Asserted-Identity) that does not belong to the contracted trunk.
- A destination that blocks calls, such as a block list or a restriction on anonymous calls.
- Destinations outside the plan, such as international numbers, when the carrier responds 403 instead of 404.
- A person on a mobile phone who rejects the call, on networks that use 21 in that case.

How to diagnose it

1. Look at the original SIP response in `sngrep`: 401 or 407 points to authentication, 403 to permissions, 603 to a refusal by the destination.
2. Check `pjsip show registrations` (or `sip show registry` in `chan_sip`) and the trunk status in FreePBX, Issabel or 3CX.
3. Compare the A number sent in From and P-Asserted-Identity with the numbers in the contract.
4. Test another destination over the same trunk. If they all fail, it is the trunk; if only one fails, it is the destination.
5. Read the Warning or the Reason text: many carriers explain the reason for the rejection there.

What to send the carrier

- Date, time and time zone of the attempts.
- The A and B numbers exactly as they went out, including From and P-Asserted-Identity.
- The Call-ID and the `SIP/2.0 403 Forbidden` line (or 603), with the Reason and the Warning.
- The public IP address the call went out from.
- Whether the trunk was registered at the time of the failure.

Frequently asked questions

Why does a password error show up as cause 21?

Because RFC 3398 maps SIP 401 Unauthorized and 407 Proxy Authentication Required to Q.850 cause 21 when authentication is not successfully retried. The PBX or gateway, when it translates the rejection into telephony terms, shows a rejected call. The real cause is in the SIP response, so it pays to look at the signaling before blaming the destination.

What is the difference between 403 and 603 for cause 21?

RFC 3398 maps cause 21 to SIP 403 Forbidden by default. When the cause comes with the user location, meaning the destination itself refused, the gateway may use SIP 603 Decline. In practice, a 403 suggests a block by the network or the trunk, and a 603 suggests a refusal by the destination.

Q.850 CAUSES**Q.850 cause 27 (Destination out of order): what it means**

By 2N5 team Updated October 20, 2026

Short answer: Q.850 cause 27 (Destination out of order) means the destination cannot be reached because the interface to it is not working: a physical link down, a PBX switched off, or a customer trunk not responding. It is a real failure, common with 0800 numbers and IVRs that terminate on the company's own PBX. Check the destination's equipment and link.

Cause 27 in the reference

Cause	27
Official name (ITU-T Q.850)	Destination out of order
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 502 Bad Gateway

How Monitor treats cause 27

When the carrier gives the Q.850 cause in the **Reason** header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 27 reads: "the destination is out of order".

Cause 27 also decides how a **SIP 480** is read, since on its own a 480 does not say whether the destination rang or nothing answered:

What should answer the number	SIP 480 with cause 27
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 27 mean in practice?

Q.850 cause 27, Destination out of order, means the destination cannot be reached because the interface to it is not working properly. The signaling message could not be delivered to the remote user, for example because of a failure at the physical or data link layer, or because the user's equipment is disconnected. Unlike the busy or absent causes, there is a real failure here. Under RFC 3398 (section 7.2.4.1), cause 27 becomes [SIP 502 Bad Gateway](#). In the other direction, the same RFC maps 502 to cause 38, so the cause can change when it crosses two gateways. It arrives in the ISUP REL, the ISDN DISCONNECT and a **Reason: Q.850;cause=27**. In Asterisk, `$_HANGUPCAUSE` is set to 27. In LTE, the "cause code 27" in ESM messages (Missing or unknown APN) has nothing to do with this cause.

When it appears

- The E1 on the customer's PBX is in alarm or out of sync: the carrier's switch sees the link as down.

- The SIP trunk between the carrier and the customer's PBX is not responding (OPTIONS unanswered) and is marked out of service.
- The internet link at the customer's site is down, for an 0800 number or an IVR that terminates on the local PBX.
- The PBX is switched off, restarting, or its SIP service has stopped.
- The customer's firewall is blocking the carrier's IP address after a rule change.

How to diagnose it

1. Confirm on the destination PBX that it is up and the service is running (`core show uptime` in Asterisk).
2. On an E1 through DAHDI, use `dahdi show status` to see alarms and `pri show spans` for the state of the D channel.
3. On a SIP trunk, check whether the carrier's OPTIONS are arriving and being answered with 200 OK; in `sngrep`, filter by the OPTIONS method.
4. Test the site's internet link and the firewall rules for the carrier's signaling and media IP addresses.
5. In 3CX, check the trunk status on the dashboard and the Activity Log at the time of the failure.
6. If everything is up on your side, ask the carrier for the state of the circuit as seen from its switch.

What to send the carrier

- Date, time and time zone when the failure started, and whether it is still going on.
- The called number (for example, the 0800 number) and the A numbers used in the tests.
- The Call-ID and the response received (`SIP/2.0 502 Bad Gateway`), with the Reason.
- The E1 circuit ID or the SIP trunk identifier.
- The state on your side: E1 alarms, OPTIONS status, current public IP address.

Frequently asked questions

Is cause 27 my fault or the carrier's?

It depends on where the destination is. If the number terminates on your own PBX, the cause means the network cannot deliver the call to it: a link, equipment or firewall on your side, or the circuit between the two. If the destination belongs to someone else, the failure is in their equipment. The trace and the cause's location field help tell them apart.

Does a real test call catch cause 27 on an 0800 number?

Yes, because it is a real call placed to the number: it gets the same final SIP response a customer would get when the destination is out of service. A failure is retried before it becomes an alert, so nobody is notified over an isolated event, and the alert goes out by email.

Q.850 CAUSES

Q.850 cause 31 (Normal, unspecified): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 31 (Normal, unspecified) is a generic cause: it reports a normal event when no other cause in the normal class applies. On its own it does not indicate a failure. The trouble is that many gateways use it as a default value and hide the real reason. Look for the original SIP response and the Reason on the previous leg.

Cause 31 in the reference

Cause	31
Official name (ITU-T Q.850)	Normal, unspecified
Class	Normal event (1 to 31)
SIP response per RFC 3398	SIP 480 Temporarily Unavailable

How Monitor treats cause 31

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 31 reads: "an unspecified network fault".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 31 mean in practice?

Q.850 cause 31, Normal, unspecified, is meant for reporting a normal event only when no other cause in the normal class applies. In other words, it is the cause used by whoever does not know, or does not say, the reason. Under RFC 3398 (section 7.2.4.1), cause 31 becomes [SIP 480 Temporarily Unavailable](#). In the other direction it is even more common: the RFC sets 31 as the default value for any SIP response not in the table, and also for 488 and 606 when the Warning is not about bearer capability. That is why a cause 31 on your PBX can come from dozens of different SIP responses. It appears in the ISUP REL, in a `Reason: Q.850;cause=31` and, in Asterisk, in `${HANGUPCAUSE}`. In LTE, the "cause code 31" in ESM messages (Request rejected, unspecified) comes from a different table.

When it appears

- A SIP-TDM gateway that receives a SIP response with no match in RFC 3398 and applies the default 31.
- A media negotiation failure (488 or 606) without a Warning about the bearer.
- Carriers that use 31 as a generic response to any rejection.
- Calls to a mobile phone that is switched off, on networks that do not give the reason in detail.
- Older equipment that does not implement more specific causes.

How to diagnose it

1. Do not stop at cause 31: find the original SIP response with `pjsip set logger on` (or `sip set debug on`) or in `sngrep`.
2. If the 31 came from the carrier in a Reason, ask it for the original cause or response from the destination network.
3. Group the cause 31 calls by destination and time. A single destination points to the far end; many destinations point to the route.

4. See whether there was a 488 or 606 with a Warning: then the problem is the codec, not the network.
5. In Asterisk 11 or newer, the `HANGUPCAUSE` function with the `tech` type keeps the SIP response for each dialed channel; write it to the log next to the cause.
6. In 3CX, turn on Verbose logging so the Activity Log shows the full SIP response.

What to send the carrier

- Date, time and time zone of the attempts.
- The A and B numbers.
- The Call-ID, the SIP response received and the Reason with cause=31.
- An explicit request for the original cause on its leg of the call, since 31 does not identify the reason.
- The pcap, if you have it, and the list of affected destinations.

Frequently asked questions

Is cause 31 an error?

By definition, no: it belongs to the normal class and only says there is no more specific cause. In practice, it shows up a lot in real failures, because gateways use it as a default value when they do not know how to translate a SIP response. If calls with cause 31 are not completing, treat it as a failure and look for the original response.

What is the difference between cause 16 and cause 31?

Cause 16 (Normal call clearing) says someone hung up an established call. Cause 31 (Normal, unspecified) says there was a normal event with no specific cause, and it usually appears on calls that never even completed. 16 is the expected end of a conversation; 31 calls for investigation when it repeats.

Q.850 CAUSES

Q.850 cause 34 (No circuit/channel available): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 34 (No circuit/channel available) means no suitable circuit or channel was free to complete the call at that moment. It is a lack of resources, not a wrong number. On SIP and E1 trunks it usually means a concurrent channel limit has been reached, yours or the carrier's. Compare active calls with what you contracted and try again.

Cause 34 in the reference

Cause	34
Official name (ITU-T Q.850)	No circuit/channel available
Class	Resource unavailable (34 to 47)
SIP response per RFC 3398	SIP 503 Service Unavailable

How Monitor treats cause 34

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 34 reads: "no circuit or channel was available".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 34 mean in practice?

Q.850 cause 34, No circuit/channel available, means there is no appropriate circuit or channel available to handle the call at that moment. It opens the resource unavailable class: the number may be right and the route may exist, but some leg ran out of capacity. Under RFC 3398 (section 7.2.4.1), cause 34 becomes `SIP 503 Service Unavailable`, the same response as causes 38, 41, 42 and 47. On E1 with SS7 it arrives in the ISUP REL and, on an ISDN E1, in the DISCONNECT; in SIP, in a `Reason: Q.850;cause=34` (RFC 6432), sometimes together with `Retry-After`. In Asterisk, `_${HANGUPCAUSE}` is set to 34, the CDR shows hangupcause 34 when it records the field, and `DIALSTATUS` is usually CONGESTION. In 3CX, the failure appears in the Activity Log with the 503 response received from the trunk.

When it appears

- A SIP trunk with every contracted channel in use.
- An E1 with all 30 channels busy, or with some of its channels blocked.
- Dialer or campaign peaks that exceed the trunk's limit.
- A carrier interconnection route congested for a specific destination.
- A GSM or E1 gateway with ports unavailable.

How to diagnose it

1. At the time of the failure, run `core show channels` in Asterisk and compare the total number of calls with the channels in the contract.
2. See whose 503 it is: your SBC's, your PBX's or the carrier's. In `sngrep`, look at the User-Agent or Server in the response.
3. Check whether the response carries `Retry-After` and respect the interval before trying again.
4. On an E1, `pri show spans` and `dahdi show status` show whether there are blocked channels or alarms.
5. Split by destination: if only one prefix fails, the shortage of channels is at the interconnection, not on your trunk.

What to send the carrier

- Date, time and time zone, ideally with the start and end of the period with failures.
- The A and B numbers and how many attempts failed.
- The Call-ID, the `SIP/2.0 503 Service Unavailable` line and the Reason with cause=34.
- How many concurrent calls you had at the time and how many the contract provides for.
- The affected destination prefixes.

Frequently asked questions

Does cause 34 mean I need to buy more channels?

Sometimes. If you were using every contracted channel at the time of the failure, then yes, the limit is yours. If there were channels to spare and the cause came from the carrier, the missing circuit is inside its network or at an interconnection, and it is the carrier's to fix. That is why it pays to note how many calls were active when the failure happened.

What is the difference between cause 34 and cause 41?

Both become SIP 503. Cause 34 says specifically that no free circuit or channel was available, in other words capacity. Cause 41 says there was a temporary network failure, unrelated to capacity. In practice, 34 tends to come with traffic peaks, while 41 appears during instability and also as the translation of SIP 500 and 503.

Q.850 CAUSES

Q.850 cause 38 (Network out of order): what it means

By [2N5 team](#) · Updated October 7, 2026

Short answer: Q.850 cause 38 (Network out of order) means the network is not working correctly and the condition is likely to last a relatively long time, so retrying right away is unlikely to help. It is a real failure, not one at the destination. It may be on your link, on the SBC or in the carrier's network. Find out who generated the cause and open a ticket.

Cause 38 in the reference

Cause	38
Official name (ITU-T Q.850)	Network out of order
Class	Resource unavailable (34 to 47)
SIP response per RFC 3398	SIP 503 Service Unavailable

How Monitor treats cause 38

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 38 reads: "the network is out of order".

Cause 38 also decides how a **SIP 480** is read, since on its own a 480 does not say whether the destination rang or nothing answered:

What should answer the number	SIP 480 with cause 38
A greeting or recorded message (the IVR)	Down
A test tone or echo service	Down
Other: it may ring or be answered	Down

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 38 mean in practice?

Q.850 cause 38, Network out of order, means the network is not working correctly and the condition is likely to persist for a relatively long period. The definition itself says that retrying immediately will probably not work. That sets it apart from [cause 41](#), which describes a passing failure. Under RFC 3398 (section 7.2.4.1), cause 38 becomes [SIP 503 Service Unavailable](#); in the other direction, [SIP 502 Bad Gateway](#) maps back to 38. It arrives in the ISUP REL on SS7 links, in the ISDN DISCONNECT, and in a `Reason: Q.850;cause=38` on SIP trunks. In Asterisk, `$_HANGUPCAUSE` is set to 38 and the CDR shows hangupcause 38 when the field is recorded. In LTE, the “cause code 38” in ESM messages comes from a different 3GPP table and has nothing to do with this cause.

When it appears

- A fiber cut or transmission failure in the carrier’s network.
- A carrier SBC or switch down, with no alternative route for the traffic.
- An E1 in alarm on the carrier’s side.
- An intermediate gateway that receives a 502 from another and converts it into 38.
- Your site’s internet link down, when the local SBC responds on behalf of the network.

How to diagnose it

1. See who generated the cause: in `sngrep`, the 503 or 502 response shows the User-Agent or Server of whoever responded.
2. Test several destinations. If they all fail over the same trunk, the failure is in the network, not at the destination.
3. Check the link: packet loss and latency to the carrier’s signaling IP address, with `ping` and `mtr`.
4. Check `pjsip show registrations` and the trunk status in FreePBX, Issabel or 3CX.
5. On an E1, `dahdi show status` shows red alarms (local loss of signal) or yellow alarms (the far end is not receiving your signal).

What to send the carrier

- Date, time and time zone when the failure started, and whether it is still going on.
- The A and B test numbers and the list of destinations that failed.
- The Call-ID, the response received and the Reason with cause=38.
- The `mtr` result to the carrier’s signaling IP address.
- The circuit ID or the trunk identifier, and the E1 alarms, if any.

Frequently asked questions

Should I retry the call when I get cause 38?

Retrying once does no harm, but the definition of cause 38 says the condition is likely to last, so immediate back-to-back attempts tend to fail the same way. If there is an alternative route, such as another trunk or another carrier, it is better to divert traffic to it while the failure is being handled.

What is the difference between cause 38 and cause 27?

Cause 27 (Destination out of order) says the interface to the destination failed, so the problem is close to the receiving end. Cause 38 says the network itself has a fault. Under RFC 3398, 27 becomes SIP 502 and 38 becomes SIP 503. On an 0800 number that terminates on your PBX, 27 points to your side;

38, to the network.

Q.850 CAUSES

Q.850 cause 41 (Temporary failure): what it means

By [2N5 team](#) Updated October 20, 2026

Short answer: Q.850 cause 41 (Temporary failure) means the network is not working correctly, but the condition should be short-lived, so a new attempt right afterward may complete. It is one of the most common causes because SIP 400, 481, 500 and 503 all become 41. Look for the original SIP response to find the real reason.

Cause 41 in the reference

Cause	41
Official name (ITU-T Q.850)	Temporary failure
Class	Resource unavailable (34 to 47)
SIP response per RFC 3398	SIP 503 Service Unavailable

How Monitor treats cause 41

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 41 reads: "a temporary network failure".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 41 mean in practice?

Q.850 cause 41, Temporary failure, means the network is not working correctly and the condition is probably not going to last long: the user can try another call almost immediately. Under RFC 3398 (section 7.2.4.1), cause 41 becomes [SIP 503 Service Unavailable](#). What makes it so frequent is the other direction: the same RFC maps [SIP 400 Bad Request](#), [481 Call/Transaction Does Not Exist](#), [500 Server Internal Error](#) and 503 Service Unavailable to 41. So a cause 41 on your PBX may originally be a message format error or an internal failure on a SIP server. It arrives in the ISUP REL, the ISDN DISCONNECT and a Reason: `Q.850;cause=41`. In Asterisk, `${HANGUPCAUSE}` is set to 41, and the CDR shows hangupcause 41 when it records the field.

When it appears

- A restart or momentary overload of a carrier SBC or proxy.
- A 500 error on some SIP server along the path, translated into 41 by a gateway.
- A malformed SIP message (400) generated by your PBX, for example with an invalid header.
- A carrier server switchover during maintenance.
- A brief flap of the link between your site and the carrier.

How to diagnose it

1. Try again in a few seconds. If it completes, note the time and keep an eye on the frequency.

- 2. Find the original SIP response in `sngrep` or with `pjsip set logger on: 400, 481, 500 and 503` lead to different investigations.
- 3. If it is a 400, read the status line and the Warning: many carriers say there which header they rejected.
- 4. If it is a 503, see whether there is a `Retry-After` and respect the interval.
- 5. In 3CX, turn on Verbose logging to record the full response in the Activity Log.
- 6. If the failures repeat at the same time of day, cross-check them with the carrier's maintenance windows.

What to send the carrier

- Date, time and time zone of each failure, and the frequency (for example, how many per hour).
- The A and B numbers.
- The Call-ID and the full SIP status line received, with the Reason and the Warning.
- Whether a new attempt right afterward completed.
- The pcap of the call that failed.

Frequently asked questions

If cause 41 is temporary, do I need to open a ticket?

For an isolated occurrence that completes on the next attempt, usually not. If cause 41 repeats several times a day, or always on the same destination, it is no longer passing in practice, and it is worth opening a ticket with the times, the Call-IDs and the original SIP response for each failure.

Does a real test call tell a passing failure apart from one that lasts?

A real test call gets the network's final SIP response, like any call. When it fails, the attempt is retried before it becomes an alert. So a failure that disappears on retry raises no alarm, and one that persists becomes an alert sent by email.

Q.850 CAUSES

Q.850 cause 47 (Resource unavailable, unspecified): what it means

By 2N5 team Updated October 7, 2026

Short answer: Q.850 cause 47 (Resource unavailable, unspecified) means some resource needed for the call was unavailable, and no other cause in the resource class applies. It is a generic shortage: media channels, transcoding, licenses or a gateway's capacity. Find out which device generated the cause and check its limits.

Cause 47 in the reference

Cause	47
Official name (ITU-T Q.850)	Resource unavailable, unspecified
Class	Resource unavailable (34 to 47)
SIP response per RFC 3398	SIP 503 Service Unavailable

How Monitor treats cause 47

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier.

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 47 mean in practice?

Q.850 cause 47, Resource unavailable, unspecified, is used to report a resource unavailable event when no other cause in that class applies. It is the generic cause in the group that includes 34 (no circuit), 38 (network out of order), 41 (temporary failure) and 42 (congestion). Under RFC 3398 (section 7.2.4.1), cause 47 becomes `SIP 503 Service Unavailable`, so only the `Reason: Q.850;cause=47` header (RFC 6432) or the TDM trace tells it apart from the others. It arrives in the ISUP REL on SS7 links and in the ISDN DISCONNECT. In Asterisk, `$_HANGUPCAUSE` is set to 47, and the CDR shows hangupcause 47 when the field is recorded. Because it is unspecific, the most useful step is finding out which device generated it.

When it appears

- A gateway or SBC without the DSP resources to transcode, for example from G.729 to G.711.
- A media server or IVR with no free ports to play the announcement.
- Concurrent call licenses used up on a gateway or SBC.
- A lack of memory or free RTP ports on an overloaded device.
- A GSM or E1 gateway that uses 47 for any unavailable port.

How to diagnose it

1. In `sngrep`, see who answered with the 503 carrying cause=47, from the User-Agent or the Server.
2. If it was your equipment, check transcoding: in Asterisk, `core show translation` shows the paths between codecs and `core show channels` the active calls.
3. Compare the codecs offered in the SDP with those accepted on the other side; using the same codec at both ends avoids transcoding.
4. Check the license limits and the RTP port range (`rtpstart` and `rtpend` in `rtp.conf`).
5. Correlate the failures with peaks in concurrent calls.

What to send the carrier

- Date, time and time zone of the failures.
- The A and B numbers.
- The Call-ID, the `SIP/2.0 503 Service Unavailable` line and the Reason with cause=47.
- The codecs offered in your INVITE (the `m=` and `a=rtpmap` lines of the SDP).
- How many concurrent calls there were on your side at the time.

Frequently asked questions

What is the difference between cause 34 and cause 47?

Cause 34 says specifically that a circuit or channel was missing. Cause 47 says some resource was missing, without saying which. Both become SIP 503 under RFC 3398. When the device knows the problem is a channel, it should use 34; 47 appears when the limit is something else, such as transcoding or a license.

Can cause 47 be a codec problem?

Indirectly, yes. If there is no codec in common and the gateway has to transcode, it spends DSP or CPU on that. When that resource runs out, the call can fail with cause 47. When there is no acceptable codec at all, the expected cause is a different one, such as 65, which becomes SIP 488 Not Acceptable Here.

Q.850 CAUSES

Q.850 cause 65 (Bearer capability not implemented): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 65 (Bearer capability not implemented) means the equipment that cleared the call does not support the requested bearer capability, such as a data call or a media format. On SIP trunks, it is almost always a missing common codec or T.38 fax not being accepted. Compare the SDP in the INVITE with what the carrier accepts.

Cause 65 in the reference

Cause	65
Official name (ITU-T Q.850)	Bearer capability not implemented
Class	Service or option not implemented (65 to 79)
SIP response per RFC 3398	SIP 488 Not Acceptable Here

How Monitor treats cause 65

When the carrier gives the Q.850 cause in the **Reason** header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier.

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 65 mean in practice?

Q.850 cause 65, Bearer capability not implemented, means the equipment sending the cause does not implement the requested bearer capability. In TDM telephony, the bearer says whether the call is speech, 3.1 kHz audio or unrestricted digital data. In SIP, that role falls to the media description in the SDP. Under RFC 3398 (section 7.2.4.1), cause 65 becomes [SIP 488 Not Acceptable Here](#). In the other direction, a 488 or a [606 Not Acceptable](#) whose Warning is about bearer capability maps back to 65; without that Warning, the default is 31. The cause appears in the ISUP REL, the ISDN DISCONNECT and a **Reason: Q.850; cause=65**. In Asterisk, `_${HANGUPCAUSE}` is set to 65, and the CDR shows hangupcause 65 when it records the field.

When it appears

- No codec in common: the PBX offers only G.729 and the carrier accepts only G.711 A-law (PCMA), the standard in Brazil.
- A re-INVITE for T.38 fax on a trunk that does not accept T.38.
- An ISDN data call (64 kbit/s unrestricted) sent to a voice-only destination.
- Video or other media offered on an audio-only trunk, without a compatible audio line.
- An older gateway that does not implement the requested bearer format.

How to diagnose it

1. In `sngrep`, open the INVITE and look at the `m=audio` and `a=rtpmap` lines in the SDP.
2. Read the 488's Warning: codes 304 (Media type not available) and 305 (Incompatible media format) point to the media.
3. In Asterisk, check the endpoint's `allow` and `disallow` with `pjsip show endpoint trunk-name` and make sure `alaw` is on the list.
4. In 3CX, check the codec order configured on the trunk and compare it with the carrier's list.
5. For fax, test with T.38 turned off (G.711 passthrough) or confirm with the carrier whether T.38 is supported.
6. On an ISDN E1, `pri set debug on span 1` shows the Bearer Capability sent in the SETUP.

What to send the carrier

- Date, time and time zone of the attempts.
- The A and B numbers and the type of call (voice, fax, data).
- The Call-ID, the `SIP/2.0 488 Not Acceptable Here` line, the Warning and the Reason.
- The full SDP of the rejected INVITE or re-INVITE.
- An explicit request for the codec list and for T.38 support on the trunk.

Frequently asked questions

Is cause 65 the same as SIP 488?

Almost. Under RFC 3398, cause 65 becomes SIP 488 Not Acceptable Here, but a 488 only maps back to cause 65 when the Warning is about bearer capability; without that, it becomes cause 31. So a 488 can reach the telephone network as 65 or as 31, depending on what the rejecting equipment wrote in the Warning.

Why does fax fail with cause 65 while voice works?

Because fax usually switches media mid-call, with a re-INVITE asking for T.38. If the trunk or gateway on the other side does not accept T.38, it rejects the new media, and the cause can be 65. Voice in G.711 keeps working because it uses a format both sides accept.

Q.850 CAUSES

Q.850 cause 88 (Incompatible destination): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 88 (Incompatible destination) means the destination received a call request with compatibility attributes it cannot handle, such as the terminal type or the data rate. It is common when voice, fax and data get mixed up. Check the type of call sent, the bearer requested, and what kind of terminal is connected to the number.

Cause 88 in the reference

Cause	88
Official name (ITU-T Q.850)	Incompatible destination
Class	Invalid message (81 to 95)
SIP response per RFC 3398	SIP 503 Service Unavailable

How Monitor treats cause 88

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier.

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 88 mean in practice?

Q.850 cause 88, Incompatible destination, means the destination equipment received a call request with low layer compatibility, high layer compatibility or other attributes, such as the data rate, that it cannot accommodate. Unlike [cause 65](#), which is about a bearer that is not implemented, 88 points to a mismatch between what the call asks for and the type of terminal at the destination. It belongs to the Q.850 invalid message class. Under RFC 3398 (section 7.2.4.1), cause 88 becomes [SIP 503 Service Unavailable](#), so on the SIP side it gets mixed up with capacity failures. It arrives in the ISUP REL, the ISDN DISCONNECT and a `Reason: Q.850;cause=88` (RFC 6432). In Asterisk, `$_HANGUPCAUSE` is set to 88, and the CDR shows hangupcause 88 when it records the field.

When it appears

- An ISDN data call (64 kbit/s unrestricted) to an analog phone or a voice-only destination.
- A voice call to a line configured for fax or data only.
- An ISDN gateway that passes on a high layer compatibility (HLC) the destination terminal does not accept.
- Older ISDN terminals with a specific configuration, such as group 4 fax or a videophone.
- SIP-ISDN gateways that build the SETUP with a different bearer from the one the destination expects.

How to diagnose it

1. On an ISDN E1, turn on `pri set debug on span 1` and read the Bearer Capability, Low Layer Compatibility and High Layer Compatibility fields in the SETUP.
2. Confirm with the number's owner what type of terminal is connected to it.
3. On the SIP-ISDN gateway, check how it fills in the bearer: many have an option to force speech or 3.1 kHz audio.

4. In Asterisk with DAHDI, the `TRANSFERCAPABILITY` channel variable sets the bearer for outbound calls; confirm the value in use.
5. In `sngrep`, see whether the 503 came with `Reason: Q.850;cause=88`, to tell it apart from a 503 caused by a lack of capacity.

What to send the carrier

- Date, time and time zone of the attempts.
- The A and B numbers and the intended type of call (voice, fax, data).
- The Call-ID, the `SIP/2.0 503 Service Unavailable` response and the Reason with cause=88.
- The trace of the SETUP with the Bearer Capability, if the call goes out over an ISDN E1.
- A request to confirm the type of terminal provisioned on the destination number.

Frequently asked questions

Why does cause 88 become SIP 503 and not 488?

Because that is what the RFC 3398 table defines: cause 88 becomes SIP 503 Service Unavailable, while cause 65 becomes 488. In practice, this makes a compatibility problem look like a capacity problem on the SIP side. Only the Reason header with cause=88, or the TDM trace, shows the real reason.

Does cause 88 appear on pure SIP trunks?

Rarely. It originates in networks with ISDN or ISUP signaling, where the call carries compatibility attributes. On a SIP trunk, you usually see it when the call crosses a gateway into the TDM network and the destination there rejects the attributes. If the path is SIP end to end, the usual result is a 488 with a Warning.

Q.850 CAUSES

Q.850 cause 102 (Recovery on timer expiry): what it means

By [2N5 team](#) Updated October 20, 2026

Short answer: Q.850 cause 102 (Recovery on timer expiry) means an error recovery procedure started because a protocol timer expired: some expected message did not arrive in time. On SIP trunks it is usually an unanswered INVITE, caused by a firewall, NAT, SIP ALG or a UDP packet that is too large. Follow the signaling and see where it stops.

Cause 102 in the reference

Cause	102
Official name (ITU-T Q.850)	Recovery on timer expiry
Class	Protocol error (96 to 111)
SIP response per RFC 3398	SIP 504 Server Time-out

How Monitor treats cause 102

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier. On the English screens, cause 102 reads: "a protocol timer expired".

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 102 mean in practice?

Q.850 cause 102, Recovery on timer expiry, means a procedure was started by the expiry of a timer, as part of the error handling procedures. In other words, a device waited for a message that never came and gave up. It belongs to the protocol error class. Under RFC 3398 (section 7.2.4.1), cause 102 becomes [SIP 504 Server Time-out](#) (the RFC writes "504 Gateway timeout"); in the other direction, [SIP 408 Request Timeout](#) and 504 map back to 102. In SIP, the typical timer is the INVITE's Timer B, at 64 times T1, or 32 seconds with the default T1 of 500 ms (RFC 3261). The cause arrives in the ISUP REL and in a Reason: Q.850;cause=102. In Asterisk, `$_HANGUPCAUSE` is set to 102, and the CDR shows hangupcause 102 when it records the field.

When it appears

- A firewall blocking the carrier's response or the outbound INVITE itself.
- NAT with no mapping for the SIP port, or the router's SIP ALG altering the messages.
- An INVITE with an SDP that is too large (many codecs), which exceeds the MTU, gets fragmented over UDP, and loses a fragment.
- The trunk's DNS resolving to an old carrier IP address.
- An intermediate SIP server that does not respond and does not even send 100 Trying.

How to diagnose it

1. In `sngrep`, see whether the INVITE was retransmitted several times with no response: that confirms nothing came back.
2. Capture on the external interface with `tcpdump -i any -s 0 -w call.pcap port 5060` to find out whether the INVITE left and whether any response arrived.
3. Disable the router's SIP ALG and check the NAT and firewall rules for the carrier's IP addresses.
4. If the INVITE is larger than about 1,300 bytes, reduce the codecs offered or use TCP on the trunk.
5. Check the DNS for the trunk's domain and the IP address the INVITE is actually going to.
6. In 3CX, the Activity Log with Verbose logging shows whether the trunk received any response before the timeout.

What to send the carrier

- Date, time and time zone of the attempts.
- The A and B numbers.
- The Call-ID and what you received: no response, `SIP/2.0 408 Request Timeout` or 504, with the Reason.
- The source public IP address and the destination IP address the INVITE was sent to.
- The pcap showing the INVITE retransmissions.

Frequently asked questions

Why does the call take about 32 seconds before failing with cause 102?

Because a SIP INVITE over UDP is retransmitted until Timer B expires, and Timer B is 64 times T1. With the default T1 of 500 ms, that comes to 32 seconds. If no final response arrives within that interval, the transaction fails, and the gateway reports cause 102 to the telephony side.

Does a real test call help with timeout failures?

Yes, because it follows the same signaling path as an ordinary call and gets the final SIP response, or the lack of one. A failure is retried before it becomes an alert, which avoids an alarm over a single lost packet, and the alert goes out by email.

Q.850 CAUSES

Q.850 cause 127 (Interworking, unspecified): what it means

By [2N5 team](#) Updated October 7, 2026

Short answer: Q.850 cause 127 (Interworking, unspecified) means the call passed through an interconnection with a network that does not provide causes for its actions, so the exact reason is unknown. It does not point to a specific defect. It is common on gateways between SIP, ISDN and R2 signaling. Find out on which leg the call failed and which response came from there.

Cause 127 in the reference

Cause	127
Official name (ITU-T Q.850)	Interworking, unspecified
Class	Interworking (127)
SIP response per RFC 3398	SIP 500 Server Internal Error

How Monitor treats cause 127

When the carrier gives the Q.850 cause in the `Reason` header of the SIP response, Monitor keeps it beside the SIP code on every test call, and it appears in the incident's evidence, ready for the ticket with the carrier.

Like any failure, a test call read as Down is repeated 3 times, every 2 minutes, and only 4 failures in a row open an incident.

What does cause 127 mean in practice?

Q.850 cause 127, Interworking, unspecified, means there was interworking with a network that does not provide causes for the actions it takes, so the exact cause of the message is not known. It is the only cause in the interworking class in Q.850. Under RFC 3398 (section 7.2.4.1), cause 127 becomes [SIP 500 Server Internal Error](#). In the other direction, several SIP responses with no telephony equivalent map back to 127: 413, 414, 416, 420, 421, 423, 505 and 513. In other words, problems with SIP extensions (the `Require` header), message size or version can reach the telephone network as 127. The cause appears in the ISUP REL, in a `Reason: Q.850;cause=127` (RFC 6432) and, in Asterisk, in `${HANGUPCAUSE}`, as well as the CDR's `hangupcause` field when it is recorded.

When it appears

- A call that crosses a leg with R2 MFC signaling, still common on E1 in Brazil, which does not carry Q.850 causes.
- A gateway that receives SIP 420 Bad Extension or 421 Extension Required because of `Require: 100rel` or `Require: timer`.
- A SIP message that is too large, answered with 513 Message Too Large.
- A Request-URI that is too long (414) or an unsupported URI scheme (416).
- Equipment from different vendors in sequence, each one translating the previous one's cause.

How to diagnose it

1. Follow the call leg by leg in `sngrep` or in the pcap and find the first error response.
2. If the first response is 420 or 421, check the INVITE's Require and Supported headers; in PJSIP, test with `100rel` and `timers` turned off on the endpoint.
3. If it is 513, reduce the size of the INVITE (fewer codecs or headers) or use TCP.
4. On an E1 with R2, check on the gateway which B signal the destination returned and how the gateway translated it.
5. In 3CX, the Activity Log with Verbose logging shows the trunk's exact response before translation.
6. Ask whoever is responsible for each gateway for the cause or response it received before generating the 127.

What to send the carrier

- Date, time and time zone of the attempts.
- The A and B numbers.
- The Call-ID, the `SIP/2.0 500 Server Internal Error` line and the Reason with cause=127.
- The full INVITE, with Require, Supported and the message size.
- An explicit request for the original response or signal on the leg where the carrier did the interconnection.

Frequently asked questions

Does cause 127 mean the carrier has an internal error?

Not necessarily. Under RFC 3398, cause 127 becomes SIP 500 Server Internal Error, which looks like an internal failure, but the cause only says the call passed through a network that did not report the reason. The real error can be on any leg, including a rejected SIP extension or a message that is too large.

Why do I see cause 127 on E1 trunks with R2?

Because R2 MFC signaling does not carry Q.850 causes: it uses B signals, such as subscriber free, busy or congestion. The gateway that connects R2 to ISDN or SIP has to translate those signals, and when there is no clear translation, or when the destination simply releases the call, it may use cause 127.

Glossary

Short definitions and examples of telephony and voice over IP terms: MOS, MOS-LQO, jitter, SIP ALG, IVR, DID, 0800, SIP trunk, SLA and more.

GLOSSARY

0300 and 0303 numbers: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: 0300 and 0303 numbers are Brazilian nationwide shared-cost access codes, where the cost of the call is split between the caller and the number's subscriber, under Anatel's numbering plan.

In Portuguese: Números 0300 e 0303

What are 0300 and 0303 numbers?

They are nationwide shared-cost access codes: the cost of the call is split between the caller and the number's subscriber. The rule is in Anatel Resolution No. 749/2022 (Resolução Anatel nº 749/2022), art. 18, the same rule that makes 0800 free for the caller.

The two serve different purposes. 0300 is a number for receiving calls: customer service, support and information lines, where a company wants a single nationwide number without bearing the whole cost as it would with 0800. 0303 identifies high-volume callers. Under Anatel Act No. 12,712/2024 (Ato Anatel nº 12.712/2024), it is mandatory for anyone who makes more than 10,000 calls on at least one day of the month, as in telemarketing, debt collection and fundraising. The number appears on the screen of the person receiving the call.

Like 0800, a 0300 number is translated by the carrier's network to a real destination, usually one of the company's trunks, which delivers the call to the PBX and the IVR. Routing at the carrier is a point of failure that PBX monitoring does not see.

In practice

A comparison of the most common access numbers:

Number	Who pays	Typical use
0800	Subscriber	SAC (the customer-service channel Brazilian law regulates), ombudsman, support
0300	Shared cost	Customer service on a nationwide number
0303	Shared cost	Mass calling: telemarketing, debt collection, fundraising
4004, 4003, 3003	The caller, usually as a local call, according to the carriers	Single nationwide number

The 4004, 4003 and 3003 numbers are not a category in Anatel's numbering plan; see 4004, 4003 and 3003 numbers. Choosing between them comes down to who the company wants to bear the cost of the call and whether the number will receive calls or make them.

How Monitor uses this

Any dialable number can be monitored, including 0300 and 0303. The test dials the number itself, just as a customer would, and goes through the carrier's translation, the trunk and the IVR, testing the full path the customer uses. See [toll-free and IVR monitoring](#).

Frequently asked questions

Who is required to use 0303?

Under Anatel Act No. 12,712/2024 (Ato Anatel nº 12.712/2024), anyone who makes more than 10,000 calls on at least one day of the month must use the 0303 code. That covers activities such as telemarketing, debt collection and fundraising. With 0303 on the screen, the person receiving the call knows up front what kind of call it is and can decide whether to pick up.

What is the difference between 0300 and 0800?

Who pays. With 0800 (toll-free), the call is free for the caller and the subscriber pays everything. With 0300, the cost is shared between the caller and the subscriber, under Anatel Resolution No. 749/2022 (Resolução Anatel nº 749/2022), art. 18. Both are nationwide numbers that the carrier routes to the company's real destination, and both can fail in routing without the PBX noticing.

GLOSSARY

0800 number: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: An 0800 number is a non-geographic access code where the call is free for the caller and the number's subscriber pays for it, under Anatel's numbering plan.

In Portuguese: Número 0800

What is an 0800 number?

An 0800 number is the access code where the caller pays nothing and the number's subscriber pays for the call. The rule is in Anatel Resolution No. 749/2022 (Resolução Anatel nº 749/2022), art. 18. It is usually written as 0800 123 4567 and is used for SAC (the customer-service channel Brazilian law regulates), ombudsman lines, technical support and sales.

An 0800 number is not geographic: it does not belong to a city or to a physical line. When someone dials it, the carrier's network translates the 0800 to a real destination, usually one of the company's [SIP trunks](#) or E1 trunks, which delivers the call to the PBX and the [IVR](#).

That translation creates a point of failure that internal monitoring cannot see. The trunk can be registered, the PBX up and the IVR working, and the 0800 can still fail because the routing at the carrier was changed, the number was deactivated by mistake or the configured destination is wrong. The caller hears a carrier message, and the company only notices when queue volume drops or complaints arrive through other channels.

Because it is free, 0800 is also the number customers call when something has already gone wrong. An outage on it tends to hit exactly the people who most need help.

In practice

There are two ways to test an 0800, and they do not measure the same thing:

Test	Does the carrier translate the 0800?	Trunk and PBX	IVR
Call the destination number directly	No	Yes	Yes
Call the 0800	Yes	Yes	Yes

Example: 0800 123 4567 points to 11 3456-7000. If the carrier mistakenly changes the destination to 11 3456-7100, the test on the direct number keeps passing, and only a call to the 0800 reveals the failure. On the company's side nothing changes: the trunk stays registered and the IVR answers whoever manages to get through.

How Monitor uses this

Any dialable number can be monitored, including 0800. The test dials the 0800 itself, just as a customer would, and so goes through the carrier's translation, the trunk, the PBX and the IVR, testing the full path the customer uses. See [toll-free and IVR monitoring](#).

Frequently asked questions

Who pays for a call to an 0800 number?

The number's subscriber, meaning the company that contracted the 0800. Under Anatel Resolution No. 749/2022 (Resolução Anatel nº 749/2022), art. 18, a call to 0800 is free for the caller. That is why 0800 is common for SAC (the customer-service channel Brazilian law regulates), ombudsman lines and support. The company's cost grows with call volume and duration, one more reason to make sure every call gets answered.

Why does my 0800 work from one phone and fail from another?

Because 0800 is routed by the carrier's network, and that routing can vary with where the call comes from, such as the caller's network or region. A misconfiguration can affect only some of the calls. It can also happen that the destination accepts the call from one network and rejects it from another because of a codec or a channel limit.

GLOSSARY

4004, 4003 and 3003 numbers: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: 4004, 4003 and 3003 numbers are eight-digit fixed-line access codes that Brazilian carriers sell as a single nationwide number, dialed without an area code in the locations the carrier serves.

In Portuguese: Números 4004, 4003 e 3003

What are 4004, 4003 and 3003 numbers?

They are eight-digit fixed-line numbers, in the format 4004-1234, that Brazilian carriers sell as a single nationwide number. The company advertises one number, and callers dial it without an area code in the locations the carrier serves. Unlike [0800](#) and [0300](#), they are not a category in Anatel's numbering plan: they are a carrier product. According to the carriers, the caller usually pays the price of a local call.

Behind the number, the carrier configures the same code in each local area where it offers the service and forwards every call to the company's real destination, usually a [SIP trunk](#) or E1 trunk that delivers the call to the PBX and the [IVR](#).

That design has a consequence for anyone monitoring it. There is one number, but routing is configured per location. A wrong change can make the 4004 number work for callers in São Paulo and fail for callers in Recife. There can also be differences between calls from landlines and from mobile phones, depending on the dialing format the carrier accepts in each case.

In practice

Three ways to publish a customer service number:

Number	Who pays	Format	Reach
0800 123 4567	Subscriber	Non-geographic code	Nationwide
4004-1234	The caller, usually as a local call	Eight digits, no area code	The carrier's locations
(11) 3456-7000	The caller, depending on distance	Fixed-line number with area code	One local area

Example: a company moves its 4004 destination from an old trunk to a new one. If the carrier updates 20 locations and forgets one, calls from that city fail while the rest work.

How Monitor uses this

Any dialable number can be monitored, including 4004, 4003 and 3003. The test dials the number itself, just as a customer would, and goes through the carrier's routing, the trunk and the IVR, testing the full path the customer uses. See [toll-free and IVR monitoring](#).

Frequently asked questions

Is 4004 the same as 0800?

No. 0800 is a category in Anatel's numbering plan and is free for the caller. The 4004, 4003 and 3003 numbers are not an Anatel category: they are eight-digit fixed-line numbers that carriers sell as a single nationwide number. Instead of the company paying for the call, the caller usually pays for a local call, as the carriers describe the service.

Is calling a 4004 number free?

Generally not. Carriers describe the service as a number where the caller pays the price of a local call, even from another city. The exact cost depends on the caller's plan and where the call comes from. For free customer service, the right number is 0800 (toll-free), paid by the subscriber.

ASR and ACD: definition and a practical example

By [2N5 team](#) Updated October 14, 2026

Short answer: ASR (Answer Seizure Ratio) is the percentage of call attempts that were answered, and ACD (Average Call Duration) is the average length of the answered calls. Both come from the carrier's or PBX's call detail records (CDR).

In Portuguese: ASR e ACD

What are ASR and ACD?

ASR and ACD are the two oldest quality indicators of a voice route, calculated from call detail records, the CDRs.

ASR (Answer Seizure Ratio, from ITU-T Recommendation E.411) is the percentage of call attempts that were answered: $\text{answered calls} \div \text{attempted calls} \times 100$. **ACD** (Average Call Duration) is the average length of the answered calls: $\text{total talk time} \div \text{number of answered calls}$.

Carriers define ASR in slightly different ways. Some leave out of the calculation the calls the destination rejected, such as busy, and the difference changes the number. Check the definition in the contract before comparing routes.

What a change indicates

- **ASR dropping sharply:** the route may have stopped completing calls, wholly or in part. It is the cue to look at the response codes, such as [SIP 503](#).
- **A very short ACD:** calls are answered and drop soon after, which usually points to silent audio, [SIP ALG](#) or a codec problem.
- **A very long ACD:** stuck calls, for example a call that did not hang up at the far end.

Both look at the whole, not at a single call. They hide the failure of the 0800 that nobody called overnight, and they do not say what the person who called heard.

In practice

In one day, a [SIP trunk](#) made 1,000 call attempts and 600 were answered, with 30,000 seconds of talk time in total (illustrative numbers).

Indicator	Calculation	Result
ASR	$600 \div 1,000 \times 100$	60%
ACD	$30,000 \text{ s} \div 600$	50 s

If tomorrow the ASR falls to 20% with no change in traffic, the route has a problem. If it falls to 20% because a campaign dialed a poor list, the trunk is fine.

How Monitor uses this

Monitor does not calculate ASR or ACD, because it does not analyze your customers' traffic. It works the other way round: it calls a number you control, from the outside, and reports whether the call connected, with the SIP code when it does not. Use the ASR in your CDRs to see that something changed and the test call to know whether the circuit still completes calls. See [synthetic monitoring](#), [post dial delay](#) and [SIP trunk and VoIP monitoring](#).

Frequently asked questions

What is a good ASR?

There is no good number for everyone. ASR depends on the kind of traffic: calls to customers who do not pick up lower the figure with no network fault at all. Compare a route's ASR with its own on similar days, and investigate when it drops suddenly, because the sharp drop is the useful signal.

Does a low ASR mean the trunk has a problem?

Not always. A low ASR can come from a trunk fault, but also from a busy called number, a campaign to a poor list, or a cheap and unstable route. ASR tells you something changed; a test call to a number you control shows whether the trunk still completes calls.

GLOSSARY

DID: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: DID (Direct Inward Dialing, called DDR in Brazil) is the service that maps public numbers to the extensions of a PBX, so outside callers reach an extension directly without going through an operator or an IVR.

In Portuguese: DDR

What is DID?

DID (Direct Inward Dialing, known in Brazil as DDR) is the carrier service that gives a company a range of public numbers, such as (11) 3456-7000 to 7099, and lets each number reach an extension on the [IP PBX](#) directly. Someone who calls (11) 3456-7012 talks to extension 7012 without going through the front desk or the [IVR](#).

An extension is the internal number of each phone, softphone or group inside the PBX. The DID is the bridge between the public number and that extension.

The DID range reaches the company over a [SIP trunk](#) or an [E1 trunk](#). On each call, the carrier states which number was dialed: in SIP, in the Request-URI and the To header; on E1, in the called (B) number. The PBX reads that number, usually the last four digits, and looks up the matching extension in the dial plan.

The most common failure lies in that agreement on digits. If the carrier changes the format it sends, the PBX stops recognizing the number: the call falls to the default rule, lands on the operator or is rejected. The trunk stays up, the extension stays registered, and only a call from outside reveals the problem.

In practice

A 100-number DID range with the PBX reading four digits:

Number dialed	Digits received	Destination on the PBX
(11) 3456-7000	7000	Front desk IVR
(11) 3456-7012	7012	Extension 7012, finance

Number dialed	Digits received	Destination on the PBX
(11) 3456-7050	7050	Support group
(11) 3456-7012, after a change at the carrier	+551134567012	Default rule: operator

In the last row, nothing changed inside the company, yet finance stopped receiving direct calls. Internal calls to 7012 keep working, which delays noticing the problem. That is why the test has to come from outside, through the public number.

How Monitor uses this

Any dialable number can be monitored, including extensions reachable by DID. The test dials the extension's public number, just as a customer would, and goes through the carrier, the trunk and the PBX dial plan. See [SIP trunk monitoring](#).

Frequently asked questions

What is the difference between a DID and an extension?

An extension is the internal number of a phone inside the PBX, such as 7012, dialed between colleagues. A DID is the public number that leads to that extension, such as (11) 3456-7012, dialed by people outside. Every DID points to an extension, a group or an IVR, but not every extension has a DID: many only receive calls transferred by the front desk.

Why does a call to a DID land on the operator?

Usually because the PBX did not recognize the number it received and fell back to the default inbound rule. That happens when the carrier starts sending a different number of digits, for example the full number with +55 instead of the last four, or when the extension was removed without updating the dial plan. The PBX call log shows the number that arrived.

GLOSSARY

DTMF: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: DTMF (Dual-Tone Multi-Frequency) is the system that represents each phone key as the sum of two frequencies, used for dialing and for choosing options in IVRs.

In Portuguese: DTMF

What is DTMF?

DTMF (Dual-Tone Multi-Frequency) is the system of phone keypad tones. Each key generates two frequencies at the same time, one from a low group and one from a high group, and the receiver identifies the key from the combination. DTMF is how people navigate an [IVR](#), dial an extension or confirm a PIN over the phone.

	1209 Hz	1336 Hz	1477 Hz	1633 Hz
697 Hz	1	2	3	A
770 Hz	4	5	6	B

	1209 Hz	1336 Hz	1477 Hz	1633 Hz
852 Hz	7	8	9	C
941 Hz	*	0	#	D

In traditional telephony, the tones traveled inside the audio itself. In VoIP there are three ways to carry them:

- In the audio (in-band): the tone travels as sound inside [RTP](#). It works well with G.711, but compressed [codecs](#) such as G.729 distort the tone and the key gets lost.
- RFC 4733 (telephone-event): each key becomes its own event in the RTP stream, with a payload type agreed in the SDP. It replaces RFC 2833 and is the most common method on [SIP](#) trunks.
- SIP INFO: the key travels as a SIP message, outside the audio. It depends on every device in the path passing the message along.

In practice

Pressing key 5 makes the phone generate 770 Hz and 1336 Hz together. A customer calls the company's 0800 (toll-free) number and the IVR says "for billing, press 2". What happens depends on the method:

Phone sends	IVR expects	Result
RFC 4733	RFC 4733	Key 2 recognized
In the audio, G.711	In the audio	Key 2 recognized
In the audio, G.729	In the audio	Tone distorted, key ignored
RFC 4733	In the audio	Key never arrives

How Monitor uses this

On test-tone or echo circuits, Monitor sends test tones and checks that they come back. On IVRs, the test call listens for the expected greeting. See [toll-free and IVR monitoring](#).

Frequently asked questions

Why doesn't the IVR recognize the keys I press?

Almost always because the two ends use different DTMF methods. If the phone sends RFC 4733 and the IVR expects tones inside the audio, or the other way around, the digits get lost. Compressed codecs such as G.729 also distort tones sent in the audio. Check in the SDP whether telephone-event was negotiated and use the same method on both ends.

What is the best DTMF method for VoIP?

On most SIP trunks, RFC 4733, which sends each key as an event in the RTP stream, with a payload type negotiated in the SDP. It does not depend on the codec and arrives in step with the audio. In-band DTMF is only reliable with G.711, and SIP INFO requires every device in the path to pass it along.

GLOSSARY

E1 trunk: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: An E1 trunk is a 2.048 Mbit/s digital circuit between the carrier and the PBX that carries up to 30 simultaneous voice channels, with ISDN or R2 signaling.

In Portuguese: Tronco E1

What is an E1 trunk?

An E1 trunk is a dedicated digital connection between the carrier and the company's [PBX](#), at a rate of 2.048 Mbit/s. It is divided into 32 time slots of 64 kbit/s: slot 0 handles synchronization, slot 16 carries signaling and the remaining 30 carry one voice channel each. That is why an E1 holds up to 30 simultaneous calls.

There are two kinds of signaling. With ISDN (PRI), slot 16 carries messages such as SETUP, CONNECT and RELEASE, and the reason each call ended comes as a Q.850 cause in the release message, such as 17 (busy) or 34 (no circuit available). With R2, widely used in Brazil, signaling is done with line and register signals, which carry less information about why a call failed.

E1 predates IP telephony but is still present in many companies, often connected to an IP PBX through a gateway. ISDN's Q.850 causes are the same ones [SIP](#) carries in the Reason header, which makes it easier to compare failures across the two trunk types.

In practice

E1 compared with a [SIP trunk](#):

Characteristic	E1	SIP trunk
Medium	Dedicated physical circuit	IP network or internet
Capacity	30 channels per E1	Set by contract
Growth	In blocks of 30 channels	By changing the contract
Signaling	ISDN or R2	SIP
Failure reason	Q.850 cause (ISDN)	SIP response and Q.850 cause
Weak points	Synchronization, cabling, equipment	Link, NAT, SIP ALG, codecs

With all 30 channels busy, the next incoming call gets busy, with Q.850 cause 34 on ISDN. On R2, the same situation shows up as a busy signal, without the detailed cause.

How Monitor uses this

Monitor calls the number, so it works the same way whether the number is on an E1 or a SIP trunk. The test does not depend on the trunk type: it shows whether the call completed and why it failed. See [SIP trunk monitoring](#).

Frequently asked questions

E1 or SIP: which should I choose?

E1 has a fixed capacity of 30 channels, a dedicated physical circuit and predictable quality, but it requires installation and grows in blocks of 30 channels. A SIP trunk uses the data network, adds channels by contract and usually costs less, but it depends on the link, the firewall and the NAT configuration. Many companies move to SIP and keep the E1 as a backup.

How many simultaneous calls fit on an E1?

Up to 30. An E1 has 32 time slots of 64 kbit/s: slot 0 is used for synchronization, slot 16 carries signaling and the other 30 carry voice. If a company needs 45 simultaneous calls, it needs two E1s, or one E1 plus a SIP trunk. The 31st call gets a busy signal.

GLOSSARY

Early media: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: Early media is audio sent to the caller before the call is answered, usually after a SIP 183 Session Progress response, such as ringback tone or a carrier message.

In Portuguese: Early media

What is early media?

Early media is the audio the caller hears before the call is answered. In SIP, it usually starts with a 183 Session Progress response carrying an SDP: the network opens the RTP stream before the 200 OK and plays something to the caller. The model is described in RFC 3960.

The content can be ringback generated by the carrier, hold music, a notice such as “this call may be recorded” or an error message, such as “the number you have called does not exist” or “this phone is switched off or out of the coverage area”. To the human ear, all of it is sound. To the signaling, the call has not been answered yet.

That sets a trap for any automated test. If the test treats “I heard audio” as success, a carrier message saying the number is out of service passes as a good call. What defines the answer is the 200 OK, and what defines whether the right destination responded is the content of the audio.

In practice

A call to a number whose trunk is down might play out like this:

Time	Message	What the caller hears
0.0 s	INVITE sent	Nothing
0.2 s	100 Trying	Nothing
1.1 s	183 Session Progress with SDP	“The number you have called is not available”
9.0 s	480 Temporarily Unavailable	Call ends

There was never a 200 OK. A test that only looks for audio would see 8 seconds of speech and mark the line as working. The final response, 480, tells the real story. Some carriers even answer with 200 OK just to play the message, and then only the content of the audio reveals the failure.

How Monitor uses this

For a greeting circuit, a recorded message in place of ringback is usually the carrier speaking for a number it could not reach, and Monitor counts that as a failure. See toll-free and IVR monitoring.

Frequently asked questions

What is the difference between 180 Ringing and 183 Session Progress?

180 Ringing only signals that the destination is ringing; the caller's side generates the ringback tone locally. 183 Session Progress usually carries an SDP and opens the audio before the answer, so the network can play its own tone, an announcement or an error message. Neither one means the call was answered.

Does hearing audio mean the call was answered?

No. With early media, the caller hears sound before any 200 OK. That sound can be ringback, music or a recording such as "the number you have called is not available". The call has only been answered when the 200 OK response to the INVITE arrives.

GLOSSARY

IP PBX: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: An IP PBX is a company's private telephone exchange running over an IP network, connecting extensions, trunks and IVR and routing internal and external calls with SIP and RTP.

In Portuguese: PABX IP

What is an IP PBX?

An IP PBX is a company's private telephone exchange that uses the data network instead of telephone wiring. It registers extensions (IP phones, softphones and adapters for analog handsets), connects the company to the carrier through [SIP trunks](#) or, via a gateway, [E1 trunks](#), and decides where each call goes.

Signaling between extensions, the PBX and the carrier uses [SIP](#), and audio travels in RTP packets. On top of that, the PBX offers the usual features: [IVR](#), call queues, transfer, call forwarding, voicemail, recording and the mapping of [DID](#) numbers to extensions.

An IP PBX can be a server or appliance on the company's premises, a virtual machine in a data center or a hosted service. In every case, it is only one part of the path. A call coming from outside goes through the carrier, the trunk, the firewall or [SBC](#), the PBX, and only then reaches the extension or IVR. Monitoring the PBX's CPU, memory and processes shows that the server is on, but not that a customer's call completes with audio in both directions.

In practice

An office with 40 extensions, a SIP trunk with 10 channels and a DID range from 11 3456-7000 to 7039. An outside call to 11 3456-7012 goes through:

Segment	What can fail
Carrier	Number routing, number portability, channels
SIP trunk	Registration, authentication, channel limit
Firewall or SBC	NAT, SIP ALG, blocked RTP ports

Segment	What can fail
IP PBX	Dial plan, DID mapping, codecs
Extension 7012	Extension unregistered, phone switched off

With the PBX up, any one of these segments can drop the call.

How Monitor uses this

Monitor works with any PBX because it calls numbers, just as a customer would. There is no integration, agent or configuration on the PBX: the test dials the public number, and the result shows whether the call completed across the whole path. See [SIP trunk monitoring](#).

Frequently asked questions

Do I need to integrate call monitoring with my IP PBX?

Not when the test is done with real calls. A test call dials the public number, just as a customer would, and travels through the carrier, the trunk, the PBX and the extension or IVR. So it does not depend on the PBX's brand, version or API, and it also catches failures outside the PBX, such as routing at the carrier or problems on the internet link.

What is the difference between an IP PBX and a cloud PBX?

The function is the same: extensions, queues, IVR and routing. With an on-premises IP PBX, the server sits at the company and the team takes care of hardware, updates and the link. With a cloud PBX, the provider hosts the system and extensions register over the internet. Either way, the call depends on the trunk, the network and the configuration, and it can fail even while the PBX is up.

GLOSSARY

IVR: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: An IVR (Interactive Voice Response, called URA in Brazil) is the system that answers calls without a human agent, plays recorded messages and routes callers based on DTMF keypresses or speech.

In Portuguese: URA

What is an IVR?

An IVR (Interactive Voice Response, known in Brazil as URA) is the system that answers a call without a human agent, plays a recorded greeting and offers a menu such as “press 1 for billing, 2 for support”. Callers respond with the phone's keys, sent as [DTMF](#) tones, or by speaking, on IVRs with speech recognition. The IVR then transfers the call to a queue, an extension or an automated service, such as reissuing a payment slip.

The IVR usually lives inside the [IP PBX](#) or a contact center platform, behind an [0800](#) (toll-free), 0300 or 4004 number. So a failure can be in several places: in the number's routing at the carrier, in the trunk, in the PBX or in the IVR application itself. To the caller, these failures all sound alike: a silent call, a busy signal, a carrier message or the wrong greeting.

One trait helps a lot with testing: an IVR answers every call. It does not get busy like a person or miss a call because it went to lunch. If the call gets a busy signal or rings until it drops, that alone is a sign of trouble.

In practice

A company publishes its customer service on an 0800 number and the IVR answers on the PBX. What the caller hears, and what it usually means:

What the caller hears	Typical SIP response	Situation
The company's greeting	200 OK	IVR working
"Your call could not be completed"	183 Session Progress , no 200	Carrier message; the IVR was never reached
Busy signal	486 Busy Here	Channels exhausted or IVR frozen
Rings until it drops	480 Temporarily Unavailable	IVR did not answer
Silence after answering	200 OK with no audio	Audio (RTP) failure

For example, if the IVR freezes for 30 minutes at a time of day with 2,000 calls per hour, about 1,000 customers go unserved before anyone notices the drop in queue volume.

How Monitor uses this

An IVR circuit passes when the call is answered and the chosen greeting is heard. If the carrier is the one speaking, with an error message in place of the greeting, the test fails. A busy signal or no answer also counts as a failure, because an IVR answers every call. See how [toll-free and IVR monitoring](#) works.

Frequently asked questions

How do I know whether the IVR is down?

Call it as a customer would and listen to what plays. A healthy IVR answers and plays the right greeting. A busy signal, ringing with no answer, silence after it picks up or a carrier message in place of the greeting all point to a failure. Looking only at the IVR server is not enough, because the problem can be in the number, the trunk or the carrier's routing.

Are URA and IVR the same thing?

Yes. URA (Unidade de Resposta Audível) is the name used in Brazil and IVR (Interactive Voice Response) is the English term. Both describe the system that answers the call, plays recorded messages and directs callers through keypresses or speech recognition. In vendor and contact center documentation, the two names appear side by side.

GLOSSARY

Jitter: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: Jitter is the variation in the arrival interval of a voice call's packets, which forces the receiver to hold packets in a buffer before playing the audio.

In Portuguese: Jitter

What is jitter?

Jitter is the variation in the arrival time of voice packets. On a VoIP call, the phone sends an [RTP](#) packet at fixed intervals, usually every 20 ms. If the network delivered every packet with the same delay, they would also arrive every 20 ms. In practice, router queues, Wi-Fi and congested links make some arrive early and others late. That irregularity is jitter.

To play the voice without stutters, the receiver holds packets in a jitter buffer and releases them at the right pace. The buffer handles small variations, but it has a price: every millisecond of buffer is one more millisecond of delay in the conversation. If a packet arrives after the moment it should have played, it is discarded and the listener hears a gap, just like [packet loss](#).

RTCP, the control protocol that accompanies RTP, reports an estimate of interarrival jitter calculated with the method in RFC 3550.

Latency

Latency is the delay of a packet from one end to the other, regardless of variation. ITU-T G.114 recommends not exceeding 400 ms of one-way delay for general network planning and notes that below 150 ms most applications have essentially transparent interactivity. Above that, people start talking over each other and waiting for the other side to respond. Total delay adds up the network, the codec and the jitter buffer.

In practice

A phone sends four packets, one every 20 ms. Here is when they arrive:

Packet	Sent (ms)	Arrived (ms)	Delay (ms)
1	0	40	40
2	20	75	55
3	40	72	32
4	60	110	50

The delay ranges from 32 to 55 ms, a swing of 23 ms. With a 30 ms buffer, all of them play on time, and the perceived delay stays around 70 ms, well under G.114's 150 ms. If a fifth packet arrived with 100 ms of delay, it would overflow the buffer and become a gap in the audio.

How Monitor uses this

The data circuit agent that 2N5 installs at each site measures latency, packet loss and jitter on that location's internet link, and shows charts for the last 100 hours. That way you can see whether the link was unstable at the time voice quality dropped. See [data circuits](#).

Frequently asked questions

What is acceptable jitter for VoIP?

No standard sets a single limit. What matters is whether the jitter fits in the phone's or the PBX's buffer. A 60 ms buffer absorbs variations smaller than that, but adds 60 ms to the call's delay. When jitter exceeds the buffer size, packets arrive too late, get discarded, and the effect is the same as packet loss.

Are jitter and latency the same thing?

No. Latency is how long a packet takes to travel from one end to the other. Jitter is how much that time varies from one packet to the next. A network can have high, stable latency with little jitter, or low latency with a lot of jitter. The first gets in the way of conversation; the second chops the audio.

GLOSSARY

Maintenance window: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: A maintenance window is a scheduled, agreed period in which a service may be unavailable for planned work, without failures during that interval being treated as incidents.

In Portuguese: Janela de manutenção

What is a maintenance window?

A maintenance window is a period booked in advance for work that may take the service down: upgrading the IP PBX, replacing the firewall, moving a SIP trunk to another carrier or changing the routing of an 0800 (toll-free) number. During the window an outage is expected and should not be treated as an incident.

The window solves two problems. The first is operational: without it, every test that fails during maintenance fires alerts, wakes the on-call engineer and trains the team to ignore warnings. The second is contractual: many SLA contracts separate scheduled downtime, announced in advance, from unscheduled downtime.

A good window has a defined start and end, covers only the affected circuits and ends with a check. Covering everything “just in case” hides real failures on lines that had nothing to do with the maintenance. And a window that ends without a test misses the most common case: the maintenance is over, but the line did not come back.

In practice

The carrier will migrate the headquarters SIP trunk on Saturday, from 11:00 p.m. to 1:00 a.m. The company has three monitored circuits:

Circuit	In the window?	Failure at 12:30 a.m.	Failure at 1:15 a.m.
0800 arriving over the headquarters trunk	Yes	No incident	Incident
DID extension at headquarters	Yes	No incident	Incident
Branch office IVR, on another trunk	No	Incident	Incident

The branch office failure at 12:30 a.m. is real and must alert, even with maintenance under way at headquarters. The 1:15 a.m. failure on the headquarters circuits, on the other hand, shows that the migration ended without the trunk coming back.

How Monitor uses this

A maintenance window covers the chosen circuits during a defined time slot. Failures that happen inside it do not open incidents, and the other circuits keep alerting as normal. When the window ends, the following tests open incidents again if the failure persists. See [alerts and escalation](#).

Frequently asked questions

Does scheduled maintenance count against the SLA?

It depends on the contract. Scheduled maintenance, announced with the notice set in the agreement and carried out within an agreed time slot, is commonly excluded from the availability calculation. Outside those conditions, downtime usually counts as normal. Read the clause carefully: notice period, maximum duration, how many windows per month and who can request them.

What is the best time for a maintenance window?

The time with the fewest calls on the affected circuits, which is not always the middle of the night. A 24-hour support 0800 line may be busy at night, while a scheduling IVR gets almost no calls on Sunday. Check call volume before booking, leave a margin for overruns and test the lines right after the window ends.

GLOSSARY

MOS: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: MOS (Mean Opinion Score) is the 1-to-5 score, defined by ITU-T P.800, that expresses the voice quality perceived by someone listening to a call.

In Portuguese: MOS

What is MOS?

MOS (Mean Opinion Score) is a 1-to-5 score that sums up how people perceive the quality of a voice call. The scale comes from ITU-T P.800: 5 is excellent, 4 good, 3 fair, 2 poor and 1 bad. Originally, the score came from tests with real listeners who rated audio samples; the average of those ratings is the MOS.

Score	Quality (ITU-T P.800)
5	Excellent
4	Good
3	Fair
2	Poor
1	Bad

Nobody convenes a listening panel for every call. That is why ITU-T P.800.1 separates the types of MOS by where the score comes from: MOS-LQS is the subjective listening score, given by people; [MOS-LQO](#) is the objective score, calculated by an algorithm that compares the received audio with the original; MOS-LQE is a listening estimate; and MOS-CQE is the conversational estimate that comes out

of a network model. When a report shows “MOS 4.1”, the first question is: which MOS? A score estimated from network numbers and a score measured on the audio can diverge sharply on the same call.

R-factor and the E-model

The E-model, from ITU-T G.107, is a network planning model. It starts with an R-factor and subtracts penalties for delay, [packet loss](#), [codec](#) and echo. With the default parameters, R comes out around 93.2. The standard provides the formula that converts R to MOS:

R-factor	MOS-CQE
90	4.34
80	4.03
70	3.60
60	3.10
50	2.58

MOS-CQE

MOS-CQE is the estimated conversational score, the E-model's output expressed on the MOS scale. The maximum is 4.5. It never listens to the audio: it calculates the score from parameters such as delay, loss and codec. It works well for sizing a network, but it cannot notice a silent call, one side with no audio or a recorded carrier message in place of the person being called.

In practice

Suppose a network device reports R = 80 for a call. By the G.107 table, that equals a MOS-CQE of 4.03, a “good” call. Now suppose that call landed on a carrier recording saying the number is not available. The packets arrived in order, with no loss and little [jitter](#), so the E-model keeps showing 4.03. For the caller, the call failed.

The reverse happens too. A network with 0% loss can deliver audio chopped up by a misconfigured transcoder. The E-model does not see the problem, because it lives in the content of the audio, not in the packets. Only a score measured on the received sound, compared with the expected sound, reveals both cases.

How Monitor uses this

Monitor never estimates MOS from network numbers. It measures MOS-LQO on the audio actually received on test calls. The details of how the score is calculated are in [voice quality with MOS-LQO](#) and in the [MOS-LQO](#) entry.

Frequently asked questions

What is a good MOS for VoIP?

On the ITU-T P.800 scale, 4 means good and 3 means fair. The E-model's MOS-CQE tops out at 4.5, so a score above 4 is already close to the best the model allows. Scores measured on the audio, such as MOS-LQO, depend on the algorithm and the codec's bandwidth, so it only makes sense to compare scores of the same type.

Are MOS and R-factor the same thing?

No. The R-factor is the output of the E-model (ITU-T G.107), on a scale that reaches about 93.2 with the default parameters. It is converted to MOS-CQE with a formula from the standard itself: R 80 equals 4.03 and R 70 equals 3.60. The R-factor is a network planning estimate, not a score taken from the audio.

GLOSSARY

MOS-LQO: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: MOS-LQO is the 1-to-5 score, defined in ITU-T P.800.1, that an algorithm calculates by comparing the audio received on a call with the reference audio that should have arrived.

In Portuguese: MOS-LQO

What is MOS-LQO?

MOS-LQO (Mean Opinion Score, Listening Quality Objective) is a 1-to-5 score that an algorithm calculates from the actual audio of a call. The algorithm takes two files, the reference audio and the audio that arrived at the other end, and measures how far the second has drifted from the first. The result is expressed on the same scale as the ITU-T P.800 MOS, where 5 is excellent and 1 is bad.

The acronym comes from ITU-T P.800.1, which classifies MOS scores by their source. "LQ" stands for listening quality and "O" for objective, meaning calculated by a machine from the sound. It contrasts with LQS, given by listeners, and with estimated scores, such as the E-model's MOS-CQE, which never listen to the audio.

Measured or estimated

The practical difference is what each score can see. An estimated score looks at the packets: delay, loss and jitter. A measured score looks at the sound. If the call arrives silent, if only one side has audio, if the carrier plays a message instead of reaching the destination or if a transcoder distorts the voice, the packets can be perfect and the estimated score stays high. MOS-LQO drops, because the received sound does not resemble what was expected.

The best-known algorithms for MOS-LQO are PESQ and POLQA, both ITU standards, and the open-source ViSQOL.

In practice

Picture three calls to the same number, all with 0% packet loss and jitter under 5 ms:

Call	What happened	Network	MOS-LQO
A	Clean audio	Normal	High
B	Audio went through two transcoders	Normal	Lower
C	Carrier played "number unavailable"	Normal	Very low

On the network, the three calls are identical. In the audio, only call A delivered what it should. That difference is why it pays to measure instead of estimate.

How Monitor uses this

Monitor measures MOS-LQO on every answered test call. It compares the received audio with the expected recording, the reference recorded when the circuit is set up, using a reference-based perceptual listening-quality model (ITU-T P.800.1). The default thresholds: from 3.8 up the call is healthy; below 3.0 it counts as degraded; between 3.0 and 3.8 it passes with a note. See [voice quality with MOS-LQO](#).

Frequently asked questions

What is the difference between MOS-LQO and MOS-CQE?

MOS-LQO is measured: an algorithm compares the received audio with the original audio and produces the score. MOS-CQE is estimated: the ITU-T G.107 E-model calculates the score from delay, loss and codec, without listening to anything. That is why MOS-CQE can give a good score to a call that landed on a carrier recording, and MOS-LQO cannot.

Is a MOS-LQO from one algorithm valid for another?

Not directly. PESQ, POLQA and ViSQOL produce scores on the same 1-to-5 scale, but each has its own mapping and its own ceiling. ViSQOL in speech mode reaches 5.0, while audio mode tops out near 4.75. Compare scores from the same algorithm, in the same mode, before drawing conclusions.

GLOSSARY

MTTR and MTTA: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: MTTR and MTTA are per-incident time averages: MTTA measures how long the team takes to acknowledge the alert, and MTTR how long it takes to restore the service.

In Portuguese: MTTR e MTTA

What are MTTR and MTTA?

MTTA (Mean Time to Acknowledge) is the average time between an incident's alert and the moment someone acknowledges it and takes ownership. MTTR is the average time until the service is working again. The acronym MTTR has variations (Mean Time to Repair, to Restore, to Recover, to Resolve), so every company needs to state where the clock starts and where it stops.

The two numbers answer different questions. MTTA shows whether the alert reaches anyone: a trunk that goes down at 2 a.m. and is only acknowledged at 7 a.m. has an alerting problem, not a repair problem. MTTR shows how long the end customer goes without service, adding up reaction, diagnosis, contact with the carrier and the fix.

In contracts, they show up as response and repair times alongside the availability [SLA](#). A common source of confusion: MTTR can be measured from the actual start of the failure or from detection. If detection depends on a customer complaining, the official MTTR looks short while the real outage was much longer.

In practice

Four incidents in one month, counted from detection:

Incident	Until acknowledged	Until service restored
A, SIP trunk, 8:00 a.m.	6 min	40 min
B, IVR, 10:30 a.m.	2 min	15 min
C, 0800, 2:10 a.m.	25 min	95 min
D, DID extension, 4:45 p.m.	3 min	10 min

$MTTA = (6 + 2 + 25 + 3) \div 4 = 9$ minutes. $MTTR = (40 + 15 + 95 + 10) \div 4 = 40$ minutes. Incident C, in the middle of the night, pulls both averages up. Without it, MTTA would drop to about 4 minutes, a sign that the problem lies in after-hours alerting.

How Monitor uses this

The report shows the average time to first acknowledgment of each incident, which is MTTA, and the downtime of each incident. That way you can see whether the delay is in someone noticing the alert or in fixing the failure. See [alerts and escalation](#).

Frequently asked questions

What is the difference between MTTR and MTTA?

MTTA ends when someone takes ownership of the incident, that is, acknowledges the alert. MTTR ends when the service is working again. MTTA measures how the team and the alerting process react; MTTR measures the reaction plus diagnosis and repair. A high MTTA almost always points to an alert nobody saw, sent to the wrong person or outside business hours.

How do you reduce MTTA?

Send the alert to someone who can act, through a channel that person sees outside business hours, and escalate to the next person on the list when nobody acknowledges within a few minutes. Cutting false alarms helps too, because a team that gets many alerts with no real problem starts ignoring them. Measuring MTTA per incident shows where the process gets stuck.

GLOSSARY

Packet loss: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: Packet loss is the fraction of a VoIP call's packets that never reach the destination or arrive too late to be played, causing gaps in the audio.

In Portuguese: Perda de pacotes

What is packet loss?

Packet loss is the share of voice packets that leave one end of the call and never reach the other. On a VoIP call, speech is cut into 20 ms chunks and each chunk travels in an [RTP](#) packet. If a packet is lost, that stretch of sound does not exist at the destination. The rate is expressed as a percentage: packets

lost divided by packets sent.

Packets get lost through congestion, when a router discards what does not fit in its queue; through errors on radio or Wi-Fi links; and through excessive [jitter](#), when a packet arrives after its time to play and the receiver throws it away. Because RTP runs over UDP, there is no retransmission: in voice, a late packet is no longer useful.

Codecs try to mask the failure with packet loss concealment, which repeats or interpolates the neighboring sound. That works for isolated packets. In bursts, when several packets in a row go missing, concealment has nothing to copy and the listener hears clipped words.

In practice

With 20 ms packets, each direction of the call sends 50 packets per second, or 3,000 per minute.

Loss	Packets lost per minute	Audio lost per minute
0.5%	15	0.3 s
1%	30	0.6 s
3%	90	1.8 s
5%	150	3.0 s

The same 30 packets lost in a minute sound different depending on how they are spread. Scattered, they are 20 ms gaps that concealment hides. Together, they are 0.6 s of silence all at once, enough to wipe out a word like “confirmed”.

How Monitor uses this

The data circuit agent that 2N5 installs at each site measures packet loss, latency and jitter on that location’s internet link, with charts for the last 100 hours. On a voice call, lost packets show up as choppy audio and pull down the measured [MOS-LQO](#). See [data circuits](#).

Frequently asked questions

How much packet loss can a VoIP call handle?

It depends on the codec and on how the loss is distributed. Codecs with loss concealment mask isolated packets, and 1% spread across the call can go almost unnoticed. The same rate concentrated in bursts, several packets in a row, wipes out whole syllables. That is why the average rate alone is not enough: you have to look at the audio that arrived.

Are packet loss and jitter related?

Yes. When jitter exceeds the receiver’s buffer size, packets that arrive late are discarded. To the listener, a packet that arrived too late and a packet that never arrived have the same result: a gap. Many devices count both together in total loss.

GLOSSARY

POLQA and PESQ: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: POLQA (ITU-T P.863) and PESQ (ITU-T P.862) are ITU standards that calculate a MOS-LQO score by comparing the audio received on a call with the reference audio.

In Portuguese: POLQA e PESQ

What are POLQA and PESQ?

POLQA and PESQ are algorithms standardized by the ITU to calculate a call's voice quality from its audio. Both are full-reference: they compare the audio that arrived at the other end with the original audio and return a **MOS-LQO** score, on the 1-to-5 **MOS** scale.

PESQ (Perceptual Evaluation of Speech Quality) is ITU-T P.862, created for narrowband speech, classic 8 kHz telephony. For many years it was the industry reference for measuring voice quality objectively. On January 5, 2024 the ITU withdrew P.862 and began pointing to P.863 in its place.

POLQA (Perceptual Objective Listening Quality Analysis) is ITU-T P.863, the successor. It also covers wideband and super-wideband speech, which matters for modern codecs, and handles distortions that PESQ assessed poorly. POLQA is distributed under a commercial license.

Both rely on a model of how the human ear perceives distortion. The score does not say what went wrong; it says how far the received sound drifted from the original to the ears of an average listener.

In practice

Algorithm	Source	Status	Access
PESQ	ITU-T P.862	Withdrawn on 1/5/2024	Older implementations still around
POLQA	ITU-T P.863	Standard in force	Commercial license
ViSQOL	Google Research	Maintained, version v3	Open source

An example of why scores should not be mixed: a 2019 report shows a MOS of 3.9 calculated with PESQ, and a test today shows 3.6 with another algorithm. The 0.3 drop may mean nothing, because each algorithm has its own mapping and its own ceiling. To track a line over time, the algorithm and the mode must always be the same.

How Monitor uses this

Monitor uses neither POLQA nor PESQ. It measures MOS-LQO with a reference-based perceptual listening-quality model, comparing the audio received on each answered test call with the circuit's reference recording. The reason to explain all three options here is that many scores on the market come from PESQ or POLQA. See [voice quality with MOS-LQO](#).

Frequently asked questions

Can PESQ still be used?

PESQ implementations are still around, and many older reports use its scores. But the ITU withdrew ITU-T P.862 on January 5, 2024 and points to P.863, POLQA, as its replacement. For a new process, PESQ is no longer the standard in force, and its scores should not be compared directly with POLQA's.

Do POLQA and ViSQOL give the same score?

Not necessarily. Both compare the received audio with the reference and return a 1-to-5 MOS-LQO, but they use different perception models and different mappings to the scale. The same call can get different scores. Always compare scores from the same algorithm and the same mode.

GLOSSARY

Post dial delay (PDD): definition and a practical example

By 2N5 team Updated October 14, 2026

Short answer: Post dial delay (PDD) is the interval between the moment the call is requested, when dialing ends, and the first sign that the destination is being called: the ringing or the opening message the caller hears.

In Portuguese: Post dial delay (PDD)

What is post dial delay?

Post dial delay, or PDD, is how long the caller waits in silence after dialing. It starts when the call is requested, as dialing ends, and stops when the first sign arrives that the destination is being called: the ringing tone, or the opening message of an IVR or of the carrier. In SIP, it is the time between the INVITE and the first provisional response, such as 180 Ringing or 183 Session Progress. RFC 6076, which defines performance metrics for SIP, calls this time Session Request Delay.

PDD matters more than it seems. After a few seconds of silence, the caller assumes the call is not working and hangs up. For anyone with an 0800 number or a contact center, each abandonment like that is a customer who tried and could not get through.

What increases PDD

- **Routes with several carriers in sequence**, each adding a little time.
- **A slow number portability lookup** before the call is routed.
- **An alternate route:** when the first destination fails, the equipment tries the next one before the caller hears anything.
- **Slow DNS resolution** or an overloaded edge device.

A SIP trunk can be registered and complete every call and still have a bad PDD, which is why PDD shows up as a customer complaint before it shows up as an outage.

In practice

Here is what the caller experiences on two calls to the same number (illustrative values):

Call	Last digit	First ring	PDD	What the caller notices
A	0 s	2 s	2 s	The call starts normally
B	0 s	11 s	11 s	A long silence; many people hang up

Both calls complete. Only measuring the time to the first ring shows the difference.

How Monitor uses this

Monitor calls each voice circuit the way a customer does, and a call that takes too long to connect ends up counted as a failure, with the SIP code and the Q.850 cause the route returned. Monitor does not publish PDD as a separate metric. Use the PDD in your call detail records to watch the trend, and the test call to confirm, from the outside, that the number still connects. See [SIP trunk and VoIP monitoring](#) and [synthetic monitoring](#).

Frequently asked questions

What is post-dial delay and what is normal?

Post-dial delay (PDD) is the silence between the last dialed digit and the first ring. There is no single value in a standard for every kind of call: it is usually a few seconds and varies with the number of carriers on the path and with the number portability lookup. What counts is the circuit's own baseline: if a number that took a few seconds starts taking much longer, something changed on the route.

What increases post dial delay on a SIP trunk?

Routes with several carriers in sequence, a slow number portability lookup, a failure at one destination that forces a retry on an alternate route, slow DNS resolution and overloaded equipment. In each case the caller hears silence between the last digit and the first ring, and many hang up before it ends.

GLOSSARY

RTP: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: RTP (Real-time Transport Protocol), defined in RFC 3550, is the protocol that carries a VoIP call's audio in numbered, timestamped packets.

In Portuguese: RTP

What is RTP?

RTP (Real-time Transport Protocol) is the protocol that carries the voice on a VoIP call. While [SIP](#) sets up the call, RTP carries the audio itself, in small packets sent at a fixed pace, usually one every 20 ms. It is defined in RFC 3550 and runs over UDP.

Each RTP packet has a 12-byte header with four main pieces of information: the payload type, which identifies the [codec](#); the sequence number, which makes lost or out-of-order packets visible; the timestamp, used to play the audio at the right pace and to calculate [jitter](#); and the SSRC, which identifies the source of the stream. Each direction of the call is a separate stream, with its own ports, agreed in the SDP.

RTCP

RTCP (RTP Control Protocol), defined in the same RFC 3550, travels alongside RTP and carries no voice. It sends periodic reports with packets sent, fraction [lost](#) and interarrival jitter. Traditionally it uses the next odd port after RTP's.

SRTP

SRTP (Secure RTP), from RFC 3711, encrypts the audio payload and authenticates the packets. Anyone who captures the traffic sees the headers but cannot hear the conversation.

In practice

A test call of about 7 seconds with 20 ms packets generates some 350 packets per direction. A log line such as "RTP 343 sent / 340 received" shows:

Measure	Value
Packets sent	343
Packets received	340
Lost	3 (0.9%)
Audio lost	about 60 ms

If the received side showed 0, the call would have been answered with no audio in that direction, a typical NAT or [SIP ALG](#) case.

How Monitor uses this

The evidence for each incident shows the RTP packets sent and received on each test call, for example "RTP 343 sent / 340 received". Monitor checks that audio flows in both directions, so a call that is answered but silent on one side does not pass as healthy. See [voice quality with MOS-LQO](#).

Frequently asked questions

Why does the call connect but have no audio?

Because SIP and RTP take different paths. SIP can set up the call normally while the RTP packets are blocked by a firewall, sent to a private IP because of NAT or rewritten by a SIP ALG. The result is an answered call that is silent in both directions or in just one.

What is the difference between RTP and SRTP?

SRTP (RFC 3711) is RTP with encryption, authentication and replay protection. The packet format and pace are the same; the difference is that the audio payload travels encrypted and cannot be heard by anyone who captures the traffic. The keys are agreed in the call signaling.

GLOSSARY

SBC: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: An SBC (Session Border Controller) is the device or software at the edge of a VoIP network that controls, protects and adapts the SIP signaling and audio going in and out.

In Portuguese: SBC

What is an SBC?

An SBC (Session Border Controller) is the device that sits on the border between two VoIP networks, usually between the company's PBX and the carrier's SIP trunk. Every call coming in or going out passes through it, both the [SIP](#) signaling and the [RTP](#) audio. It can be a physical appliance, a virtual machine or software on the telephony server itself.

The most common functions of an SBC are:

- Topology hiding: replaces internal IPs with a public address, so the carrier cannot see the company's network.
- NAT traversal: fixes addresses in SIP and SDP, the job a [SIP ALG](#) tries to do and often does badly.
- Security: blocks scans, fraud attempts and message floods.
- Normalization: adjusts headers so equipment from different vendors can understand each other.
- Transcoding: converts between [codecs](#) when the two ends have none in common.
- Admission control: limits how many simultaneous calls get through.

Because it sits in the middle of everything, the SBC is also a place where a call can fail: a wrong rule, an exhausted license or an expired certificate blocks traffic.

In practice

An INVITE leaves the PBX and passes through the SBC before reaching the carrier:

Field	Leaving the PBX	Leaving the SBC
IP in Contact	10.0.0.5	203.0.113.10
Media IP in SDP	10.0.0.5	203.0.113.10
RTP port	12000	40000
Codecs offered	G.711, G.729	G.711

The carrier only sees the public address and the codec the SBC let through. If port 40000 is blocked on the firewall, the call is answered and stays silent.

How Monitor uses this

There is nothing to install on the SBC. Monitor calls the number from outside, as any customer would, so the test call passes through the SBC exactly like a real call. See [SIP trunk monitoring](#).

Frequently asked questions

Do I need an SBC to have a SIP trunk?

Not always. Many IP PBXs talk directly to the carrier. An SBC becomes worthwhile when the company needs to hide its internal network, control how many calls come in, connect equipment that speaks different SIP dialects or transcode between codecs. Some carriers also put their own SBC in front of the trunk.

Can an SBC cause calls with no audio?

Yes. If the SBC rewrites the SDP with a wrong IP or port, or if the firewall in front of it blocks the RTP port range, the call is answered and the audio does not get through in one or both directions. The signaling shows 200 OK, so only a test that checks the audio reveals the problem.

GLOSSARY

SIP: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: SIP (Session Initiation Protocol), defined in RFC 3261, is the signaling protocol that sets up, modifies and ends voice calls over IP networks.

In Portuguese: SIP

What is SIP?

SIP (Session Initiation Protocol) is the protocol that opens, controls and ends VoIP calls. It plays the role of signaling in traditional telephony: it announces a call, makes the destination ring, reports the answer and marks the end. SIP is defined in RFC 3261 and uses text messages similar to HTTP's, usually sent over port 5060.

SIP requests are called methods. The most used are INVITE (start a call), ACK (confirm the final response), BYE (end the call), CANCEL (give up before the answer), REGISTER (register a phone or trunk) and OPTIONS (check whether the other side responds).

Each request gets numbered responses in six classes: 1xx provisional, such as 180 Ringing; 2xx success, such as 200 OK; 3xx redirection; 4xx errors on the caller's or the destination's side, such as 486 Busy Here; 5xx server failure, such as 503; and 6xx global failure. A call's final response is the one that best tells what happened. The full list is in [SIP codes](#).

SDP

SIP does not carry audio. Inside the INVITE and the response, it carries a media description in SDP (Session Description Protocol), from RFC 4566, now replaced by RFC 8866. The SDP gives the IP and port where each side expects to receive RTP and which [codecs](#) it accepts. One side offers, the other answers, and the voice starts flowing along the agreed path.

In practice

A successful call exchanges these messages:

Step	From	Message
1	Caller	INVITE with SDP
2	Destination	100 Trying
3	Destination	180 Ringing
4	Destination	200 OK with SDP
5	Caller	ACK
6	Either side	BYE
7	The other side	200 OK

If the destination is busy, step 4 becomes 486 Busy Here, and the response may carry Q.850 cause 17, user busy, in the Reason header, as described in RFC 6432.

How Monitor uses this

Monitor reads the final SIP response and the Q.850 cause of every test call. With both, the incident says whether the destination refused the call, the network failed or the number does not exist. See [SIP trunk monitoring](#).

Frequently asked questions

What is the difference between SIP and RTP?

SIP is the signaling: it says who is calling whom, whether the destination rings, whether it answers and when the call ends. RTP carries the audio. The two take different paths and ports, so a call can be answered in SIP and still have no voice, when RTP is blocked or diverted along the way.

What is the Q.850 cause in a SIP response?

It is a numeric code inherited from ISDN telephony that explains why the call ended, such as 16 for normal call clearing or 17 for user busy. Many carriers send it in the Reason header of the SIP response, as described in RFC 6432. It complements the SIP code, which is sometimes too generic to diagnose the failure.

GLOSSARY

SIP ALG: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP ALG (Application Layer Gateway) is a feature of routers and firewalls that rewrites addresses inside SIP messages to get through NAT, often breaking the audio on calls.

In Portuguese: SIP ALG

What is SIP ALG?

SIP ALG (Application Layer Gateway) is a feature of routers and firewalls that opens [SIP](#) messages and replaces the private IP addresses inside them with the connection's public IP. The idea is to help the call get through NAT. In practice, many implementations rewrite the wrong fields, leave others uncorrected or break messages that the PBX or the [SBC](#) had already adjusted. The classic result is one-way silent audio, calls that drop on their own and extensions that lose their registration.

SIP ALG often comes enabled out of the box on home and small-business routers. When IP telephony starts failing after a router or link change, the ALG is one of the first suspects.

NAT

NAT (Network Address Translation) is what lets several devices with private IPs, such as 192.168.0.20, reach the internet through a single public IP. The router swaps the address in the packet header. The problem is that SIP repeats addresses inside the message itself, in the Via and Contact headers and in the SDP, which says where to send the [RTP](#). If those internal addresses reach the carrier uncorrected, it sends the audio to a private IP that does not exist on the internet.

In practice

Symptom	Likely cause
Call answered, one side cannot hear	Private IP or wrong port in the SDP
Call drops after about 32 seconds	ACK does not arrive; the destination gives up after $64 \times T1$ (RFC 3261)
Extension loses registration from time to time	ALG closes or alters the REGISTER session
Works on one link and fails on another	Different routers, one with ALG enabled

To disable it, look on the router's or firewall's page for SIP ALG, SIP Helper or SIP inspection, turn it off and restart. On Linux-based devices, the feature is usually the SIP conntrack module.

How Monitor uses this

A test call that connects but has no audio in one direction is detected as one-way audio and counts as a failure, even with a SIP 200 OK response. See [SIP trunk monitoring](#).

Frequently asked questions

How do I disable SIP ALG on my router?

Open the router's or firewall's admin page and look under NAT, firewall or advanced settings for SIP ALG, SIP Helper, SIP Passthrough or SIP inspection. Uncheck the option, save and restart the device. If the router was supplied by the carrier and the option does not appear, ask their support to turn it off or to set bridge mode. Then make a test call in both directions.

Should I always disable SIP ALG?

In most cases, yes, as long as something else handles NAT correctly: the PBX configured with the public IP, phones using STUN or an SBC. Without any of these, disabling the ALG can leave phones behind NAT with no audio. Test before and after the change.

GLOSSARY

SIP registration: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: SIP registration is the process in which a phone, PBX or trunk sends a REGISTER message to the registrar to report its current address and be able to receive calls for its account.

In Portuguese: Registro SIP

What is SIP registration?

SIP registration is how a phone, PBX or [SIP trunk](#) tells the registrar where it is. The device sends a REGISTER request, defined in RFC 3261, with its account and its IP address and port. The registrar keeps that binding for a period, given in the Expires field, and forwards calls meant for the account to that address.

The usual flow has four steps: the device sends REGISTER; the registrar answers 401 Unauthorized with a challenge; the device resends the REGISTER with a response calculated from the password; the registrar confirms with 200 OK. Before the period runs out, the device registers again. If the registration

expires and is not renewed, inbound calls stop arriving.

Many devices also use SIP OPTIONS (RFC 3261, section 11) as a keep-alive: a periodic query that, when answered with 200 OK, shows the other side is alive. The regular traffic also keeps the NAT mapping open on the router, which avoids problems with [SIP ALG](#) and firewalls.

Registration and OPTIONS test the signaling up to the registrar. They do not go through the carrier's routing, the channel limit, the codecs or the audio.

In practice

What each check proves about a trunk:

Check	Expected response	Proves that
REGISTER	200 OK, with Expires of 3600 s, for example	Credentials accepted and registrar reachable
OPTIONS	200 OK	The other side responds to signaling
Real call	200 OK and audio in both directions	Calls complete with voice

A trunk can pass the first two rows all day and fail the third, for example with [503 Service Unavailable](#) when the channels run out. On the PBX, the trunk keeps showing as registered throughout the failure.

How Monitor uses this

A trunk can be registered and still not complete calls, so the test is a real call. When 2N5's own calling account cannot register, the test is marked "monitoring blind" and is never blamed on the customer's line. See [SIP trunk monitoring](#).

Frequently asked questions

Does a registered trunk mean it is working?

Not necessarily. Registration shows that the PBX can talk to the carrier's registrar and that the credentials were accepted. Calls can still fail because of a channel limit, an incompatible codec, wrong routing at the carrier or audio blocked by NAT. Only a real call, with a final SIP response and audio in both directions, shows that the trunk works.

What is SIP OPTIONS for?

It is a request defined in RFC 3261, section 11, for asking another device which capabilities it supports. In practice, PBXs and SBCs send OPTIONS periodically as a keep-alive: if they get 200 OK, they consider the other side up. Sending it regularly also keeps the NAT mapping open on the router. An answered OPTIONS does not prove that calls complete.

GLOSSARY

SIP trunk: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: A SIP trunk is the voice connection between a company's PBX and the carrier over an IP network, with SIP signaling and audio in RTP, in place of physical lines or E1.

In Portuguese: Tronco SIP

What is a SIP trunk?

A SIP trunk is the connection between the company's [IP PBX](#) and the carrier over the data network, either across the internet or over a dedicated link. Signaling uses [SIP](#) and voice travels in RTP packets. It replaces analog lines and the [E1 trunk](#): channels stop being physical and are set by contract instead.

The carrier identifies the trunk in one of two ways: by [SIP registration](#), with a username and password, or by IP address, when it only accepts calls from known IPs. Between the PBX and the internet there is usually a firewall or an [SBC](#), which handles NAT, security and message normalization.

Every call depends on several things at once: the internet link, authentication, the channel limit, [codec](#) negotiation and RTP getting through in both directions. That is why a SIP trunk can show as registered on the PBX and still fail to complete calls, or complete them with no audio. The final SIP response of each call shows at which of these stages the failure happened.

In practice

Common SIP trunk failures and the response they tend to produce:

Symptom	Typical SIP response	Likely cause
Call rejected right at the start	401 Unauthorized	Credentials or authentication challenge
Rejected at peak time	503 Service Unavailable	Channel limit or carrier unavailable
Rejected right after the offer	488 Not Acceptable Here	No codec in common
Answers, but one side cannot hear	200 OK	NAT or SIP ALG altering the SDP
Destination does not answer	480 Temporarily Unavailable	Extension or destination down

The full list is in [SIP codes](#).

How Monitor uses this

A real test call to a number on the trunk goes through everything at once: authentication, routing, the channel limit, codecs and audio in both directions. The final SIP response of that call says where it stopped. See [SIP trunk monitoring](#).

Frequently asked questions

How do I know whether the SIP trunk is working?

Make a real call over it and check three things: the final SIP response, audio in both directions and voice quality. A registered status on the PBX only shows that signaling reached the registrar. The trunk can be registered and still reject calls because of a channel limit, an incompatible codec or a routing error at the carrier.

How many channels do I need on a SIP trunk?

The number of simultaneous calls at peak time, with headroom. If the real peak is 18 calls at once, a 20-channel contract is at its limit; once that limit is hit, the carrier rejects the next call, often with 503 or 486. The PBX call log shows the real peak for each day.

GLOSSARY

SLA: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: An SLA (Service Level Agreement) is the contractual commitment that defines the minimum level of a service, such as 99.9% availability, and how it is measured and calculated.

In Portuguese: SLA

What is an SLA?

An SLA (Service Level Agreement) is the part of the contract that says how much the service has to work and what happens when it does not. In telephony, the central number is availability: the percentage of time the line, the trunk or the 0800 (toll-free) number received and completed calls.

The formula is simple: $\text{availability} = (\text{total time} - \text{unavailable time}) \div \text{total time} \times 100$. The hard part is the rest of the agreement. A well-written SLA defines the measurement period (month or year), what counts as unavailable (no dial tone, no audio, call rejected), who measures, what is excluded (such as a scheduled [maintenance window](#)) and the compensation, usually a credit on the bill.

Other indicators often appear alongside availability, such as response and repair time, measured by [MTTA](#) and [MTTR](#). Without independent measurement, the company depends on the carrier's own report to know whether the SLA was met.

In practice

How much downtime each target allows:

SLA	Maximum per month (30 days)	Maximum per year (365 days)
99%	7 h 12 min	3 days 15 h 36 min
99.5%	3 h 36 min	43 h 48 min
99.9%	43 min 12 s	8 h 45 min 36 s
99.99%	4 min 19 s	52 min 34 s

The difference between yearly and monthly matters: with a 99.9% yearly SLA, a 6-hour outage in March does not breach the contract, but it would breach a 99.9% monthly SLA. Run the numbers in the [SLA calculator](#) and see the full method in the [telephony SLA guide](#).

How Monitor uses this

The [SLA reports](#) are monthly and show the figure for the last 12 months, which decides a yearly SLA, and the month's figure as a reference. Periods when monitoring had no visibility because of 2N5 are left out of the calculation, so they count neither against the customer's line nor in its favor.

Frequently asked questions

How do you calculate a trunk's availability?

Divide the time the trunk was working by the total time in the period and multiply by 100. In a 30-day month (43,200 minutes) with 90 minutes of downtime, availability is $(43,200 - 90) \div 43,200 \times 100 = 99.79\%$. The deciding point is how downtime is measured: by real calls, by opened tickets or by the

carrier's own report.

Yearly or monthly SLA: which is better for the customer?

Monthly usually protects more. At 99.9%, a monthly SLA tolerates about 43 minutes per month, while a yearly one tolerates about 8 hours and 45 minutes across the whole year. A single 6-hour outage fits within the yearly SLA but breaks the monthly one by a wide margin. With a yearly SLA, the calculation has to track the last 12 months to know whether the target still holds.

GLOSSARY

STIR/SHAKEN: definition and a practical example

By 2N5 team Updated October 7, 2026

Short answer: STIR/SHAKEN is the set of standards that digitally signs the calling number of a SIP call, so the terminating network can verify whether the caller ID is legitimate.

In Portuguese: STIR/SHAKEN

What is STIR/SHAKEN?

STIR/SHAKEN is a set of standards for fighting calling-number spoofing on SIP calls. The originating carrier digitally signs the caller's number, and the terminating network checks that signature before delivering the call. STIR is the IETF part, with the protocols; SHAKEN is the usage framework that defines how carriers apply those protocols.

Four RFCs form the foundation. RFC 8224 defines the Identity header, which carries the signature inside the INVITE message. RFC 8225 defines PASSporT, a signed token with the calling number, the called number and the time. RFC 8226 defines the certificates that prove the signer has rights to the numbers. RFC 8588 adds the SHAKEN extension, with the attestation level and an origination identifier.

During verification, the terminating network downloads the certificate referenced in the header, checks the signature, confirms the time is recent and compares the numbers in the token with those in the call. The result can become a verified badge on the screen, a suspicious-call warning or a block, depending on each carrier's policy.

In practice

Main fields of a PASSporT with the SHAKEN extension:

Field	Example	Meaning
attest	A	Attestation level: A, B or C
orig	5511345670000	Calling number
dest	5521987650000	Called number
iat	1759838400	Time of signing, in Unix seconds
origid	unique identifier	Point in the network where the call entered

If the token arrives with an old iat or with an orig that differs from the call's number, verification fails, even with a correct signature. That prevents a valid token from being copied and reused on another call.

How Monitor uses this

Monitor does not sign or authenticate caller ID. Test calls check whether the number answers, what the final SIP response was and whether audio gets through. STIR/SHAKEN signing stays with the carrier that originates the call. See [SIP trunk monitoring](#).

Frequently asked questions

What do attestation levels A, B and C mean?

They are the confidence levels the originating carrier declares, defined in RFC 8588. A (full): the carrier knows the customer and knows they are allowed to use the number. B (partial): it knows the customer but cannot vouch for the number. C (gateway): the call came in from another network, such as an international gateway, and the carrier does not know where it came from.

Does STIR/SHAKEN stop calls with spoofed numbers?

Not on its own. It gives the terminating network verifiable information about the calling number and about the confidence level of the carrier that signed it. What to do with that information, such as showing a verified badge, flagging the call as suspicious or blocking it, is up to each carrier. A call with a spoofed number can still arrive with C attestation.

GLOSSARY

Synthetic monitoring: definition and a practical example

By [2N5 team](#) Updated October 14, 2026

Short answer: Synthetic monitoring, or synthetic call monitoring, is testing in which a system places real calls on a schedule to a number or circuit and records whether the call connected, what answered and how the audio sounded.

In Portuguese: Monitoramento sintético

What is synthetic monitoring?

Synthetic monitoring means monitoring a service by doing what a user would do, instead of waiting for a user to show up. In telephony these are synthetic calls: real test calls, placed by a system on a schedule, to a [toll-free \(0800\) number](#), an [IVR](#) or a [SIP trunk](#). Each call records whether it connected, who answered, whether audio flowed both ways and what the quality was.

The adjective comes from website and network monitoring, where a test simulates a user's visit. Among people who run voice networks, the more common name is active monitoring. Both describe the same idea, and both stand opposite passive monitoring, which analyzes the calls customers have already made, through SIP and RTP capture or call detail records.

Synthetic or passive

Question	Passive	Synthetic
Analyzes real customer calls	Yes	No: it places its own calls
Sees the 0800 that nobody called overnight	No, it only sees what passes	Yes, the test call dials

Question	Passive	Synthetic
Measures the quality of the received audio	Through the network, usually estimated	From the audio of each test call
Shows whose fault a failure is	Depends on what was captured	Brings the route's answer on every call

One does not replace the other. Passive shows what customers experienced; synthetic shows what a customer would experience right now, even when nobody called.

In practice

A circuit is tested every 15 minutes, for example. At 3:10 a.m. the carrier route stops completing calls. With passive monitoring, nothing shows until someone calls, and in the middle of the night that can be late. With synthetic monitoring, the 3:15 call fails and is retried, and four consecutive failed calls open an incident with the SIP code received, before the first customer complains.

How Monitor uses this

Monitor is a synthetic monitoring service: 2N5 calls each voice circuit at the agreed pace, checks the response, the audio and the message, retries a failed call, and only then opens the incident. See the [guide to synthetic call monitoring](#) and [SIP trunk and VoIP monitoring](#) for how it works.

Frequently asked questions

Is synthetic monitoring the same as active monitoring?

In telephony, in practice, yes. Both names describe tests that the monitoring itself starts, as opposed to passive monitoring, which only observes the calls customers have already made. Synthetic is the term inherited from website and network monitoring; active is the one voice engineers use more often.

Does a synthetic call count as a real call?

Yes, in the sense that matters: it crosses the carrier, the route and the answering point like a customer call, so it finds failures the equipment cannot see. What makes it synthetic is only who starts it, a system, and the script, which is always the same, so each call can be compared with the one before.

GLOSSARY

ViSQOL: definition and a practical example

By [2N5 team](#) Updated October 7, 2026

Short answer: ViSQOL (Virtual Speech Quality Objective Listener) is an open-source algorithm from Google Research that compares received audio with its reference and returns a 1-to-5 MOS-LQO score.

In Portuguese: ViSQOL

What is ViSQOL?

ViSQOL (Virtual Speech Quality Objective Listener) is an algorithm that scores the quality of audio by comparing it with the original. It was created by Google Research, is open source and is published at github.com/google/visqol. The current version is described in the paper “ViSQOL v3” by Chinen et al., presented at the QoMEX conference in 2020.

ViSQOL is a full-reference algorithm: it needs two files, the reference audio and the degraded audio. It aligns the two in time, compares their spectrograms segment by segment and measures the spectro-temporal similarity between them. That similarity is converted into a **MOS-LQO** score, on the same 1-to-5 scale as **MOS**.

The algorithm has two modes. Speech mode works with 16 kHz audio and returns MOS-LQO from 1 to 5, with the scale tuned to reach 5.0. Audio mode was designed for music and general audio, and its maximum score is near 4.75. The two scales do not mix: a 4.5 in audio mode is not the same as a 4.5 in speech mode.

In practice

Mode	Use	Sample rate	Maximum score
Speech	Voice, telephony	16 kHz	5.0
Audio	Music, general audio	48 kHz	about 4.75

On a call using the G.711 **codec**, audio is sampled at 8 kHz. To use speech mode, the reference and the received audio are brought to 16 kHz before the comparison. The reference has to be the same content that should arrive: if the destination plays a greeting, the reference is that greeting, recorded. If the received audio is something else, such as a carrier message, similarity drops and the score follows.

Comparing scores

Scores from different models don't mix: each algorithm has its own mapping and its own ceiling. When you read a MOS-LQO report, check which model and which mode produced the score. See [voice quality with MOS-LQO](#).

Frequently asked questions

Is ViSQOL free?

Yes. ViSQOL is published by Google Research as open source, in the github.com/google/visqol repository. Anyone can download, build and run the algorithm on their own audio files to check a score. That sets it apart from POLQA, the ITU-T P.863 standard, which is distributed under a commercial license.

Which ViSQOL mode should I use for phone calls?

Speech mode, which works with 16 kHz audio and returns MOS-LQO on a 1-to-5 scale, with a ceiling of 5.0. Audio mode was built for music and general audio at 48 kHz, and its maximum score is near 4.75. For telephone voice, speech mode is the right one.

GLOSSARY

Voice codecs: definition and a practical example

By [2N5 team](#) · Updated October 7, 2026

Short answer: Voice codecs are the algorithms that encode speech into digital data for transmission on a call and decode it at the destination, such as G.711, G.729 and Opus.

In Portuguese: Codecs de voz

What are voice codecs?

A voice codec is the algorithm that turns the sound of a voice into bits to travel across the network and rebuilds the sound on the other side. Every VoIP call uses a codec, and both ends need to use the same one. The choice determines how much bandwidth the call consumes, the audio quality and how well it holds up against [packet loss](#).

The three codecs most often found in telephony are:

- G.711 (ITU-T): uncompressed, samples speech at 8 kHz and produces 64 kbit/s. It has two variants, μ -law and A-law; Brazil uses A-law (PCMA). It is the de facto standard on SIP trunks.
- G.729 (ITU-T): compresses to 8 kbit/s. It saves bandwidth but loses quality and distorts DTMF tones sent inside the audio.
- Opus (RFC 6716): an open, adaptive codec from 6 to 510 kbit/s, with audio bandwidth up to 48 kHz. Common in WebRTC and apps.

The codec is agreed at the start of the call through SDP, inside the [SIP](#) message. The caller offers a list of codecs; the called party picks one from the list. If none works, the call fails.

In practice

Bandwidth per direction at the IP layer, with 20 ms packets (40 bytes of IP, UDP and RTP header per packet, not counting Ethernet):

Codec	Audio	Payload per packet	IP bandwidth per direction
G.711	64 kbit/s	160 bytes	80 kbit/s
G.729	8 kbit/s	20 bytes	24 kbit/s
Opus	6 to 510 kbit/s	variable	variable

A link carrying 10 simultaneous G.711 calls needs about 800 kbit/s in each direction just for voice. With G.729, about 240 kbit/s. Note that with G.729 the header weighs twice as much as the audio.

How Monitor uses this

When the destination accepts none of the codecs offered on the test call, the call ends with [SIP 488](#) and Monitor counts the circuit as down. See [SIP trunk monitoring](#).

Frequently asked questions

Which codec should I use for VoIP: G.711 or G.729?

G.711 does not compress speech and delivers classic telephone quality, at the cost of 64 kbit/s of audio in each direction. G.729 uses 8 kbit/s but loses quality and interferes with DTMF sent inside the audio. With bandwidth to spare, G.711 is the simpler choice. With tight bandwidth, G.729 fits better. What

matters is that both ends accept the same one.

What happens when the two ends have no codec in common?

The SDP negotiation fails and the call does not complete. The destination usually answers with SIP 488 Not Acceptable Here. An SBC in the path can solve it by transcoding, but every transcoding step costs processing and a little quality. For a permanent fix, align the PBX's codec list with the carrier's.

Guides

GUIDE

Your 0800 went down and nobody told you: how to monitor 0800 and IVR

An 0800 (toll-free) number can be down all night with every dashboard green, because the failure happens before the call reaches your equipment. This guide shows why that happens, how to test the 0800 and the IVR from outside, what the SAC rules require and what to do when the number won't connect.

By [2N5 team](#) Published October 7, 2026 Updated October 20, 2026

Summary

- An 0800 can fail before it reaches your PBX, on the carrier's route, in number portability or on the IVR platform, and in that case none of your equipment records anything.
- Only a call placed from outside, the way a customer calls, shows whether the 0800 connects and whether the IVR answers with the right message.
- An 0800 that plays the carrier's announcement, or an IVR swapped without notice, is down for the caller even though the call connects.
- Decree 11,034/2022 (Decreto nº 11.034/2022) applies until November 1, 2026 and requires SAC access 24 hours a day, 7 days a week, through at least one channel (not necessarily the phone), and mandatory phone support, with at least 8 hours a day of live-agent support. From November 2, 2026, Decree 13,135/2026 applies, which revokes it and requires phone support with live agents from 8 a.m. to 8 p.m. and an electronic channel available 24 hours.
- For telecom providers, Anatel's RGC (Anatel Resolution No. 765/2023) requires free phone support available 24 hours a day.

How does an 0800 stay down all night without anyone noticing?

An 0800 stays down all night with no alarm when the failure happens before the call reaches your equipment. The PBX receives nothing, so it has nothing to complain about; the SIP trunk stays registered; and in the small hours, few calls is normal. Only the callers notice, and callers give up.

One night, hour by hour

The timeline below is an illustration, built to show the mechanism, not a real case. The times and the cause are hypothetical; the pattern is what matters.

Time	What happens	What the dashboards show	What the caller hears
01:50	A change at the carrier breaks the 0800's route to your trunk	All green	Nothing yet
02:00	First call after the change	No calls on the PBX; SIP trunk registered	Silence, then the call drops

Time	What happens	What the dashboards show	What the caller hears
02:00 to 07:59	Customers try and give up	Low volume, normal for the small hours	Silence
08:00	Live-agent support opens	Empty queues, idle agents	Silence
09:00	A supervisor finds the empty queue odd and calls from their own mobile	Nothing has changed	The supervisor hears silence
09:30	A ticket is opened with the carrier, with no start time and no error code	Nothing has changed	Silence

The customer who called at 2:00 a.m. left no trace on any of the company's equipment. The outage started at 01:50 and was noticed at 09:00: 7 hours and 10 minutes. And the ticket opened at 9:30 starts weak, because nobody can say when the failure began or what the route was responding.

On the same night, with a test call every 15 minutes at Monitor's default pace, the 02:00 test would fail, so would the retries at 02:02, 02:04 and 02:06, and at 02:06 an incident would open with the SIP code of every call, alerting the first tier of the escalation ladder by email. The difference is 6 hours and 54 minutes, from 02:06 to 09:00.

Why the dashboards stayed green

Each dashboard was answering a different question from the one that mattered:

- **the PBX** said it was up, and it was; it has no way of knowing about calls the carrier stopped delivering;
- **the SIP trunk** showed as registered, and it was; registration tells the carrier where to deliver calls, but the 0800 translation happens at another point in its network;
- **the volume chart** showed few calls, which is normal in the small hours.

None of these measurements is wrong. They just don't call the 0800. The [voice circuit monitoring](#) guide compares what each method can see.

0800 down, 0800 not working, 0800 not connecting: is it the same thing?

For the caller it is: the customer does not reach the company. For whoever investigates, they are different symptoms. "Down" and "not working" describe what the customer perceives: the number does not ring, drops, gives a busy signal or plays the wrong message. "Not connecting" describes what the signaling shows: the call was rejected or never reached its destination. An 0800 also "does not work" when it connects and plays the carrier's announcement or a swapped IVR.

"Not working from a mobile phone" is a case of its own: the number answers from some origins and not from others. The causes are in [why an 0800 won't connect from a mobile phone](#). In every case the first step is the same: call from outside, from more than one origin, and write down the time, what you heard and the SIP code.

Why does an 0800 stop connecting?

An 0800 stops connecting because of failures anywhere between the caller and whoever answers: the carrier's route, number portability, the IVR platform, the SIP trunk behind it or the message itself, swapped without notice. Almost all of these points sit outside your PBX and raise no alarm on it.

The carrier's route

An **0800** is not a geographic number. It is free for the caller, as Anatel Resolution 749/2022 sets out in art. 18, and the carrier translates it into a destination: your trunk, a **DID** or the platform that hosts the IVR. That translation is a setting in the carrier's network. A botched destination change, a mistake during maintenance or an interconnection problem between networks breaks the route without anything changing on your side.

The route can also vary with the origin. A mobile call and a landline call, or calls from different carriers, can take different paths to the same 0800, and a failure can affect only one of them. The guide [why an 0800 won't connect from a mobile phone](#) covers that case.

Number portability

When the 0800 or the destination number moves to another carrier, every network has to start routing calls to the new one. Until that happens everywhere, the number works from some origins and fails from others. It is a partial, treacherous failure: whoever tests from their own phone may be on an origin that has already been updated, and conclude that everything is fine.

The IVR platform

The **IVR** can run on the PBX, on a contact center platform or in the cloud. When that platform goes down, runs out of channels or gets a flow published with an error, the call rings with nobody answering, is rejected, or answers and drops. If the platform belongs to a third party, your PBX takes no part in the call and does not see the failure.

The SIP trunk behind the 0800

Behind the IVR or the 0800's destination there is a SIP trunk, and it has its own failure points: exhausted channels, which usually show up as **SIP 503** with Q.850 cause 34 (no circuit/channel available); a lost **SIP registration**; a faulty **SBC**; or NAT and a **SIP ALG** that let the signaling through and block the audio, which produces a silent call. The guide [silent calls and one-way audio](#) covers that last case in detail.

The greeting swapped without notice

A new version of the IVR flow, a holiday recording left switched on, a forgotten test message, or a route change that sends the 0800 to another customer's IVR on the platform. The call connects and, as far as the signaling is concerned, everything worked. The caller, however, hears the wrong message, and for them the 0800 is down.

Early media: a recorded message instead of ringing

Before answering, a route can send audio: this is **early media**, described in RFC 3960, which usually arrives with the **183 Session Progress** response. It is meant for playing the ringback tone, but it is also used for carrier announcements such as "this number does not exist" or "this number is temporarily unavailable".

To the caller, it feels like a call that went somewhere: a voice was heard. But your company never answered, and your PBX never saw the call. Monitoring that only checks whether there was audio would pass this test. That is why testing an 0800 has to compare what was heard with the expected greeting, not just detect sound.

How do I test the 0800 and the IVR from outside?

Test the 0800 from outside by calling the way a customer calls: from another carrier, from a mobile phone and a landline, and at the hours when nobody on the team is watching. For the IVR, check that the greeting that answers is the expected one. A test placed from inside your own network does not go through the carrier's route and cannot see the carrier's failure.

The manual test

A good manual test always follows the same steps, so the result can serve as evidence:

1. **Call from at least two origins:** a mobile phone and a landline, ideally on carriers other than the one that serves the 0800.
2. **Write down the time to the second** and the calling number of each attempt.
3. **Listen to what happens before the answer:** silence, ringback, a busy signal or a carrier announcement.
4. **Check the greeting:** is it your company's, the current version, with the right first menu?
5. **Test the IVR menu:** press an option from the first menu and see whether the IVR recognizes the tone. [DTMF](#) usually travels inside RTP, as RFC 4733 describes, and a negotiation failure makes the IVR ignore the keys even while voice gets through.
6. **Repeat outside business hours** at least once, because the IVR may switch to a different flow at night.

The manual test has an obvious limit: nobody calls at 3:00 a.m. on a Sunday, every 15 minutes. It is for investigating and for checking after a change, not for keeping watch.

The automated test

The automated test is a scheduled test call, placed from a platform outside your network, that calls the number and judges the result on its own. In Monitor, calls go out from the 2N5 platform to support numbers, 0800, 0300, 4004 and landline numbers, and each number has its own schedule, at intervals or at fixed times, including nights and weekends.

During setup, 2N5 records each number's reference greeting. Every test call compares what it heard with that recording: a different message, such as a carrier saying the number does not exist, does not pass. For an IVR, busy, no answer or ringing until timeout also counts as a failure, because it should answer every call on its own. The same comparison measures voice quality, [MOS-LQO](#), a score from 1 to 5 in the terminology of ITU-T P.800.1.

Two rules protect people. The audio of a test call is kept only when what answers is the recording or the IVR on file, encrypted and kept for the period 2N5 sets; a person's voice is never kept. And if the circuit could be answered by a person, they hear a message saying it is an automated test, and the call ends.

On cost: the test calls are part of the service. For an 0800, as for any inbound call, the termination rate follows your contract with the carrier. See [0800 and IVR monitoring](#).

What a test call does not see

A call that hears the greeting and hangs up does not navigate the menu. It does not check whether each option leads to the right queue, whether DTMF is recognized or whether live agents are available at that hour. Those points need the manual test after every flow change.

And a test that always starts from the same origin cannot see a failure that applies only to another one, such as an 0800 that won't connect from a mobile phone on a particular carrier. The automated test watches the whole path, around the clock; the manual test from other origins completes the picture.

What do the SAC decree and Anatel require of an 0800?

Until November 1, 2026, Decree 11,034/2022 (Decreto nº 11.034/2022) requires the SAC, the customer-service channel Brazilian law regulates, to be reachable 24 hours a day, 7 days a week, through at least one channel (not necessarily the phone), and makes phone support mandatory, with at least 8 hours a day of live-agent support. From November 2, 2026, Decree 13,135/2026 revokes 11,034 and keeps phone support mandatory, with live agents from 8 a.m. to 8 p.m. every day (art. 14), and requires at least one electronic channel available 24 hours (art. 12). Neither decree says the 0800 must run 24 hours. For telecom providers, Anatel's RGC requires free phone support available 24 hours a day. Measuring the 0800's availability is how you show those rules were met.

Rule	What it requires
Decree 11,034/2022 (until 11/1/2026), art. 4	SAC access 24 hours a day, 7 days a week, through at least one channel
Decree 11,034/2022 (until 11/1/2026), art. 4, § 2	Phone support is mandatory
Decree 11,034/2022 (until 11/1/2026), art. 5, I	At least 8 hours a day of live-agent phone support
Decree 11,034/2022 (until 11/1/2026), art. 5, II	Complaint and cancellation options in the first menu
Decree 11,034/2022 (until 11/1/2026), art. 12, § 3	Call recordings kept for at least 90 days
Decree 13,135/2026 (from 11/2/2026), art. 14	Phone support is mandatory; live agents from 8 a.m. to 8 p.m., every day
Decree 13,135/2026 (from 11/2/2026), art. 12	At least one electronic channel, 24 hours a day, 7 days a week
Decree 13,135/2026 (from 11/2/2026), art. 15, I	Wait of up to 20 seconds for 80% of the demands and 300 seconds for the rest, unless the regulator has a rule
Decree 13,135/2026 (from 11/2/2026), art. 19	Call recordings kept for at least 90 days
Anatel, RGC (Res. 765/2023), art. 15	Free phone support available 24 hours a day; a live agent 24 hours for urgent requests and from 6:00 a.m. to 10:00 p.m. for everything else
Anatel, RGC (Res. 765/2023), art. 18	First menu with automatic cancellation, complaints and live-agent support
Anatel, Res. 749/2022, art. 18	0800 free for the caller; 0300 and 0303 with shared cost

Decree 11,034 has applied to SACs in regulated sectors since October 3, 2022, until November 1, 2026, and Decree 13,135 applies from November 2, 2026. The RGC has applied to telecom providers since September 1, 2025.

The rules say what has to be available; showing that it was is the job of whoever runs the 0800. A test call every 15 minutes, with each result kept, is an hour-by-hour record that the number connected and the IVR answered with the expected greeting. Monitor's [SLA reports](#) show each number's availability over the year and the month, with every failure and the moment someone took ownership. The content of the first menu, for its part, is checked by whoever listens to the IVR after each change.

What does the caller hear on each 0800 failure, and what does the test record?

Each 0800 failure sounds different to the caller and leaves a different trace on a real test call. The table below links the likely cause, what the customer hears and what the test records, so that your first message to the carrier already says where to look.

Failure	What the caller hears	What a real test call records
The carrier cannot find the 0800's destination	Silence and a drop, or the number-does-not-exist announcement	SIP 404 , or early media with the carrier's announcement, which does not match the reference
The route rejects the call	Silence, an error tone or an announcement	SIP 403 or SIP 603 , with the Q.850 cause when one is sent
The trunk behind the 0800 has no channels	A busy signal or a congestion announcement	SIP 503 , for example with Q.850 cause 34
The IVR platform is down	Rings without answering, or drops	SIP 480 or SIP 408 ; ringing until timeout is a failure for an IVR
The IVR is busy	A busy signal	SIP 486 : a failure for an IVR
Incompatible codecs on the route	The call does not connect	SIP 488
The greeting swapped without notice	A different message: a holiday one, an old one or another company's	Call answered, but the greeting does not match the reference recording: failure
Silent call	It answers and stays silent	Call answered without the expected greeting: failure
One-way audio	The caller hears but is not heard	On a circuit with a test tone or echo, the tones sent do not come back: failure
Choppy voice	The greeting with gaps and noise	MOS-LQO below 3.0: the call counts as degraded
The IVR does not recognize the keys	The IVR repeats the menu or says the option is invalid	Does not show up on a call that only listens to the greeting: needs the manual test

Two readings apply to every row. An 0800 that rings without answering and a route that cancels the call both look like no answer to the tester, sometimes with [SIP 487](#) when the test itself gave up waiting; a [SIP 180](#) followed by nothing is not a success. And the same SIP response can come with different Q.850 causes, which the carrier uses to find the point in its network. The [list of Q.850 causes](#) and the [SIP code lookup](#) help read each one.

What should I do when the 0800 goes down?

When the 0800 goes down, confirm the failure from another origin, find out whether calls are reaching your equipment, gather the evidence and open the ticket with the carrier using it. While the carrier works on it, tell the support team and reinforce the other SAC channels. Afterwards, record the start and end of the 0800 outage.

1. **Confirm from outside.** Call from a mobile phone and a landline, on different carriers. If one origin connects and another doesn't, note which: that points to the route or to portability.
2. **See whether calls reach you.** On the SBC or the PBX, look for the INVITEs of the attempts at the time they were made. If none arrived, the failure is upstream of you, at the carrier or the IVR platform. If they arrived and were rejected, the failure is on your side.

3. **Gather the evidence.** The number called, the origins, the time of each attempt with seconds and time zone, the SIP code, the Q.850 cause when one is sent, the Call-ID and a recording of the message that played, if any did. Keep the last call that worked as well.
4. **Open the ticket with everything in the first message.** Ask for the ticket number and ask at what point on the route the call stops. A ticket with a time and an error code is harder to close with “no fault found”.
5. **Protect customer service.** Tell the SAC team and the supervisors, and reinforce the other support channels while the 0800 is down. If the carrier contract provides an alternate destination for the 0800, ask for the diversion.
6. **Close with numbers.** Record the start and end of the outage, because that is what goes into the month’s and the year’s SLA. The guide [telephony SLA: yearly and monthly calculation](#) shows the math. And ask why your team heard about it from a customer and not from an alert.

In Monitor, steps 1 and 3 happen on their own. Four failures in a row, by default (the first one and three retries), open an incident with a block ready to paste into the ticket: the circuit, the carrier, the first symptom, the SIP code and Q.850 cause of each call and the last healthy call. The escalation ladder alerts the first tier by email and climbs until someone takes ownership with one tap on the email link. When two good calls in a row confirm recovery, the incident closes and the customer’s contacts are notified. See [alerts and escalation](#).

One caveat applies to any monitoring: if the problem is on the tester’s side, the test says nothing about the 0800. In Monitor, a test that could not run because of a problem on 2N5’s side is marked monitoring blind, does not count against the number, stays out of the SLA and does not wake your team.

How often should I test the 0800?

Test the 0800 as often as the availability target demands and at the hours when nobody is watching. With one call every 15 minutes and retries every 2 minutes, an overnight 0800 outage becomes an incident within 21 minutes; with one call an hour, a whole hour can pass before the first sign.

The worst-case math: an outage that starts right after a good test takes up to 15 minutes to be seen, plus 6 minutes for the three retries to confirm it, or $15 + 6 = 21$ minutes to the incident. At the 15-minute pace, each number gets $24 \times 4 = 96$ test calls a day, each holding a channel for a few seconds, within a call limit that never competes with customers for a channel.

Three things help when choosing the schedule:

- **Nights and weekends count more.** That is when the 0800 is on its own. A 24-hour 0800 needs 24-hour testing.
- **Fixed times complement the interval.** A test at a set time, for example just before live-agent support opens, confirms the number at the moment the most customers are about to call.
- **Each support number is a circuit.** If the company has an 0800 for customers, a 4004 and a collections number, each has its own route and its own IVR, and each needs its own test. The [voice circuit monitoring](#) guide shows how to choose what to monitor.

The goal is simple to state: the next time the 0800 goes down at 2:00 a.m., the first person to know should be someone on your team, at 2:04 a.m., with the SIP code in hand, not a customer at 9:00 a.m.

Frequently asked questions

Why won't my 0800 connect?

The failure can be outside your PBX: the carrier stopped routing the 0800 to your destination, a number port was not updated on every network, the IVR platform went down or the SIP trunk behind it ran out of channels. The SIP code and Q.850 cause of a test call show where the call stopped.

How do I test an 0800 in the middle of the night without anyone having to call?

With scheduled test calls, placed from outside your network, that call the 0800 like a customer and check what answered. In Monitor, each number has its own schedule, at intervals or at fixed times, including nights and weekends, and each call compares the greeting it heard with the reference recording.

The 0800 connected, but it plays the carrier's announcement. Is it down?

For the caller, yes: the customer never reaches your company. That announcement usually arrives as early media, before any answer, or on a call answered by the carrier itself. Monitoring that only checks whether there was audio lets this through; you have to compare what was heard with the IVR's expected greeting.

What does the SAC decree require of an 0800?

Until November 1, 2026, Decree 11,034/2022 requires SAC access 24 hours a day through at least one channel, mandatory phone support, at least 8 hours a day of live-agent support and complaint and cancellation options in the first menu. From November 2, Decree 13,135/2026 revokes it and requires phone support with live agents from 8 a.m. to 8 p.m. and an electronic channel available 24 hours. Neither one mentions the 0800.

What should I do first when the 0800 goes down?

Confirm it from another origin, calling from a mobile phone and a landline on different carriers. Then check whether the calls reach your SBC or PBX: if none do, the failure is upstream of you. Gather the time, the SIP code, the Q.850 cause and the Call-ID of each attempt and open the ticket with the carrier using that evidence.

Is my 0800 down, or does it just not work from mobile phones?

Test from a landline and from mobile phones on at least two carriers, at the same time. If all of them fail, the number is down for callers, and the problem is usually the carrier's route, number portability, the IVR or the trunk. If only mobile phones fail, suspect an origin restriction or a mobile route; the guide on why an 0800 won't connect from a mobile phone lists the causes.

Sources

1. Decree No. 11,034/2022 (Decreto nº 11.034/2022): SAC rules, including 24-hour access and phone support (applies until November 1, 2026)
2. Decree No. 13,135/2026, arts. 12, 14, 15 and 19: the new SAC decree, which revokes 11,034 and applies from November 2, 2026
3. Anatel, Resolution No. 765/2023 (RGC), arts. 15 and 18: phone support and the first menu

4. Anatel, Resolution No. 749/2022, art. 18: 0800 free for the caller, 0300 and 0303 with shared cost
5. RFC 3261: SIP and its response codes
6. RFC 3960: early media, the audio before the answer
7. RFC 4733: DTMF carried in RTP
8. RFC 6432: Q.850 causes in the Reason header of SIP responses
9. ITU-T P.800.1: MOS terminology, including MOS-LQO

GUIDE

Failing calls: how to diagnose them by SIP code and Q.850 cause

When the calls on a SIP trunk or an 0800 start to fail, the SIP code and Q.850 cause of the first failed call say where it stopped and whose next step it is. This guide gathers the codes that show up most, 503, 486, 480, 487 and 603, and the causes 16, 34 and 127, with the right question to ask for each.

By [2N5 team](#) Published October 14, 2026 Updated October 14, 2026

Summary

- The SIP code says what the server decided to do with the call; the Q.850 cause, in the Reason header, says why, and it is almost always what the carrier uses to locate the fault.
- SIP 503 with cause 34 is a lack of channels or routes, and SIP 486 with cause 17 is usually the simultaneous channel limit being exhausted.
- SIP 480 is a destination that is unregistered or unavailable, SIP 487 is a call canceled before it was answered and SIP 603 is an explicit decline.
- Cause 16 is normal clearing and not an error; cause 127 is a reason the network did not report.
- To open a ticket with the carrier, gather the time with time zone, the number called, the SIP code, the Q.850 cause and the Call-ID of each attempt.

Where do you start when calls fail?

With the code of the first call that failed, not with the complaint. Each rejected call returns a SIP code, and many also return a Q.850 cause in the Reason header ([RFC 6432](#)). Together they say where the call stopped and who can fix it.

A short path, in order:

1. **Gather the attempts:** time with time zone, origin number, number called, SIP code, Q.850 cause and Call-ID of each one.
2. **See whether the call reached your SBC or PBX.** If none did, the failure is before you, on the carrier's route.
3. **Read the code in the table below** and follow its page for the detail.
4. **Repeat from other origins:** a landline and mobile phones on different carriers show whether the failure applies to all of them.

A scheduled test call helps a lot here: it already brings the time, the code and the cause of each attempt, including the ones nobody saw. See [what is synthetic call monitoring](#).

What does each SIP code mean?

Code	What it means	Most common cause on a trunk or 0800	Next step
503 Service Unavailable	The server is temporarily unable to process the call	Lack of channels or routes at the carrier, or an SBC that is down	Read the Retry-After and the Q.850 cause; open a ticket with the Call-ID if it persists
486 Busy Here	The destination was reached but cannot accept another call	Simultaneous channel limit exhausted	Compare the active calls at that moment with the contracted limit
480 Temporarily Unavailable	The destination was located but is not available now	Destination unregistered, on Do Not Disturb or not answering	Check the destination's registration and the Retry-After
487 Request Terminated	The request was ended by a CANCEL before it was answered	The caller hung up while it rang, or the PBX gave up on a timeout	Not a destination failure; investigate only when excessive
603 Decline	The destination was contacted and explicitly declined the call	The person rejected it, or a call block	See whether the decline came from the destination, the PBX or the carrier

For the other codes, the [SIP code index](#) has a page for each, and the [SIP code lookup](#) opens the code you type.

And the Q.850 causes?

The Q.850 cause is the reason the network gives for ending the call, inherited from traditional telephony and carried by SIP in the Reason header. The three that raise the most questions:

Cause	What it indicates	When to investigate
16, Normal call clearing	Normal clearing: one side hung up after the call was established	Only when many calls end with 16 within a few seconds, with no audio or right after the IVR answers
34, No circuit/channel available	There was no free circuit or channel to complete the call	Always: it is usually the channel limit exhausted, yours or the carrier's
127, Interworking, unspecified	The call crossed an interconnection that does not report the reason	Find out where the call failed and what response came from there

The [Q.850 cause index](#) lists the rest.

How does the SIP code connect to the Q.850 cause?

The two fit together. [SIP 503](#) often appears with cause 34 (no channel available), but also with others, such as 38, 41 and 88. [SIP 486](#) usually comes with cause 17, destination busy. [SIP 603](#) goes with 21, call rejected. [SIP 480](#) comes with causes 19, 20 and 31, and [SIP 487](#) has no associated cause, because it is a cancellation. The mapping between the two families is in [RFC 3398](#).

Which failure patterns deserve attention?

- **Many 503s with cause 34 at a peak hour:** trunk capacity exhausted. Compare the peak of simultaneous calls with the contracted channels.
- **486 and cause 17 on an IVR or 0800:** usually the same channel limit, showing up as “busy”.

- **480 on an extension or an IVR that should answer:** destination unregistered. Check the SIP registration.
- **Many 487s in a row:** the PBX is giving up before the call is answered, or callers hang up while it rings. See how many seconds passed before the CANCEL.
- **Cause 16 on very short calls:** the clearing is normal, but the problem may be earlier. See who sent the BYE and whether there was audio both ways.
- **Cause 127 on calls that cross a gateway:** the reason was lost in the conversion between networks. Ask the carrier for the originating code.

What should you take to the carrier?

A table, one row per attempt: time with time zone, origin number, number called, SIP code, Q.850 cause, Call-ID and what was heard. That is what the carrier can locate in its own records. The guide [How to enforce a carrier SLA with proof](#) shows how to turn these records into a claim.

How does Monitor by 2N5 use these codes?

Each Monitor by 2N5 test call records the SIP code and Q.850 cause returned, and the four failures in a row that open an incident reach the alert with that evidence. See [SIP trunk monitoring](#) and [0800 and IVR monitoring](#).

Frequently asked questions

Why are my calls failing on the SIP trunk?

Look at the SIP code of the first failure and the Q.850 cause in the Reason header. A 503 with cause 34 points to a lack of channels or routes; a 486, to the channel limit; a 480, to an unregistered destination; a 603, to an explicit decline. Without the code, any diagnosis is a guess. Record the time, the number called and the Call-ID of each attempt.

What does SIP 503 Service Unavailable mean?

SIP 503 Service Unavailable means the server is temporarily unable to process the call, because of overload or maintenance. On trunks, the most common reasons are the carrier running out of channels or routes, or an SBC being down. It matters when it repeats: read the Retry-After and the Q.850 cause and open a ticket with the Call-ID.

Is Q.850 cause 16 an error?

No. Cause 16, Normal call clearing, is the normal end of a call: one side hung up after the call was established. It is the most common value in call logs. It only deserves investigation when many calls end with 16 within a few seconds, with no audio or right after the IVR answers, because then the normal clearing may be a consequence of another problem.

What is the difference between the SIP code and the Q.850 cause?

The SIP code is the protocol's response, such as 503 or 486, and has classes: 4xx is a request failure, 5xx a server failure and 6xx a global failure. The Q.850 cause comes from traditional telephony and gives the reason more precisely, such as 34 for no channel available. RFC 6432 defines how the cause travels in the Reason header.

How do I know whether the problem is the carrier's or my network's?

See whether the call reaches your SBC or PBX. If none does, the failure is before you, on the carrier's route, and the code and time of each test call are the proof for the ticket. If it arrives and your inbound rule rejects it, the problem is in your own configuration.

What does SIP error 503 / 486 / 480 / 408 mean?

503 means service unavailable: on trunks, usually no channel or no route at the carrier, or an SBC that is down. 486 means busy, normally the channel limit. 480 means the destination is temporarily unavailable, such as a number that is not registered. 408 means timeout: the call got no answer, because of the route or a firewall. Always read the code together with the Q.850 cause and the time of each call.

What is Q.850 cause 16 / 34 / 127?

A Q.850 cause is the disconnect reason defined by traditional telephony, carried in the Reason header of the SIP response. Cause 16 is normal call clearing, not an error. Cause 34 is no circuit or channel available. Cause 127 is interworking, unspecified: the call crossed a network that does not report the reason, and the fault can be on any leg.

Sources

1. [RFC 3261: SIP, the response classes and the 4xx, 5xx and 6xx codes](#)
2. [RFC 6432: Q.850 causes in the Reason header of SIP responses](#)
3. [RFC 3398: the mapping between Q.850 causes and SIP response codes](#)
4. [ITU-T Q.850: call clearing cause values](#)

GUIDE

How to enforce a carrier SLA with proof

To claim against a carrier's SLA, the company has to show, outage by outage, what happened and when: the time, the number called, the SIP code, the Q.850 cause and the downtime, with maintenance already subtracted. An outage a customer noticed is hard to prove; one a test recorded is hard to dispute.

By [2N5 team](#) Published October 14, 2026 Updated October 20, 2026

Summary

- The claim starts in the contract: the measurement period, what counts as unavailability, the maintenance window and who measures.
- The proof is a record made at the time, from outside the network, with the time, number called, SIP code and Q.850 cause of each test call.
- Availability is $(\text{period} - \text{maintenance} - \text{downtime}) \div (\text{period} - \text{maintenance}) \times 100$; the SLA calculator does the math for any target.
- At 99.9%, the limit is 8 h 45 min 36 s of downtime per 365-day year, or 43 min 12 s per 30-day month.
- Time when the monitoring itself had no visibility is not circuit downtime and should be left out of the calculation.

How do you enforce a carrier's SLA?

In three parts: what the contract promises, what was measured and the math between the two. Without the first, you do not know what to claim; without the second, there is nothing to show; without the third, the claim is an opinion. An [SLA](#) can only be enforced when someone measured.

What does the contract need to say?

Before accepting a target, agree in writing on:

1. **the measurement period** (year or month) and the time zone;
2. **what counts as unavailability**: only a call that does not connect, or also a silent call and quality below the limit;
3. **how maintenance is announced** and the maximum window;
4. **who measures, and how**: the carrier's measurement, yours, or an independent one;
5. **which evidence counts** to open a dispute, and the deadline to request the credit.

If the contract already exists and does not say this, the first claim also serves to agree on the criteria. The guide [Telephony SLA: yearly and monthly calculation](#) details each point.

What proof should you bring?

Good proof is a record made on the spot, from outside your network, that the carrier can check. For each outage:

Field	What it is for
Start and end time, with time zone	Measures the duration and locates the failure in the carrier's records
Number called and origin number	Says which circuit and which route failed
SIP code of each test call	Shows how the route rejected it, for example SIP 503
Q.850 cause	Gives the reason, such as cause 34
Call-ID	Lets the carrier find the exact call
What was heard	Tells silence, busy and a carrier announcement apart
Duration and result of the following calls	Shows when the circuit came back

The codes and causes are explained in [Failing calls: diagnosis](#). An outage recorded with these fields is hard to dispute; one noticed by a customer who complained is not.

How do you calculate availability?

Availability = (period – scheduled maintenance – downtime) ÷ (period – scheduled maintenance) × 100. The site's SLA calculator does this math for any target and already accounts for maintenance windows: use the [SLA calculator](#).

Two references for the conversation with the carrier, at 99.9%:

Period	Downtime allowed
365-day year	8 h 45 min 36 s
30-day month	43 min 12 s

Each outage enters the calculation by its measured time, and the measured time depends on how often the test runs: with a call every 15 minutes and retests every 2 minutes after a failure, the start and end of each outage are marked within a few minutes.

What stays out of the calculation?

- **Scheduled maintenance**, announced and within the agreed window.
- **Failures caused by the customer**, such as a PBX turned off at a branch, when the contract is with the carrier and the defect is not theirs.
- **Time when the monitoring had no visibility** because of a problem of its own. A test that could not be made says nothing about the circuit.

Putting these exceptions in the report, rather than hiding them, is what makes the number defensible: the carrier sees that the math is honest.

How do you put the claim together, step by step?

1. **Gather the report for the period** (year or month, as the contract says) with the outages and the availability.
2. **List the outages** in a table, one per row, with the fields above.
3. **Subtract the maintenance** that was announced inside the window.
4. **Compare with the target** and work out the difference in minutes.
5. **Open the ticket** with the table and the report, and ask the carrier to check the routes at the times indicated.

6. **Keep the reports** month by month. The last-12-months figure is the one that closes the year.

The credit, penalty or termination is whatever the contract provides; so are the way to request it and the deadline. That is a matter for your legal team, and this guide is not legal advice.

What does Monitor by 2N5 deliver?

Monitor by 2N5 calls each voice circuit from Brazil, on its own schedule, and keeps the result of each test call. Every month it emails an [SLA report](#) with the availability for the month and the last 12 months, each incident with its duration, the total downtime and the availability by circuit and by carrier. A period when the test could not be made because of a problem on 2N5's side is left out of the calculation and shown separately. For a consumer-service obligation, see also the [SAC decree and RGC](#).

Frequently asked questions

How do I claim against a carrier's SLA?

Open the ticket with an availability report and the list of outages: time and time zone of each failure, number called, SIP code, Q.850 cause and duration. Add up the downtime, subtract the announced maintenance and compare with the contract target. First check whether the contract defines how to dispute and the deadline to request the credit.

What proof counts in a dispute with the carrier?

What the carrier can locate in its own records: time with time zone, origin and destination numbers, SIP code, Q.850 cause and Call-ID of each attempt. A record made at the time, by test calls from outside the network, is harder to dispute than the recollection that the number was down.

How much downtime is 99.9% availability?

In a 365-day year, 8 hours, 45 minutes and 36 seconds. In a 30-day month, 43 minutes and 12 seconds. In a week, 10 minutes and 5 seconds. In a day, 1 minute and 26 seconds. The site's SLA calculator shows the value for other targets and includes maintenance windows.

Does scheduled maintenance count as an outage?

Normally not, if it was announced with the notice the contract requires and stayed within the agreed window: it comes out of the period before the calculation. What goes past the window counts as an outage from the moment the window ended. That is why the window must be recorded before it starts.

What if the carrier disputes the measurement?

Show the sequence of test calls for the period, with the result of each, and ask the carrier to check the routes at the times indicated. If the contract says who measures, use that measurement as the reference and yours as supporting evidence. Questions about how to read the contract are for your legal team.

Sources

1. [ITU-T E.800: quality of service definitions, including availability](#)
2. [RFC 3261: SIP, the response codes kept in each call's record](#)

3. RFC 6432: Q.850 causes in the Reason header of SIP responses
4. Anatel, Resolution No. 765/2023 (RGC), art. 15: phone support available 24 hours

GUIDE

Line-down alerts without false alarms

A line-down alert is only useful if it is true: someone woken at 2 a.m. by a false alarm soon stops looking. In Monitor by 2N5, a call that fails is retried before anyone is told, only four failures in a row open an incident, and the alert arrives by email, climbing layer by layer until someone confirms.

By 2N5 team Published October 14, 2026 Updated October 20, 2026

Summary

- A false alarm is an alert for a failure that did not exist, and it is the main reason a team stops trusting monitoring.
- The defense is the retest: a failed call is repeated, and an incident only opens after four failures in a row, the first and three retests.
- The alert arrives by email, with the evidence and a link; one tap on the link confirms that someone took over.
- If nobody confirms in the agreed time, the alert climbs to the next layer, and then to the company's administrators.
- The email carries the time, the SIP code and the Q.850 cause of each call, so the team can act without opening the system.

How do you get line-down alerts without false alarms?

With three defenses in sequence: redo the call that failed, require consecutive failures before opening the incident, and only then notify, by layers, until someone confirms. Each defense cuts one kind of false alarm, and together they make a notice on your phone at 2 a.m. mean the circuit really went down.

What is a false alarm in telephony?

A false alarm is a notice about a failure that did not exist, or that was already over when the notice arrived. In telephony, the most common causes are:

- **a test call that failed once** because of a hiccup on the route and worked on the next one;
- **a problem on the monitoring side**, which cannot test but blames the circuit;
- **scheduled maintenance** that nobody told the monitoring about;
- **a threshold too tight** for the network's normal fluctuation.

The cost is trust. When many alerts are false, the team starts ignoring all of them, and the true one gets lost among them.

Why test again before alerting?

Because a single failure almost never says enough. Monitor by 2N5 calls each circuit, by default, every 15 minutes. When a call fails, it is repeated three times, 2 minutes apart. The incident opens with four failures in a row, the first and three retests; two good calls in a row confirm the recovery.

Moment	What happens
The test call fails	Nobody is told; the test is repeated in 2 minutes

Moment	What happens
Second and third failures	Still nobody is told
Fourth failure in a row	The incident opens, with the evidence of each call: time, SIP code and Q.850 cause
A good call in between	The count starts over; the blip does not become an alert
Two good calls in a row	The circuit is back, and the incident closes

The math of the wait: an outage that starts right after a good test is seen 15 minutes later at most, and the fourth failure arrives 6 minutes after that (2 + 2 + 2). In exchange for those minutes, the alert becomes reliable. When the failure is on 2N5's side, the test is left without visibility, does not count against the circuit and never wakes your team.

How does the alert arrive?

By email, with the evidence, the codes and the incident's numbers, and a link that confirms the person took over the incident, without logging in to the system.

Each person can have more than one email address within the layer; if one is not available, the message goes to the next, without waiting.

How does escalation work?

The alert climbs in layers until someone takes over. Layer 1 is told when the incident opens. If nobody confirms in the time set for the layer, layer 2 is told, then layer 3. A single confirmation, with one tap on the message's link, stops the climb, and the incident shows who took over and when.

Who	When they are told
Layer 1	When the incident opens
Following layers	When the previous one does not confirm in the defined time
Company administrators	When no layer confirms
Customer contacts	When a service has a problem and when it is resolved; they are not part of the climb

The time to the first confirmation feeds the [SLA report](#) as average response, which shows how long the team takes to take over. The glossary explains [MTTR](#) and [MTTA](#).

How do you avoid alerts during maintenance?

With a [maintenance window](#): the chosen circuits, at a set time, stop opening incidents. The window must exist before it starts; maintenance declared after the outage is just an outage. Anything beyond the agreed time counts again.

What else reduces false alarms?

- **Measure the call, not the equipment.** A registered trunk does not prove that the call connects; a false green is the other side of a false alarm.
- **Compare the IVR with the reference recording**, so that the carrier's announcement does not pass as an answer.
- **Have someone per layer**, not just one person: a ladder with a single layer becomes a false silence when that person has no signal.
- **Record maintenance beforehand.**

For the full path, see [alerts and escalation](#) and [how it works](#). When the incident opens, the guide [Failing calls: diagnosis](#) helps read the codes, and [How to enforce a carrier SLA with proof](#) shows what to do with the record.

Frequently asked questions

How do I get an alert when a phone line goes down?

Add the person's email as a contact in Monitor by 2N5. When an incident is confirmed, the message arrives with the evidence and a link, and one tap confirms that you took over.

What is a retest before alerting?

It is repeating the call that failed, after a short interval, before telling anyone. A one-minute blip, a route that failed once or a lost call do not become an alert. In Monitor by 2N5, the call is repeated three times, 2 minutes apart, and only four failures in a row open the incident.

How long passes between the outage and the alert?

At the default pace, an outage is seen on the next test call, within 15 minutes, and the three retests add 6 more minutes until the fourth failure. The notice goes out right after the incident opens. The interval between tests and the number of attempts exist to trade a few minutes for a reliable alert.

What if nobody answers the alert?

It climbs: if the first layer does not confirm in the defined time, the second is told, then the third. When no layer confirms, the company's administrators are told. A single confirmation, with one tap on the link, stops the climb and records who took over and when.

Can I silence alerts during maintenance?

Yes. A maintenance window covers the chosen circuits at a set time, and failures on those circuits during the window do not open an incident. The window must be recorded before it starts; anything beyond the agreed time counts as an outage again.

Sources

1. [RFC 3261: SIP, the response codes the test call records](#)
2. [ITU-T E.800: quality of service definitions, including availability](#)
3. [RFC 6432: Q.850 causes in the Reason header of SIP responses](#)

GUIDE

Passive monitoring or active test calls: which to choose for telephony

Passive monitoring analyzes the real calls that cross your network; active test calls dial on their own, from outside, and check the result. Passive explains what happened to the people who called; active warns you when nobody can call, including when nobody tried. Most of the time you want both, and this guide shows in what order.

By [2N5 team](#) Published October 14, 2026 Updated October 14, 2026

Summary

- Passive monitoring captures or reads the SIP and RTP traffic of real calls and computes indicators such as estimated MOS, jitter and packet loss.
- Active test calls leave from outside the network, on a schedule, and check whether the number connects, who answers and how the audio sounds.
- Passive cannot see what never crossed your network: a fault on the carrier's route that stops the call from arriving leaves no record to analyze.
- Active cannot see the reason inside your network: it shows that the call failed and with which code, but it does not replace capture for investigating the cause.
- For an 0800 or an IVR, start with the active test; for a busy trunk with quality complaints, add passive.

What is the difference between passive monitoring and active test calls?

In passive monitoring, the tool stays out of the path and watches the calls customers made; in active test calls, the tool dials. The difference changes what each one can see. Passive sees everything that crossed your network and nothing that did not; active sees the result of a simulated customer call, at any hour, but only for the calls it places.

	Passive monitoring	Active test calls
What it does	Captures or reads the traffic of real calls	Dials the numbers, on a schedule
Example tools	VoIPmonitor (SIP and RTP capture); PRTG and Zabbix measure the network and the PBX with sensors, without test calls	Monitor by 2N5
Needs real traffic?	Yes: with no call, there is nothing to analyze	No: it generates its own calls
Sees a fault on the carrier route?	Only if some call gets recorded	Yes, it sees the call that does not connect
Sees a swapped IVR or the carrier's announcement?	No, unless someone listens	Yes, it compares with the reference recording
Sees the cause inside your network?	Yes, with the call capture	Only the SIP code and Q.850 cause returned
Installation	In your network (port mirroring, SBC mirroring or PBX reading)	None: it dials from outside

For the tools in the table, the facts come from the vendors' own pages, listed in the sources:

[VoIPmonitor](#) does passive SIP and RTP capture, with MOS, one-way audio detection and fraud rules, and places no calls; [PRTG](#) uses SIP OPTIONS, QoS and Cisco IP SLA sensors, with MOS estimated from the network; [Zabbix](#) reads the trunks, channels and queues of Asterisk. On the pages we checked, none of the three dials out to the public network to check an 0800 or an IVR.

When is passive monitoring the best choice?

When you need to understand calls that already happened. With capture, you can open a specific call, see the signaling and the audio, and find out whether the delay was in your SBC or at the carrier.

Passive is the right tool to:

- investigate a customer's complaint about a specific call;
- follow real traffic indicators, such as the completion rate ([ASR and ACD](#)) and [post-dial delay](#);
- detect fraud patterns and abnormal trunk use;
- measure jitter and packet loss on the calls that did go through.

When are active test calls the best choice?

When the risk is the number going down without anyone noticing. Active calls serve best for:

- **0800, 0300 and 4004 numbers**, whose problem is often before your network, in the route or in number portability, where none of your equipment looks;
- **IVRs**, to check that the right message answers and not the carrier's;
- **nights, weekends and rarely dialed numbers**, where real traffic is low or nonexistent;
- **measuring the carrier's SLA** with an independent record of each hour;
- **SIP trunks** that need to be checked from outside, and not only by registration.

An 0800 that stopped connecting at 2 a.m. does not show up in passive monitoring: there was no call to capture. The guide [Your 0800 went down and nobody told you](#) shows that case.

Can you use both?

You can, and it is the most common setup in larger operations: active warns that something stopped, with the time and the code; passive helps find out why. The tools do not compete for the same job. A reasonable order:

1. **Start with active** on the numbers that earn money or that the rules require to be answered (SAC, sales 0800, support IVR).
2. **Add passive** when there is traffic volume and quality questions that only the analysis of real calls can settle.
3. **Use the codes** of the failed test calls as the starting point. The guide [Failing calls: diagnosis](#) shows what each one means.

What does Monitor by 2N5 do, and not do?

Monitor by 2N5 is about active test calls: it dials from Brazil to 0800 numbers, IVRs and SIP trunks, on a schedule of its own for each circuit, checks what answered, measures the voice quality of the audio received (MOS-LQO) and opens an incident after four failures in a row, the first and three retests. It

does not capture your network's traffic and does not replace a capture tool for finding the cause inside it. See the comparisons with [VoIPmonitor](#), [PRTG](#) and [Zabbix](#), and the definition in [what is synthetic call monitoring](#).

Frequently asked questions

What is the difference between passive and active VoIP monitoring?

Passive monitoring observes the calls customers already made, by traffic capture or by reading logs, and computes indicators about them. Active monitoring places its own test calls, on a schedule, from outside the network. Passive explains what happened on existing calls; active detects the failure even when no real call is crossing.

Does VoIPmonitor place test calls?

No. According to the vendor's site, VoIPmonitor does passive SIP and RTP capture and computes MOS, detects one-way audio and applies fraud rules; it analyzes traffic and does not generate calls. It is a very good tool for investigating real calls, and a different category from dialing the 0800 and checking the IVR.

Does PRTG place test calls to my 0800?

From the vendor's pages, PRTG measures with sensors: SIP OPTIONS ping, QoS round trip and Cisco IP SLA, with MOS estimated from the network. We found no dialing to the public network, 0800 or IVR. It shows the health of the network and the trunk; checking what answers on the 0800 is another task.

Does Zabbix monitor SIP trunks and call quality?

Zabbix, which is open source, has an Asterisk template that reads the state of trunks and peers, channels and queues. From what the project publishes, it does not place test calls or assess audio. It is for looking at the PBX from inside, not for knowing whether a call from outside connects.

Do I need both kinds of monitoring?

It depends on the risk. If what matters is knowing that the 0800, the IVR or the trunk answer at any hour, start with active. If there is heavy traffic and quality complaints, passive shows the problem calls. Larger operations run both: active warns, passive helps find the cause.

Does PRTG / Zabbix place test calls?

Not according to the vendors' own pages. PRTG measures with sensors, such as SIP OPTIONS ping, QoS round trip and Cisco IP SLA, and Zabbix reads trunk, channel and queue status through its Asterisk template. Neither one dials the public network or checks what answers. For that you need an active test call, which dials from outside and listens to the result.

Sources

1. [VoIPmonitor: passive SIP and RTP capture, MOS and one-way audio detection](#)
2. [Paessler PRTG: sensor-based VoIP monitoring \(SIP OPTIONS, QoS, Cisco IP SLA\)](#)
3. [Zabbix: Asterisk integration \(trunk, channel and queue state\)](#)
4. [RFC 3261: SIP, the protocol both approaches observe](#)

GUIDE

The SAC decree and Anatel's RGC: how to prove your 0800 and IVR were up

The SAC decrees (Decree 11,034/2022, in force until November 1, 2026, and Decree 13,135/2026, which revokes it and applies from November 2, 2026) and Anatel's General Consumer Rights Regulation say what phone support has to offer, but not how to prove it was available. This guide summarizes the articles that cover the phone channel and shows how test calls and an availability report document the whole year.

By [2N5 team](#) Published October 7, 2026 Updated October 14, 2026

Summary

- Until November 1, 2026, Decree 11,034/2022 (Decreto nº 11.034/2022), in force since October 3, 2022 (art. 18), applies: SAC access, the customer-service channel Brazilian law regulates, 24 hours a day, 7 days a week, through at least one channel, and mandatory phone support (art. 4), with at least 8 hours a day of live-agent support (art. 5, I).
- From November 2, 2026, Decree 13,135/2026 applies. It revokes 11,034 (art. 24) and takes effect twenty days after its October 13, 2026 publication (art. 25). It keeps phone support mandatory (art. 14), with live-agent support from 8 a.m. to 8 p.m. every day, and requires at least one electronic channel available 24 hours (art. 12).
- Decree 11,034 does not set a maximum wait time: it tells the regulators to set it (art. 5, III). Decree 13,135 sets up to 20 seconds for at least 80% of the demands and up to 300 seconds for the rest, unless the regulator has a specific rule (art. 15, I and § 1).
- For telecom providers, Anatel's RGC (Anatel Resolution No. 765/2023), in force since September 1, 2025, requires free phone support available 24 hours a day (art. 15).
- Neither decree mentions the 0800, and no rule says how to prove availability; a record of test calls placed from outside, with the time and result of each one, is one way to document it.
- This guide is not legal advice: how each rule applies to your company should be confirmed with your legal team.

What is the SAC decree?

The SAC decree is Decree No. 11,034 of April 5, 2022 (Decreto nº 11.034/2022), which applies until November 1, 2026. From November 2, 2026, the SAC decree is Decree No. 13,135 of October 9, 2026, which revokes 11,034. Both regulate the Código de Defesa do Consumidor (Brazil's Consumer Protection Code) to set the rules for the SAC, the customer-service channel. Decree 11,034 applies to providers of services regulated by the federal Executive Branch, and each regulator applies those rules to its own sector, taking the provider's size into account.

[Art. 1 of Decree 11,034](#) sets the scope: the decree applies "to providers of services regulated by the federal Executive Branch", and its sole paragraph tells regulators to consider "the size of the provider of the regulated service". Art. 2 defines the SAC as support delivered through several integrated channels to handle requests such as information, questions, complaints, disputes, and the suspension or cancellation of contracts and services. Offering and selling products is left out. (Quotations from Brazilian rules in this guide are our English translations of the official Portuguese text.)

One detail changes how it reads: **Decree 11,034 does not list sectors by name.** Whoever defines how it applies to each sector is that sector's regulator (arts. 5, 13 and 14). So "does the decree apply to my company?" has two parts: whether your service is regulated by the federal Executive Branch, and what your regulator's rule says. Decree 13,135 also leaves the criteria to the competent regulator (art. 1, sole paragraph, I) and adds a named group: providers that, through platforms or other digital means, offer or provide financial, payment, credit or other financial-nature services to consumers and are subject to regulation or supervision by the Central Bank of Brazil (art. 1, sole paragraph, II).

When did Decree 11,034 take effect?

Decree 11,034 took effect on October 3, 2022. Art. 18 says it applies "one hundred and eighty days after the date of its publication", and it was published in the *Diário Oficial da União* (the federal official gazette) on April 6, 2022. The same decree revoked Decree No. 6,523/2008 (Decreto nº 6.523/2008), the previous SAC rule.

The revocation is in art. 17. Anyone still citing the 2008 rules in contracts, manuals or IVR scripts is citing a text that no longer applies. The text on the Planalto (presidential) website carries a notice that it does not replace the version published in the *Diário Oficial*; for formal use, the official publication is the reference. Decree 11,034 stops applying when Decree 13,135 takes effect, on November 2, 2026, as the next section describes.

What changes with Decree 13,135, and from when?

Decree No. 13,135, of October 9, 2026, published in the *Diário Oficial da União* on October 13, 2026, regulates the SAC again and, in art. 24, revokes Decree 11,034. Art. 25 makes it take effect twenty days after publication, which, counting the period, is November 2, 2026. So **Decree 11,034 applies until November 1, 2026, and from November 2, 2026 Decree 13,135 applies.** Confirm the date in the text published in the *Diário Oficial*.

What Decree 13,135 says about the phone and the 24-hour channel, with the articles:

- service through phone channels is mandatory (art. 14, main provision);
- live-agent phone service runs from 8 a.m. to 8 p.m., without interruption, seven days a week, and may be preceded by an electronic menu (art. 14, § 1); for urgent demands it is immediate and at any time (art. 14, § 2, and art. 17);
- if live-agent phone service is accidentally interrupted before it ends, the provider must call back within three hundred seconds (art. 14, § 3);
- the maximum wait for an agent is up to twenty seconds for at least 80% of the demands served and up to three hundred seconds for the rest, unless the regulator has a specific rule (art. 15, I and § 1);
- the option to speak to a person must exist at every level of phone service (art. 15), and immediate cancellation and live-agent service must be available even at the first level (art. 16, § 2);
- at least one electronic communication channel must work 24 hours a day, seven days a week, with real-time live-agent service from 8 a.m. to 8 p.m. (art. 12 and § 1); the phone call, defined in art. 3, § 1, is a separate channel;
- access is free on all channels (art. 8);
- call recordings are kept for at least ninety days (art. 19), and the history of each demand for at least six months after resolution (art. 18, § 4, II).

The new text does not mention the 0800 and does not require the phone to work 24 hours: what works 24 hours is an electronic channel, and only urgent demands get live-agent phone service at any time. Nothing in it requires a 24-hour 0800. The RGC is an Anatel rule with its own 24-hour requirement for telecom providers, described below; this guide has not verified whether Anatel will update the RGC because of the new decree.

Point	Decree 11,034 (until November 1, 2026)	Decree 13,135 (from November 2, 2026)
Phone	Mandatory (art. 4, § 2)	Mandatory (art. 14)
Live-agent phone service	At least 8 hours a day (art. 5, I)	8 a.m. to 8 p.m., every day (art. 14, § 1)
24-hour access	At least one of the integrated channels (art. 4)	At least one electronic channel (art. 12)
Maximum wait time	Set by the regulators (art. 5, III)	20 seconds for 80% of the demands and 300 seconds for the rest, unless the regulator has a rule (art. 15, I and § 1)
Free access	Art. 3	Art. 8
Call recordings	At least 90 days (art. 12, § 3)	At least 90 days (art. 19)
Interaction record	At least two years (art. 12, § 5)	History for at least six months after resolution (art. 18, § 4, II)
Who it covers	Services regulated by the federal Executive Branch	The same criteria and, also, financial services under Central Bank regulation or supervision (art. 1, sole paragraph)

Source: the texts of [Decree 11,034/2022](#) and [Decree 13,135/2026](#) on the Planalto website. The rest of this guide describes Decree 11,034, which applies until November 1, 2026, and Anatel's RGC. Analysis of how the new decree changes your obligation is for your legal team.

What does Decree 11,034 require of phone support, until November 1, 2026?

Until November 1, 2026, Decree 11,034 requires access to the SAC to be free (art. 3), available 24 hours a day, 7 days a week, through at least one channel (art. 4), and phone support to exist (art. 4, § 2), with at least 8 hours a day of live-agent support and complaint and cancellation options in the first menu (art. 5). From November 2, the rules of Decree 13,135, in the previous section, apply.

Article (Decree 11,034)	What it says	What it means for the 0800 and the IVR
Art. 3	Access to the SAC is free and handling requests creates no cost for the consumer	The support number cannot charge the caller
Art. 4, main text and § 1	Uninterrupted access, 24 hours a day, 7 days a week, through at least one of the integrated channels	The 24-hour channel can be another one; the decree does not require it to be the phone
Art. 4, § 2	Access through phone support is mandatory, under the terms of art. 5	The phone line has to exist and work during the hours in art. 5
Art. 4, § 4	Initial access to an agent cannot depend on the consumer providing data first	The IVR cannot demand a CPF (taxpayer ID) or contract number before letting the caller speak to someone
Art. 5, I	Support hours of no less than 8 hours a day, with live-agent support	An agent available at least 8 hours a day
Art. 5, II	The first menu must include the complaint and cancellation options	The IVR's greeting and first menu have to offer those options

Article (Decree 11,034)	What it says	What it means for the 0800 and the IVR
Art. 5, III	Maximum wait time to reach an agent and to be transferred, set by the regulators	The number comes from the sector's rule, not from the decree

Source for every row: the [text of Decree 11,034/2022 on the Planalto website](#).

Free access (art. 3)

Art. 3 of Decree 11,034 says that “access to the SAC shall be free and handling requests shall create no cost for the consumer”. In Brazilian numbering, the code that meets this by design is 0800 (toll-free): under [art. 18 of Anatel Resolution 749/2022](#), on 0800 the subscribing organization takes responsibility for payment, “making it a call at no cost to the originating user”. Under the same resolution, the 0300 and 0303 codes have their cost shared between the caller and the subscriber. The guide [Why won't an 0800 connect from a mobile phone?](#) covers the origin restrictions that tend to show up under this billing model. Art. 8 of Decree 13,135 also sets free access on all channels.

24 hours through at least one channel (art. 4)

The main text of art. 4 of Decree 11,034 requires access “uninterruptedly, twenty-four hours a day, seven days a week”, and § 1 says this will be guaranteed through “at least one of the integrated support channels”. Read literally, the 24-hour channel can be an app or a website, while the phone line, mandatory under § 2, follows the hours in art. 5. § 3 allows the SAC to be interrupted when the service is not available for use or purchase, subject to the sector's rule. §§ 5 and 6 prohibit advertising during the wait without prior consent and allow informational messages. Under Decree 13,135, the 24-hour requirement moves to an electronic channel (art. 12), distinct from the phone call (art. 3, paras. 1 and 2).

The minimum for phone support (art. 5)

Art. 5 of Decree 11,034 lists what regulators must observe on the phone line: at least 8 hours a day with live-agent support (item I), complaint and cancellation in the first menu (item II) and a maximum wait time (item III). The sole paragraph lets regulators require more than 8 hours of live-agent support. Art. 15, § 5, gives the National Consumer Secretariat (Secretaria Nacional do Consumidor) the same option for providers whose SAC performs poorly.

What does Decree 11,034 not set?

Decree 11,034 does not set the maximum wait time, the resolution deadline per sector or the deadline for technically processing a cancellation. On all of these, it leaves the decision to the regulator. Nor does it set live-agent hours above 8: it only allows the regulator to require them.

The points where the decree delegates to the regulator:

- **Maximum wait time** (art. 5, III): to reach an agent directly and to be transferred to the right department. The decree requires a maximum to exist, but does not say what it is.
- **More hours of live-agent support** (art. 5, sole paragraph).
- **Deadline for resolving requests** (art. 13, § 4). The decree sets only the response deadline: seven calendar days from registration (art. 13, main text).
- **Deadline for technically processing a cancellation** (art. 14, sole paragraph). Under art. 14, cancellation takes effect immediately, unless that processing is needed.

- **The provider's size** (art. 1, sole paragraph).

In practice, anyone looking for “the maximum SAC wait time” needs their own sector's rule. For telecom, Anatel's RGC leaves that time and the definition of “first menu” to the Implementation Group (art. 18, § 3, of Resolution 765/2023). From November 2, 2026, Decree 13,135 sets the time (up to twenty seconds for at least 80% of the demands and up to three hundred seconds for the rest, art. 15, I), but it does not apply where the regulator has a specific rule (art. 15, § 1).

What records does the SAC decree already require you to keep?

The SAC decree requires recordings of calls to the SAC to be kept for at least 90 days and a ticket number for every request. Under Decree 11,034, the record of each interaction is kept for at least two years after resolution (art. 12); under Decree 13,135, the history is kept for at least six months after resolution (art. 18, § 4, II). None of these records shows whether the 0800 was up when nobody could get through.

Art. 12 of Decree 11,034 guarantees the consumer can follow up on a request “by means of a numeric record or another kind of electronic procedure”. § 3 makes “keeping the recording of the call made to the SAC” mandatory for at least ninety days, and § 4 lets the consumer request access to it within that period. § 5 requires the interaction record to be kept for at least two years. Under Decree 13,135, the consumer follows up by ticket number (art. 18), and call recordings and message content are kept for at least ninety days (art. 19).

All of these records come from a call that got through. When the 0800 won't connect, there is no call, no recording and no ticket number. A gap in the records over a period could mean no demand or no channel, and the SAC records alone cannot tell one from the other.

What is Anatel's RGC?

The RGC is the Regulamento Geral de Direitos do Consumidor de Serviços de Telecomunicações (General Regulation on Consumer Rights for Telecommunications Services), approved by Anatel Resolution No. 765/2023. It covers fixed telephony, mobile telephony, fixed broadband and pay TV, and has been in force since September 1, 2025. It is the telecom sector's rule that applies the SAC decree.

The link to the decree is in the text itself: art. 3, V, of the Annex defines Phone Support “under the terms set out in this Regulation and in Decree No. 11,034 of April 5, 2022”. That reference is in the published text of the RGC; this guide has not verified whether Anatel will update it after 11,034 is revoked. Art. 1 of the Annex covers the STFC (fixed-line telephone service), the SMP (mobile service), the SCM (multimedia communication service, which includes broadband) and pay-TV services; art. 2 excludes IoT-only connections.

When it took effect is noted on [Anatel's Resolution 765 page](#): under Ruling (Acórdão) No. 228 of August 28, 2024, most of the regulation took effect on September 1, 2025, and the previous RGC (Resolution 632/2014) was revoked on the same date.

Art. 15: 24-hour phone support

Art. 15 says Phone Support must: “I - be free; II - take calls originating from the STFC and the SMP; III - be available uninterruptedly, 24 (twenty-four) hours a day, 7 (seven) days a week; and, IV - have a live agent available 24 (twenty-four) hours for urgent requests and, for all others, from 6 (six) to 22 (twenty-two) hours, uninterruptedly, 7 (seven) days a week.”

Here the difference from the decree is direct: for telecom providers, the **phone line** has to be available 24 hours a day, not just some channel. Item II matters too: the number has to take calls from landlines (STFC) and mobile phones (SMP). § 1 lets Anatel change the live-agent hours, and § 2 leaves the definition of “urgent requests” to the Operational Manual.

Art. 18: the first menu

Art. 18 says the first menu of Phone Support must offer, at a minimum, the options of “I - terminating the contract in an automated way, without an agent’s involvement; II - registering a complaint; and, III - live-agent support”. § 1 requires the option to speak to an agent in every submenu, and § 3 leaves the maximum wait time and the definition of first menu to the Implementation Group.

Art. 91: small providers

Art. 91 covers Small Providers (Prestadoras de Pequeno Porte): their phone support must operate “for at least 8 (eight) hours, uninterruptedly, on business days”. Art. 90 lists the articles that apply to them, and from art. 15 it includes items I and II (free calls and calls from landlines and mobile phones), but not items III and IV (24 hours and the live-agent hours).

Point	Decree 11,034/2022 (until November 1, 2026)	RGC, Resolution 765/2023
Who it covers	Providers of services regulated by the federal Executive Branch, as each sector’s regulator determines	Fixed and mobile telephony, broadband and pay-TV providers
In force since	October 3, 2022	September 1, 2025
Free of charge	Art. 3	Art. 15, I
Available 24 hours	Access to the SAC, through at least one channel (art. 4)	Phone support itself (art. 15, III)
Live-agent support	At least 8 hours a day (art. 5, I)	24 hours for urgent requests, 6:00 a.m. to 10:00 p.m. for the rest (art. 15, IV)
First menu	Complaint and cancellation (art. 5, II)	Automated cancellation, complaint and live agent (art. 18)
Maximum wait time	Set by the regulators (art. 5, III)	Defined by the Implementation Group (art. 18, § 3)
Small providers	The regulator takes size into account (art. 1)	8 hours on business days (art. 91)

From November 2, 2026, the decree column becomes Decree 13,135, in the table in the section on the new decree.

How do I prove the 0800 was up?

To prove the 0800 was up, you need a record made at the time, from outside the company, showing for each time slot whether the call connected and whether the IVR answered with the right message. Neither the decree nor the RGC says how to prove availability; scheduled test calls are one way to produce that record.

The company’s own systems tell the story from the inside. A PBX that is powered on and a trunk that is registered show the equipment was working, not that a customer outside could get through. When the problem is on the carrier’s route or in a swapped IVR, the dashboard stays green.

What a test call records

Monitor, 2N5's managed service, places real test calls to each support number: 0800, 0300, 4004 and landline numbers. Each monitored number is a **voice circuit**, with its own schedule and the people who should know when it fails. The call goes out from the 2N5 platform, listens to what answers and hangs up. For each call it records:

- the time and the number called;
- whether the call connected or was rejected, with the SIP code the route returned, such as a [SIP 503](#), and the Q.850 cause when the network provides it;
- whether what answered was what was expected: during setup, 2N5 records the [IVR's](#) reference greeting, and each call compares what it heard with that recording. A carrier announcement saying the number does not exist does not pass;
- whether audio got through, and the voice quality measured on that audio, MOS-LQO.

SIP codes follow [RFC 3261](#), and Q.850 causes arrive in the Reason header defined by [RFC 6432](#). This is the kind of data the carrier can find in its own records.

A note on personal data: the audio of a test call is kept only when what answers is the recording or the IVR on file, encrypted and for the period 2N5 sets. A person's voice is never kept; if the number could be answered by someone, that person first hears that it is an automated test.

Why the interval between calls matters

A record only shows what it measured. If the test runs once an hour, a 40-minute outage can slip between two good calls, and the period comes out as available. At Monitor's default pace, each number gets a call every 15 minutes. When a call fails, it is repeated three times, every 2 minutes; four failures in a row open an incident, and two good calls in a row confirm recovery.

The margin math: an outage that starts right after a good call is seen at most 15 minutes later, on the next call. With retries every 2 minutes, the fourth failure in a row comes 6 minutes after the first ($2 + 2 + 2$). An isolated blip, which fails once and recovers, does not open an incident and does not enter the calculation.

Fixed times, nights and weekends

The obligation in art. 15, III, of the RGC is 24 hours a day, 7 days a week. Each monitored number has its own schedule, at intervals or at fixed times, including nights and weekends. Only a schedule that covers the full obligation hours lets you say "tested throughout the period".

The report that documents the year

Every month, Monitor emails a service report to the people the company chose, in Portuguese or English. It gives each circuit's availability for the month and for the **last 12 months**, which is the figure that closes the year. The 12 monthly reports, each with its 12-month figure, cover the whole year of obligation.

Report section	What it documents
Voice availability	The month and the last 12 months, against the target
Incidents	Each incident, with its severity and duration
Downtime	The total for the period, by circuit
Average response	The time until someone took ownership of each incident

Report section	What it documents
The period, day by day	Which days had a confirmed failure
By circuit and by carrier	The availability of each number and each carrier

Two things keep the figure defensible. First: a period in which 2N5 could not test because of a problem on 2N5's side stays out of the calculation and is shown separately; the monitoring's lack of visibility never turns into circuit unavailability. Second: scheduled maintenance has to be recorded beforehand, in a maintenance window with a start and an end. The guide [Telephony SLA: yearly and monthly calculation](#) shows the math, and the [SLA calculator](#) does it for any target.

An example with demo data: 99.89% availability for the month and 99.91% over the last 12 months, a 99.90% target met, 4 incidents and 3 h 55 min of downtime in the month.

What does a test call not prove?

A test call proves that the number connected from the origin it calls from and that the IVR answered with the expected message. It does not prove the wait time to reach an agent, the live-agent hours, the content of each menu option or that the number connected from every carrier.

Requirement	Article	Does the test call show it?
Phone channel available and answering	Decree 11,034 art. 4, § 2; Decree 13,135 art. 14; RGC art. 15, III	Yes, for each tested time slot, from 2N5's origin
The right IVR answering	No specific article; it is what the caller hears	Yes, the greeting is compared with the reference recording
Free for the caller	Decree 11,034 art. 3; Decree 13,135 art. 8; RGC art. 15, I	No; it depends on the code contracted, such as 0800
Calls from landlines and mobile phones	RGC art. 15, II	Only from the origin the test calls from
Live-agent support during the required hours	Decree 11,034 art. 5, I; Decree 13,135 art. 14, § 1; RGC art. 15, IV	No; the call hears the greeting and hangs up
Maximum wait time	Decree 11,034 art. 5, III; Decree 13,135 art. 15, I; RGC art. 18, § 3	No
First-menu options	Decree 11,034 art. 5, II; Decree 13,135 art. 16, § 2; RGC art. 18	Only what is in the compared greeting
Recordings and ticket numbers	Decree 11,034 art. 12; Decree 13,135 arts. 18 and 19	No; these are the SAC's own records

Each origin follows its own route to the 0800, and one route can fail while the others work. 2N5's test does not replace a test from each mobile carrier when the suspected cause is an origin restriction.

How do I build the SAC availability record?

The SAC availability record is built from four pieces: the list of numbers consumers can dial, what should answer each one, a test schedule that covers all the obligation hours and a report kept month by month. With that, the company can answer with data when someone claims the phone line was not answering.

1. **List every published number.** Art. 7 of Decree 11,034 (and art. 9 of Decree 13,135) requires the SAC access channels to appear in the provider's documents and channels; every published number, 0800, 0300 or 4004, is a number someone will dial.

2. **Define what should answer each number.** The reference recording tells “the right IVR answered” apart from “the carrier’s announcement answered”.
3. **Cover all the obligation hours.** For a telecom provider, 24 hours a day, 7 days a week (art. 15, III, of the RGC). For other sectors, the hours the regulator’s rule sets, with live-agent phone service from 8 a.m. to 8 p.m. every day from November 2, 2026 (art. 14, § 1, of Decree 13,135).
4. **Record maintenance beforehand.** A maintenance window with a start and an end, scheduled before it begins. Whatever runs past the window is an outage.
5. **Keep the monthly reports.** The 12-month figure in each report is what closes the year.
6. **Take the evidence to the carrier.** The time, the number called, the SIP code and the Q.850 cause of each test call locate the failure on the route. The [0800 and IVR monitoring](#) page shows how Monitor handles each case, and [SLA reports](#) shows the report.

What this guide is not

This guide is not legal advice. It summarizes the text of Decrees 11,034/2022 and 13,135/2026 and Anatel Resolution 765/2023, in the versions published on the official websites and checked on October 14 and 15, 2026, when Decree 13,135/2026 had been published but was not yet in force (it applies from November 2, 2026), and shows one technical way to document the availability of the phone channel.

Three questions are for your legal team: whether your company’s service is subject to the decree; what your regulator’s rule says about phone support, wait times and live-agent support; and what weight a record of test calls carries in an administrative or court proceeding. The text that governs is the one published in the Diário Oficial da União.

Frequently asked questions

Does the SAC decree require the 0800 to work 24 hours a day?

No. Under Decree 11,034/2022 (until November 1, 2026), SAC access works 24 hours a day through at least one integrated channel, not necessarily the phone, and live-agent phone support runs at least 8 hours a day (arts. 4 and 5). Under Decree 13,135/2026 (from November 2), the 24-hour channel is electronic (art. 12), and the phone has live agents from 8 a.m. to 8 p.m. (art. 14). For telecom, the RGC requires the phone line 24 hours.

What is the maximum SAC wait time under Decrees 11,034 and 13,135?

Decree 11,034 sets no number: art. 5, III, tells the regulators to set the maximum wait time. From November 2, 2026, Decree 13,135 sets up to twenty seconds for at least 80% of the demands served and up to three hundred seconds for the rest (art. 15, I), unless the regulator has a specific rule (art. 15, § 1). For telecom, the RGC leaves that time to the Implementation Group (art. 18, § 3).

What changed with the RGC in Resolution 765/2023?

Since September 1, 2025, phone support from telephone, internet and pay-TV providers must be free, take calls from landlines and mobile phones and be available 24 hours a day, with a live agent 24 hours for urgent requests and from 6:00 a.m. to 10:00 p.m. for everything else (art. 15). The first menu must offer automated cancellation, complaints and a live agent (art. 18).

How do I prove the SAC was available overnight?

With records made at the time, from outside the company. A test call to the 0800 at short intervals, overnight included, leaves for each time slot the call's result, the SIP code returned and whether the IVR answered with the expected message. Added up, those records give the period's availability.

Does this guide replace a legal review?

No. It summarizes the official text of the decrees and the RGC and shows one way to document the availability of the phone channel. Whether your company is subject to each rule, and how your regulator's rule applies the decree to your sector, are questions for your legal team.

Has Decree 11,034 been revoked?

It was revoked by Decree 13,135, of October 9, 2026 (art. 24), published in the Diário Oficial on October 13 and in force twenty days later (art. 25), that is, from November 2, 2026. Decree 11,034 applies until November 1, 2026; from November 2, Decree 13,135 applies, which keeps phone support mandatory and live agents from 8 a.m. to 8 p.m. (art. 14).

Sources

1. Decree No. 11,034/2022 (Decreto nº 11.034/2022): guidelines and rules for the SAC (arts. 1 to 18)
2. Decree No. 13,135/2026: the new SAC decree, which revokes 11,034 (arts. 1, 8, 12, 14 to 16, 18, 19, 24 and 25)
3. Anatel, Resolution No. 765/2023 (RGC): phone support, arts. 15, 18 and 91, and notes on when it took effect
4. Anatel, Resolution No. 749/2022 (Regulamento de Numeração, the numbering regulation), art. 18: the 0800 code is free for the caller
5. RFC 3261: SIP and the classes of response codes
6. RFC 6432: Q.850 causes in the Reason header of SIP responses

GUIDE

Silent calls and one-way audio in VoIP: causes and how to fix them

The call connects, the timer runs and one side hears nothing. Almost always the SIP signaling worked and the RTP audio got lost along the way, in a NAT, a firewall or a wrong address in the SDP. This guide covers each cause, where to check in Asterisk and in the firewall, and how to prove that audio is flowing in both directions again.

By [2N5 team](#) Published October 7, 2026 Updated October 20, 2026

Summary

- Silent calls (no audio in either direction) and one-way audio are RTP failures, not SIP failures: the call connects and the audio does not arrive.
- The most frequent cause is a private IP address in the SDP, announced by a PBX behind NAT with no external address configured.
- The SIP ALG in routers and firewalls rewrites the SDP and tends to make things worse; the usual advice is to turn it off and handle NAT on the PBX.
- The RTP UDP port range must be open and forwarded to the PBX; in Asterisk it lives in `rtp.conf`, 10000 to 20000 in the stock sample.
- A call that drops at around 32 seconds points to a lost ACK: the 200 OK is retransmitted for $64 \times T1$, with $T1$ at 500 ms, and then the session is torn down.
- A test against an echo or tone service proves both directions; the sent and received RTP counters show which side the audio stopped on.

What is a silent call and what is one-way audio?

A silent call is a VoIP call that connects with no audio in either direction: nobody hears anybody. One-way audio is a call in which one side hears and the other doesn't. In both cases the [SIP](#) signaling worked; what failed was the audio path, the [RTP](#) stream.

The industry terms are *no-way audio* and *one-way audio*. One-way audio means one direction found its path and the other didn't, which almost always points to an address, NAT or firewall issue on one specific side. A silent call usually points to something that affects the whole stream: the codec, encryption, a port range closed in both directions or media marked as inactive.

On an 0800 (toll-free) number or an IVR, a call that connects and stays silent is worse than a rejection: it shows up as answered and no completion alarm fires.

Why does the call connect but the audio doesn't get through?

The call connects and the audio doesn't get through because a VoIP call uses two separate paths. SIP (RFC 3261) negotiates the call; the audio travels in RTP (RFC 3550), on UDP ports chosen on the fly. The 200 OK proves SIP got through, but it doesn't prove RTP can reach the announced address and port.

The audio address travels inside the SIP message itself, in the SDP body. Under the offer/answer model in RFC 3264, the caller offers an address, a port and a list of codecs; the answering side responds with its own. Each side sends audio to the address the other announced. An example of an offer from a PBX

behind NAT:

```
v=0
o=- 4021 4021 IN IP4 192.168.1.20
s=-
c=IN IP4 192.168.1.20
t=0 0
m=audio 14562 RTP/AVP 8 0
a=rtpmap:8 PCMA/8000
a=rtpmap:0 PCMU/8000
a=sendrecv
```

The `c=` line says where to send the audio, and the `m=` line says which port and which codecs. In this example, the PBX announced 192.168.1.20, a private address. The carrier, on the outside, sends the audio to that address, which doesn't exist on the internet, and the audio disappears. The PBX, in turn, sends its audio to the public address the carrier announced, and that direction works. Result: the carrier hears the extension, and the extension doesn't hear the carrier.

The `a=sendrecv` line matters too. The `sendonly`, `recvonly` and `inactive` attributes deliberately limit the direction of the audio, for example on a call on hold. An `inactive` left over from a hold or a transfer that didn't complete cleanly can leave the call silent with the entire network working.

What are the most common causes of silent calls and one-way audio?

The most common causes of silent calls and one-way audio in VoIP are eight: unhandled NAT, a SIP ALG rewriting the SDP, a blocked RTP range, the wrong external IP in the SDP, no symmetric RTP, an incompatible codec, SRTP on one side only and mishandled early media. The table summarizes the symptom of each.

Cause	Typical symptom	Where to check
Unhandled NAT	One-way audio, almost always outside-to-inside	The address on the <code>c=</code> line of the SDP the PBX sends
SIP ALG on the router or firewall	Symptoms that change from call to call, registrations that drop, calls that drop at around 32 s	The SIP or VoIP options on the router or firewall
RTP range blocked or not forwarded	Silent or one-way calls, only on calls through the trunk	Inbound rules and UDP port forwarding
Wrong external IP in the SDP	Inbound audio disappears, sometimes only after an IP or link change	<code>external_media_address</code> in PJSIP, <code>externaddr</code> in chan_sip
No symmetric RTP	One-way audio with remote phones or with the far end behind NAT	<code>rtp_symmetric</code> and <code>force_rport</code> on the endpoint
Incompatible codec	Rejection with SIP 488, or a call that connects with silence or noise	The <code>m=</code> and <code>a=rtpmap</code> lines of the offer and the answer
SRTP on one side only	Silent in both directions, or rejection with SIP 488	The <code>RTP/SAVP</code> profile in the SDP and media encryption on the endpoint
Early media	Silence before the answer: no ringback and no carrier announcement	183 Session Progress with SDP

Unhandled NAT

NAT swaps the PBX's private address for the link's public address in the IP headers, but it doesn't touch the content of the SIP message. The SDP still says 192.168.1.20, and the far end believes it. RFC 4787 describes how NATs handle UDP, with different mapping behaviors depending on the destination, and RFC 6314 collects the NAT traversal practices for SIP: announce the public address, reply to the port the message came from, keep the mapping alive with periodic messages and use symmetric RTP. Phones and adapters behind NAT can also discover their own public address with STUN (RFC 8489).

A SIP ALG rewriting the SDP

The [SIP ALG](#) is a feature of many routers and firewalls that reads SIP messages, swaps private addresses for public ones inside SIP and the SDP, and opens RTP ports on demand. When the PBX already announces the external address, the two corrections stack; when the ALG rewrites only part of the message, for example the Contact header but not the `c=` line, the result is an inconsistent SDP. The ALG also can't read SIP over TLS, because the message is encrypted. That is why the usual advice is to turn it off and handle NAT in one place only.

RTP port range blocked or not forwarded

Each RTP stream uses one UDP port, and under RFC 3550 the matching RTCP uses the next port up. The PBX picks ports from a configured range, and the firewall has to let that entire range through to it. Partial forwarding produces an intermittent fault: calls that land on a port inside the forwarded part work, and the rest are silent.

Wrong external IP in the SDP

Configuring the external address fixes NAT until the address changes. Links with a dynamic IP, a backup link that takes over when the main one fails, or an outbound route over a different link from the one announced all make the PBX announce an address that isn't the one traffic actually leaves from. The same goes for the list of local networks: Asterisk only swaps the private address for the public one for destinations outside the networks listed as local. An internal network missing from the list, such as a branch office VPN, gets the public IP and ends up with one-way audio.

Symmetric RTP

Symmetric RTP means sending and receiving audio on the same port and, on the other side, sending audio back to the address and port it actually arrived from, ignoring what the SDP said. It is the most robust way to traverse a NAT you don't know, and it is one of the practices in RFC 6314. When the far end, a remote phone or a carrier behind NAT, announces an address that doesn't work, only symmetric RTP on your side gets the audio to it.

Incompatible codec

If the offer and the answer share no [voice codec](#), the call is usually rejected with [SIP 488 Not Acceptable Here](#). The harder case is the call that connects: one side accepts a codec it can't handle, a codec transcoding fails on the PBX, or the codec list changes in a re-sent SDP. The call shows as answered and stays silent or noisy.

SRTP on one side only

With SRTP, the audio is encrypted, and the SDP announces the `RTP/SAVP` profile instead of `RTP/AVP`. If one side requires encryption and the other sends plain audio, or if the keys weren't agreed, the packets arrive and are discarded. The call is silent in both directions with normal packet counters. Under RFC 3264, a rejected media stream comes back in the answer with port 0, and that detail in the SDP is usually the clue.

Early media

Early media is audio that arrives before the answer: ringback generated by the network or the carrier's announcement saying the number doesn't exist. It comes with a `SIP 183 Session Progress` carrying SDP, and RFC 3960 describes how that early audio is handled. If the PBX, the firewall or the phone doesn't let that audio through, the caller hears silence until the answer, and a call that fails with a recorded announcement seems silent. An SDP in the 183 that differs from the SDP in the 200 OK can also leave the audio pointed at the wrong place after the answer.

How does the symptom point to the cause?

The symptom of a silent call or one-way audio points to the cause once you look at which direction the audio fails in, at what point in the call and on which calls. A fault on every inbound call is different from one that shows up on one call in three.

Symptom	Likely cause	First check
The extension can't hear the outside party; the outside party hears the extension	Private IP in the SDP, SIP ALG, inbound RTP range closed	The SDP the PBX sends and the port forwarding
The outside party can't hear the extension	Restricted outbound UDP, or the far end's SDP with an unreachable address	Outbound firewall rules and the SDP received
Nobody hears anybody	Codec, SRTP, media <code>inactive</code> or rejected with port 0	Offer and answer side by side
Audio stops after a hold or transfer	A new SDP in a re-INVITE, direct media between the endpoints	<code>direct_media</code> and the re-INVITE's SDP
Audio stops after a few minutes of silence	Short UDP session timeout on the firewall	UDP session timeouts and keepalive messages
The call drops at around 32 seconds	ACK lost because of a Contact with a private address	The Contact header of the 200 OK
Fails on some calls and not others	RTP range half forwarded, two outbound links	The RTP port of the calls that failed
Silence only before the answer	Early media blocked	The 183 with SDP and what the PBX does with it

Why does the call drop after 32 seconds?

Because the ACK never arrived. Under RFC 3261, the side that answers a call retransmits the 200 OK until it receives the ACK. If $64 \times T1$ passes without an ACK, the session must be torn down. With the default $T1$ of 500 ms, that is $64 \times 0.5 = 32$ seconds. The ACK is sent to the address in the 200 OK's Contact header; if the PBX behind NAT put its private address there, the ACK gets lost, the PBX keeps retransmitting and drops the call when the timer runs out. It isn't an audio problem, but it shows up alongside one, because it comes from the same mishandled NAT.

How do I fix silent calls on Asterisk, Issabel and FreePBX?

On Asterisk, Issabel and FreePBX, fixing silent calls comes down to four adjustments: set the external address and the local networks on the transport, turn off direct media and turn on symmetric RTP on the trunk endpoints, align the rtp.conf RTP range with the firewall, and check the codecs and media encryption on both sides.

PJSIP: the transport

In PJSIP, the external address and the local networks live on the transport. Using documentation addresses, with 203.0.113.10 as the public IP and two internal networks:

```
[transport-udp-nat]
type=transport
protocol=udp
bind=0.0.0.0:5060
external_media_address=203.0.113.10
external_signaling_address=203.0.113.10
local_net=192.168.1.0/24
local_net=10.20.0.0/16
```

`external_media_address` is the address that goes on the SDP's `c=` line;

`external_signaling_address` is the one that goes in the SIP headers, such as Contact and Via.

`local_net` lists the networks for which Asterisk keeps the private address. List every internal network, including branch office VPNs.

PJSIP: the trunk endpoint

On the endpoint for the trunk and for remote phones:

```
[carrier-trunk]
type=endpoint
transport=transport-udp-nat
direct_media=no
rtp_symmetric=yes
force_rport=yes
rewrite_contact=yes
disallow=all
allow=alaw,ulaw
```

With `direct_media=yes`, Asterisk tries to take itself out of the audio path after the answer, with a new INVITE so the two endpoints exchange RTP directly. Across NAT, this tends to fail and the audio vanishes right after the answer or a transfer; `direct_media=no` keeps the audio flowing through Asterisk. `rtp_symmetric=yes` turns on symmetric RTP, `force_rport=yes` replies to the port the message came from, and `rewrite_contact=yes` stores the address the registration came from, not the one the device claimed. The `disallow` and `allow` lines set the accepted codecs, in order of preference; they have to overlap with what the carrier accepts. The command `pjsip show endpoint carrier-trunk` shows the values in effect.

If the carrier requires SRTP, the endpoint's `media_encryption` option has to match it. If it doesn't use SRTP, the option has to be off.

chan_sip, on older installations

In `chan_sip`, which recent Asterisk versions no longer ship, the equivalents live in `sip.conf`: `externaddr` (formerly `externip`) for the public address, `localnet` for the local networks, `nat=force_rport,comedia` for symmetric RTP and `directmedia=no` to keep the audio on Asterisk.

rtp.conf: the port range

Asterisk's RTP port range lives in `rtp.conf`:

```
[general]
rtpstart=10000
rtpend=20000
```

These are the values in Asterisk's sample file and FreePBX's default. The firewall's range has to be exactly the same, or wider. If you need to narrow the range, do the math: each call leg that passes through Asterisk uses a pair of ports (RTP and RTCP), and a call between a phone and the trunk has two legs, so four ports. A range of 10,000 ports holds, roughly, $10,000 \div 4 = 2,500$ concurrent calls.

In the same file, the `strict RTP` option, on by default in current versions, makes Asterisk accept RTP only from the source it learned at the start of the call. A carrier that changes the audio source mid-call, for example after a recorded announcement, can have its audio discarded. Turning it off is a diagnostic test, not the permanent fix.

FreePBX and Issabel

In FreePBX, the equivalent fields are under **Settings** → **Asterisk SIP Settings**, on the **General SIP Settings** tab: **External Address**, **Local Networks** and **RTP Port Ranges**. Prefer the interface over the files: FreePBX rewrites the files it generates when it applies the configuration, and manual adjustments go in the files ending in `_custom.conf`. Issabel ships Asterisk with an interface derived from FreePBX, and the same three items, external address, local networks and RTP range, appear in its SIP settings. On older Issabel systems running `chan_sip`, the previous section applies.

For a first internal test, FreePBX has an echo test on code `*43`, and on plain Asterisk the dialplan's `Echo()` application does the same. That test only proves the leg between the phone and the PBX; the trunk, the NAT and the firewall are left out.

How do I set up the firewall and router for VoIP audio?

To get VoIP audio through the firewall and the router, three adjustments solve most cases: turn off the SIP ALG when the PBX already handles NAT, forward to the PBX the same UDP RTP range it uses, and give UDP sessions a lifetime longer than the interval between messages.

Turn off the SIP ALG

The name and location vary from vendor to vendor, so always check the documentation for the installed version:

- **FortiGate:** SIP handling can come from a `SIP session helper` or from the SIP ALG applied by the VoIP profile on the policies. Check both before concluding that SIP passes through unaltered.

- **MikroTik:** the *helper* is under **IP** → **Firewall** → **Service Ports**, in the **sip** entry. Disable that entry.
- **pfSense:** the point most often raised for VoIP is outbound NAT changing the source port of connections. An outbound NAT rule with a static port for the PBX's address prevents the change.
- **Broadband routers:** many offer the option as "SIP ALG" or "SIP Passthrough" in the advanced NAT settings.

After turning it off, clear the active sessions so the old entries leave the connection table.

Open and forward the RTP range

Forward the entire UDP range, 10000 to 20000 in the example, to the PBX's internal address, and allow that same range in the inbound rules. For security, restrict the source to the carrier's addresses whenever it tells you what they are, and do the same for the SIP port.

Adjust the UDP session timeouts

The firewall forgets a UDP mapping that goes without traffic. RFC 4787 requires that timer not to expire in less than 2 minutes and recommends a default of 5 minutes or more. Equipment with shorter timeouts drops the SIP registration's mapping between two renewals, and the trunk starts receiving calls that never reach the PBX. Two ways out: raise the UDP session timeout on the firewall for the voice ports, or have the PBX send messages more often, with a shorter registration renewal and the AOR's `qualify_frequency` in PJSIP, which sends a periodic OPTIONS.

What if there's an SBC?

An **SBC** at the edge takes over NAT and anchors the audio between the PBX and the carrier; the PBX talks only to the SBC, on a known network. The thing to watch is not handling NAT twice: an SBC, an ALG and a PBX all correcting the same address produce an inconsistent SDP.

How do I know whether the 0800 or the SIP trunk has silent calls?

The way to know whether an 0800, an IVR or a SIP trunk has silent calls is to call from outside and check the audio in both directions, on a real call, repeatedly. A PBX dashboard can't see the problem: the trunk shows as registered, the call is answered, and only the person on the line notices the silence.

The test that proves both directions is a call to an echo or tone service. Whatever the call sends has to come back: if the echo returns the audio, it got there and came back. An internal call, an echo test on the PBX itself or the ringback before the answer don't prove this, because they leave out the trunk, the NAT or the audio path itself after the answer. Because of intermittent faults, such as a half-forwarded RTP range, a single test isn't enough either: the fault may show up on only some of the calls.

That is what Monitor does for each **voice circuit**, meaning each monitored SIP trunk, 0800 or support number. 2N5 places real test calls from its own platform, at the agreed pace, with nothing installed on the PBX and no integration with it. When the number is answered by a test tone or an echo service, the call passes when it is answered and the tones it sends come back. Audio in only one direction, or none at all, counts as a failure even when the call connects. See what each call checks in [SIP trunk monitoring](#).

A failure is retried before any alert. At the default pace, each circuit gets a call every 15 minutes; after a failure, the call is repeated three times, every 2 minutes, four failures in a row open an incident and two good calls in a row confirm recovery. The incident alerts the escalation ladder by email, until someone

takes ownership; see [alerts and escalation](#). If the problem is on 2N5's side, the test is marked monitoring blind and does not count against the circuit.

How do the RTP counters show which side the audio stopped on?

The sent and received RTP counters show which side the audio stopped on because each direction has its own stream. On a normal call, both counters climb together. When one of them sits at zero while the other climbs, one-way audio is proven, and the stalled direction tells you which leg to investigate.

The reference math is simple. With audio packetized every 20 ms, each direction sends $1,000 \div 20 = 50$ packets per second. A 60-second call has about 3,000 packets in each direction. A received counter at zero next to 3,000 sent leaves no doubt.

In Asterisk, three commands give you that view:

- `pjsip show channelstats` shows, for each active channel, packets received and transmitted, loss, jitter and round-trip time;
- `rtp set debug on` shows each packet with its source and destination address, and immediately reveals audio being sent to a private address; `rtp set debug ip` filters by address;
- on `chan_sip`, `sip show channelstats` does the job of the first one.

Outside Asterisk, a packet capture on the PBX or the firewall, opened in Wireshark under **Telephony** → **RTP** → **RTP Streams**, lists each stream with source, destination and packet count. Capturing on both sides of the firewall shows whether the packet stops at the edge. RTCP helps too: under RFC 3550, the sender report carries the count of packets and bytes sent, and the receiver report carries the fraction lost, the cumulative number lost and the jitter. An endpoint that never receives RTCP from the other has the same problem on the return path.

Sent by the PBX	Received by the PBX	Reading
Normal	Zero	Audio from outside isn't arriving: private IP in the SDP, SIP ALG or inbound RTP range closed
Zero	Normal	The PBX isn't sending: <code>recvonly</code> or <code>inactive</code> media, codec transcoding failure, direct media
Normal	Normal, but still silent	Audio arrives and isn't used: SRTP without a key, wrong codec, the leg between PBX and phone
Zero	Zero	No stream at all: media rejected with port 0, <code>inactive</code> , channel not bridged
Climb, then stop	Climb, then stop	A new SDP in a re-INVITE, UDP session timeout expired, source change with <code>strict RTP</code>

What should I take to the carrier?

Take the carrier the time of each failed call, the number called, the Call-ID, the SIP capture with the SDP of the offer and the answer, and the RTP counters for both directions. If your PBX sends audio to the announced address and nothing comes back, the problem is on the carrier's side or between the two of you.

A silent call has no error code: the call completed with a 200 OK. In Monitor, each incident carries the circuit, the carrier, the first symptom, the SIP code and Q.850 cause of each call, and the last healthy call, in a block ready to paste into the ticket. The [SIP code lookup](#) helps read the codes of calls that were rejected along the way.

Checklist

1. Confirm which direction fails and on which calls: all of them, only inbound, only some.
2. Capture the SIP and read the SDP: is the address on the `c=` line the public one?
3. Turn off the SIP ALG on the router or firewall, and reset the active sessions.
4. Set `external_media_address`, `external_signaling_address` and every network in `local_net`.
5. Use `direct_media=no`, `rtp_symmetric=yes`, `force_rport=yes` and `rewrite_contact=yes` on the trunk and remote phones.
6. Align the `rtp.conf` range with the firewall's UDP forwarding, all of it.
7. Adjust the UDP session timeouts or shorten the registration renewal.
8. Compare the codecs and encryption of the offer and the answer.
9. Check what happens to the 183 and the early media.
10. Prove the result with real calls from outside to an echo or tone, repeatedly, watching the RTP counters for both directions.

Frequently asked questions

Why does a VoIP call connect and then stay silent?

Because SIP, which sets up the call, and RTP, which carries the audio, take different paths. The 200 OK proves SIP got through. The audio depends on the address and port announced in the SDP; if that address is private, the port is closed on the firewall or the two sides don't agree on a codec or encryption, the call is silent.

What is one-way audio?

It is a call in which one side hears and the other doesn't. The direction that fails tells you where to look: when the person inside can't hear the person outside, inbound audio isn't arriving, almost always because of a private IP in the SDP, a SIP ALG or an RTP range that isn't forwarded. When it's the other way round, look at the PBX's outbound side and the far end's SDP.

Does turning off the SIP ALG always fix it?

Not always, but it is the recommended first step when the PBX already handles NAT. The SIP ALG rewrites addresses inside SIP and the SDP, and the double correction produces mismatched addresses. After turning it off, the PBX has to announce the right external address and the RTP range has to be forwarded, or the problem stays.

Why does the call drop after about 30 seconds?

Almost always because the ACK never arrived. Under RFC 3261, the answering side retransmits the 200 OK until it receives the ACK and, after $64 \times T_1$, 32 seconds with T_1 at 500 ms, tears down the session. The ACK goes to the address in the Contact header; if that address is private, it gets lost in the NAT.

How do I test whether audio flows in both directions?

With a real call, coming in from outside through the trunk, to an echo or tone service: whatever goes out has to come back. An internal call only proves the leg between the phone and the PBX. Meanwhile, the PBX's RTP counters show whether packets are leaving and arriving in both directions.

What causes one-way audio on VoIP calls?

Almost always SIP completes the call but the RTP, which carries the audio, does not arrive. The most common causes are a private IP advertised in the SDP, SIP ALG rewriting addresses, an RTP port or range blocked at the firewall, and a codec or encryption the two ends do not agree on. The direction that fails shows where to look, and a real call from outside confirms the fix.

What is SIP ALG and why does it drop calls?

SIP ALG is a router and firewall feature that inspects and rewrites SIP traffic to help it cross NAT. It breaks calls when it rewrites the addresses in the SIP and SDP wrongly, especially if the PBX also handles NAT: the call connects but stays silent. Turning it off is the recommended first step when the PBX already handles NAT, although it does not always fix the problem.

Sources

1. RFC 3261: SIP, including retransmission of the 200 OK until the ACK ($64 \times T1$)
2. RFC 3264: the offer/answer model with SDP
3. RFC 3550: RTP and RTCP, packet counters and jitter
4. RFC 3960: early media and ringing tone generation in SIP
5. RFC 6314: NAT traversal practices for SIP applications
6. RFC 4787: NAT behavior for UDP, including mapping timeouts
7. RFC 8489: STUN

GUIDE

Voice circuit monitoring: SIP trunks, 0800 and IVR

A registered SIP trunk and a PBX with no alarms don't prove that a call connects with audio. This guide compares the ways to monitor a voice circuit, shows how to confirm an outage without false alarms and says what evidence to keep to hold the carrier accountable.

By [2N5 team](#) Published October 7, 2026 Updated October 20, 2026

Summary

- Ping and SIP OPTIONS show that a device responds; only a real test call shows that a customer's call connects, with audio in both directions.
- Passive capture on the PBX only sees the calls that reach it: when the failure is upstream, at the carrier or on the 0800's route, it records nothing.
- At Monitor's default pace, each circuit gets a test call every 15 minutes; after a failure, three retries every 2 minutes, and four failures in a row open an incident.
- Two good calls in a row confirm the circuit's recovery; an isolated failure never becomes an incident.
- When the failure is in the monitoring itself, the test is marked monitoring blind: it doesn't count against the circuit and stays out of the SLA.
- The evidence for each failure is the SIP code, the Q.850 cause, the Call-ID and the exact time of each test call.

What is voice circuit monitoring?

Voice circuit monitoring is the continuous check that every SIP trunk, 0800 (toll-free) number, branch DID and IVR connects calls with audio in both directions and with the right message. A voice circuit, here, is any number or route through which a customer reaches your company by phone.

The definition matters because it changes what gets measured. Most of the tools in a telephony environment look at the equipment: whether the [SBC](#) is up, whether the trunk is registered, whether the PBX's processor has headroom. Those questions are necessary, but they answer something else. A device can be in perfect shape while a customer's call never reaches it.

This guide deals with the question the customer asks: if I call now, will someone answer? The guides on [0800 and IVR](#), [silent calls](#), [voice quality](#) and [SLA](#) go deeper into each part.

Why does the dashboard stay green while the customer hears silence?

The dashboard stays green because it measures the equipment, not the call. The PBX is up, the SIP trunk shows as registered and the link responds, while the failure is on the carrier's route, in the 0800 translation or on the audio path. None of those points sends an alarm to your team's dashboard.

Three common situations show the problem:

- **The call never arrives.** The carrier stopped routing your [0800](#) or a DID to your trunk. The PBX receives nothing, so it has nothing to complain about.

- **The call arrives, but the audio doesn't.** The SIP signaling gets through and the call is answered, but the [RTP](#) that carries the voice stops at a NAT, a firewall or a [SIP ALG](#). The result is a silent call or one-way audio, and the PBX logs a completed call.
- **The wrong message answers.** The IVR was swapped, the holiday recording was left on or the carrier plays the number-does-not-exist announcement. As far as the signaling is concerned, everything worked.

There's an aggravating factor: in the small hours and on weekends, few calls is normal. The customer who calls at 2:00 a.m. and hears silence gives up, and the failure only surfaces when support opens and the queues are strangely empty.

Ping, SIP OPTIONS, passive capture or a test call: what does each one see?

Each method of monitoring a SIP trunk answers a different question. Ping says whether an address responds; SIP OPTIONS and registration, whether the SIP stack responds; passive capture, how the calls that arrived went. Only a real test call says whether a customer's call connects right now, with audio.

Method	What it sees	What it doesn't see	What it's for
Ping	That the SBC's or the PBX's address responds on the network	The SIP service, the carrier's route, the audio, the message; many SBCs don't even answer ping	Knowing whether the device has dropped off the network
SIP OPTIONS and registration (keep-alive)	That the SIP stack on the other side responds and that the trunk's registration is valid	Whether an INVITE to your number is routed to you, whether audio flows, who answers	Keeping the trunk and the NAT alive and noticing a lost registration
Passive monitoring (packet capture, RTCP statistics on the PBX)	SIP codes, loss and jitter of the real calls that reached the PBX	The calls that never arrived; nothing at all when there's no traffic	Diagnosing quality and investigating a specific call
Real test call	Whether the call connects, the SIP code and Q.850 cause, who answered, audio in both directions and measured quality	Origins and numbers it doesn't dial; what happens between two tests	Knowing whether customers can reach you, at any hour

The methods complement each other: the test call tells you **whether** it's working; the others help find out **why** it isn't.

Ping

Ping is the cheapest test and the most limited. It confirms that an IP address responds, which helps you know whether an SBC or a link has disappeared. But a device that answers ping can have its SIP service stopped, and many SBCs drop ping on purpose, for security.

SIP OPTIONS and registration

The OPTIONS method, defined in RFC 3261, asks a SIP element what it supports. In practice, SBCs and carriers send OPTIONS at intervals to find out whether the other side is alive and, along with other techniques described in RFC 6314, to keep NAT mappings open. REGISTER, also from RFC 3261, is how the trunk tells the carrier where to deliver calls. See [SIP registration](#).

A trunk that is registered and answering OPTIONS proves that signaling between the two devices works. It doesn't prove three things: that the carrier routes a specific number to that trunk (an 0800's translation happens somewhere else in the carrier's network); that audio flows, since RTP travels on other ports and can be blocked while OPTIONS gets through; and that the right IVR is the one answering.

Passive monitoring on the PBX

Packet capture on the PBX or the SBC records the signaling and audio of real calls. RFC 3550's RTCP reports carry the packet loss and [jitter](#) for each direction, and many devices use them to calculate an estimated MOS score with the ITU-T G.107 E-model, which ITU-T P.800.1 calls MOS-CQE. It's the right tool for investigating why a specific call sounded bad.

The limit is structural: passive monitoring only sees what arrives. If the carrier stopped delivering the 0800's calls, there is no packet to capture, and in the small hours, with no traffic, there is no data at all. A drop in call volume is no signal at 3:00 a.m. on a Sunday, when zero calls is what you'd expect.

Real test call

The test call dials the number from outside, the way a customer does, and judges the result. It crosses the carrier, the 0800 translation, the trunk, the SBC, the PBX and the IVR, and so it sees what the customer would see. It's the only method that works the same at 3:00 a.m. as at 3:00 p.m., because it doesn't depend on someone calling.

It has limits too. It tests from your origin and only the numbers it dials; a failure that affects only calls from mobile phones, for example, may not show up (see [why an 0800 won't connect from a mobile phone](#)). It doesn't see what happens between one test and the next. And it holds a channel for a few seconds, like a customer call.

What does a test call need to check?

A test call needs to check four things, in order: whether the call connected, whether whoever should answer did, whether audio flowed in both directions and whether the voice arrived with good quality. A test that only checks the first one lets through the silent 0800, the swapped IVR and the carrier's announcement.

Did the call connect?

The route's response to an INVITE says what happened. 1xx responses indicate progress, 200 OK indicates an answer, and 4xx, 5xx and 6xx responses indicate failure. The ones a test runs into most often:

SIP response	What it usually means	For monitoring
180 Ringing	The destination is ringing	Progress, not yet an answer
183 Session Progress	Progress with audio before the answer, early media (RFC 3960)	Could be ringback or a carrier announcement; not an answer
403 Forbidden	The route rejected the call	Failure
404 Not Found	The number doesn't exist for this route	Failure
408 Request Timeout	No final response in time	Failure
480 Temporarily Unavailable	The destination is unavailable right now	Failure

SIP response	What it usually means	For monitoring
486 Busy Here	Busy	Failure for an IVR, which should answer every call
487 Request Terminated	The call was canceled before it was answered	When the test gave up waiting, it's an answer failure
488 Not Acceptable Here	The media offer, such as the codecs, was not accepted (RFC 3264)	Failure
503 Service Unavailable	The service on the other side is unavailable	Failure
603 Decline	The destination declined the call	Failure

Along with the SIP code, many routes send the telephone network's Q.850 cause in the Reason header, as RFC 6432 describes. The [list of SIP codes](#), the [list of Q.850 causes](#) and the [SIP code lookup](#) explain each one.

Did the right party answer?

"Answered" isn't enough. An 0800 can connect and play the carrier's recording saying the number doesn't exist; an IVR can be swapped without notice. That's why each circuit needs a definition of what should answer, and the test needs to check it.

In Monitor, that choice is made during setup and decides what counts as working:

- **A greeting or recorded message, the IVR:** passes when the call is answered and the expected greeting is heard. 2N5 records each number's reference greeting, and each call compares what it heard with it. A carrier's error announcement doesn't pass.
- **A test tone or echo service:** passes when the call is answered and the tones sent come back.
- **Other:** may ring until timeout or be answered with audio, as agreed. If a person answers, they hear a message saying it's an automated test.

For an IVR, busy, no answer or ringing until timeout is a failure: it answers every call on its own.

Did audio flow in both directions?

A call can be answered and stay silent, or carry voice in one direction only. SIP signaling and audio take different paths: SIP agrees on the addresses and ports, and RTP carries the voice. A NAT that rewrites only one of the two, a SIP ALG that alters the message or a firewall that blocks the media ports produce a silent call, while the PBX logs a completed call. On a circuit with a test tone or echo, the test checks that the tones sent come back. The guide [silent calls and one-way audio](#) covers the causes in detail.

Did the voice arrive with good quality?

A call can connect and sound choppy. Monitor measures listening quality from the audio the call received, compared with the reference recording, using a reference-based perceptual listening-quality model, and gives a score from 1 to 5: [MOS-LQO](#), in ITU-T P.800.1 terminology. That's different from the MOS-CQE a PBX estimates from network figures with the G.107 E-model: one measures what arrived, the other calculates what should have arrived.

By default, a circuit is healthy with MOS-LQO from 3.8 up, and a score below 3.0 marks the call as degraded. A degraded call goes through the same retries as any failure. See [measured voice quality](#) and the [MOS and MOS-LQO in practice](#) guide.

How do I confirm an outage without false alarms?

A voice circuit outage is confirmed by repeating the call before alerting anyone. An isolated failure can be a route that blipped once; four failures in a row, a few minutes apart, are an outage. The same logic applies to recovery: one good call isn't enough, two in a row confirm it.

At Monitor's default pace, each circuit gets a test call every 15 minutes. When one fails, the call is repeated three times, every 2 minutes; four failures in a row open an incident, and two good calls in a row confirm recovery. As a timeline, with T as the time of the first failed test:

Moment	Call	Result	What happens
T - 15 min	Scheduled test	Connects, the IVR answers	Circuit healthy
T	Scheduled test	SIP 503	First failure; retry in 2 minutes; nobody is alerted
T + 2 min	Retry	SIP 503	Second failure in a row; retry in 2 minutes
T + 4 min	Retry	SIP 503	Third failure in a row; retry in 2 minutes
T + 6 min	Retry	SIP 503	Fourth failure in a row: incident opened, with the evidence; tier 1 is alerted by email
After that	Following calls	Fail	The incident stays open; if nobody acknowledges in time, the alert climbs to the next tier
On recovery	Call	Connects, the IVR answers	First good call; the incident doesn't close yet
Right after	Next call	Connects, the IVR answers	Second good call in a row: recovery confirmed, incident closed, contacts notified

If the T + 2 min call connects, the sequence stops there: there was an isolated failure, nobody was woken up and nothing goes into the SLA. That filter is what lets you trust an alert at 2:00 a.m.

How long does it take for an outage to become an alert?

It depends on when the outage starts relative to the test. In the worst case, it starts right after a good test: it takes up to 15 minutes for the next test to see the failure and another 6 minutes for the three retries to confirm it, so $15 + 6 = 21$ minutes to the incident. In the best case, it starts just before a test and becomes an incident in about 6 minutes.

Testing more often shortens that time, but each test holds a channel for a few seconds. The math at the default frequency: $24 \text{ hours} \times 4 \text{ tests per hour} = 96 \text{ test calls a day, per circuit, not counting retries.}$

Why four failures and two good calls?

Four failures in a row, within 6 minutes, separate a blip from an outage without delaying the alert too much. Two good calls in a row prevent the flapping of an unstable circuit, which would otherwise close and reopen the incident on every test. These are Monitor's default values.

When does the monitoring itself go blind?

Monitoring goes blind when the test can't run because of a problem on the tester's side: the calling platform not responding, the originating account down, the monitoring's own network out. In that case the test says nothing about the voice circuit, and can't count as an outage of it.

Counting blindness as an outage does three kinds of damage. The team gets woken at 2:00 a.m. for a problem that isn't theirs. The circuit's SLA comes out worse than it really was. And disputes with the carrier get weaker, because all the carrier has to do is show its side was normal at that hour to cast doubt on the other recorded outages. Blindness can't count as uptime either: if nobody could test, nobody knows what state the circuit was in.

In Monitor, when the failure is on 2N5's side, for example a problem on 2N5's side, such as our calling platform not responding, the test is marked as **no visibility**:

- it doesn't count against the circuit and doesn't open an incident for the customer;
- it stays out of the SLA: the period comes out of the calculation and is shown separately in the report;
- it doesn't wake the customer's team; 2N5 is alerted if it lasts more than 30 minutes;
- a periodic heartbeat also alerts 2N5 if the monitoring stops running.

Anyone building their own monitoring needs the same separation. A practical approach is a control destination, a number that should always connect, tested from the same origin. If the circuit and the control fail together, the problem is probably at the test's origin, not in the circuit.

What evidence should I keep for each failure?

For every failed test call, keep the SIP response code, the Q.850 cause when one is sent, the Call-ID, the exact time with time zone, the number called and the origin. With that, the carrier can find the call in its own records, and "no fault found" becomes hard to sustain.

Evidence	Where it comes from	Why it matters
SIP code	The response's status line (RFC 3261), for example 503	Says where and how the call was rejected
Q.850 cause	The Reason header, when the route sends it (RFC 6432), for example <code>Reason: Q.850;cause=34</code>	Carries the telephone network's cause behind the SIP
Call-ID	The call's Call-ID header (RFC 3261)	Identifies the call in the carrier's records
Time	The tester's clock, with seconds and time zone	Cross-checks with the carrier's records and marks the start and end of the outage
Number called and origin	The test itself	A failure may apply only to one origin or one number
Last healthy call	The test history	Bounds when the outage began
What was heard	The call recording, when the IVR or the recording answered	Proves the wrong message or the carrier's announcement

Keep the SIP code and the Q.850 cause together, call by call: the same 503 response can come with different causes, and each one points somewhere else.

In Monitor, each incident carries a block ready to paste into the carrier ticket: the circuit, the carrier, the first symptom, the SIP code and Q.850 cause of each call, and the last healthy call. An example, with demo data: *SIP 503 Service Unavailable, Q.850 cause 34 (no circuit/channel available), confirmed after 4 failed calls*. Audio is kept only when what answers is the recording or the IVR on file, encrypted and kept for the period 2N5 sets; a person's voice is never kept.

What to monitor: SIP trunks, 0800s, branch DIDs or the IVR?

Monitor what the customer dials first, then each path those calls take. In general, that means the support numbers, such as 0800 and 4004, the IVR that answers them, one DID per SIP trunk and per carrier, and one DID at each branch. Each one becomes a voice circuit with its own schedule.

What	What to dial	What should answer	What the test shows
Support number (0800, 0300, 4004)	The number itself, as the customer dials it	The IVR's greeting	The carrier's route, number portability, the IVR platform
IVR	The number that lands on it	The greeting recorded as the reference	A swapped message, the platform being down, quality
SIP trunk	A DID that comes in through that trunk	A test tone, an echo or a greeting	Carrier rejection, full channels, one-way audio
Branch DID	A number at the branch	Tone, echo, greeting, or an extension that may ring	That branch's trunk, link and PBX
Support queue	The queue's number	The greeting, or it may ring	Whether the queue receives calls

One trunk, one DID

If you have two SIP trunks from different carriers for redundancy, test each one with a DID that comes in only through it. Testing only the main number hides the first trunk's outage while the second one carries the calls, and you only find out when both go down together. Redundancy only counts if every path is tested.

Support numbers and the IVR

The support number carries the most weight, because it's the one customers know. Test the 0800 itself, not just the DID it translates to: the translation lives at the carrier, and that's where many failures happen. The guide [Your 0800 went down and nobody told you](#) covers that case in detail, and the [0800 and IVR monitoring](#) page shows how Monitor judges each number. For branches, one DID per branch, with an extension or a test tone behind it, covers each branch's trunk, link and PBX; see [SIP trunk monitoring](#).

How often should a voice circuit be tested?

The test frequency sets the resolution of a voice circuit's SLA. With one call every 15 minutes and retries every 2 minutes after a failure, the start and end of every outage are pinned down to within a few minutes. Testing once an hour can miss entire 40-minute outages.

The math starts from the target. A 99.9% target in a 30-day month allows $43,200 \text{ minutes} \times 0.001 = 43.2$ minutes of downtime. If the test runs every 60 minutes, a 40-minute outage can start and end between two tests and never make it into the report: the SLA comes out better than it really was. For the target to mean anything, the interval between tests has to be much shorter than the time it allows. The [SLA calculator](#) does the math for any target and period.

Two other decisions weigh as much as the interval:

- **The hours when nobody is watching.** A 24-hour 0800 needs testing overnight and on weekends, because that's when nobody on the team will notice on their own. In Monitor, each number has its own schedule, at intervals or at fixed times, including nights and weekends.
- **The call limit.** Each test holds a channel for a few seconds. With many circuits behind the same trunk, testing has to fit within a call limit so it never competes with customers for a channel.

What about the internet link the voice runs over?

The internet link is the data circuit that voice almost always runs over, and an outage there takes down every voice circuit at the site at once. Measuring the link alongside the test calls answers the first question of every voice outage: was it the line or the connection?

In Monitor, an agent that 2N5 installs at the site measures latency, [packet loss](#), jitter, DNS time and speed. Each voice circuit can be linked to its site's link; with the wait-for-the-link option, the team gets one alert about the link, not one per voice circuit that runs over it. See [data circuits](#).

Who should be alerted when a voice circuit goes down?

When a voice circuit really goes down, the people who need to know are the ones who can act at that hour, not an inbox nobody opens overnight. A tiered ladder that climbs until someone takes ownership ensures that a confirmed incident's alert doesn't go unanswered.

In Monitor, the incident alerts tier 1 by email. If nobody acknowledges within the set time, the alert climbs to the next tier; a single acknowledgment, one tap on the email link, stops the climb and records who took ownership. If no tier acknowledges, the company's administrators are alerted. The customer's contacts receive the problem notice and the resolution notice, outside the ladder. Maintenance windows silence the chosen circuits during the scheduled time. See [alerts and escalation](#).

The time to the first acknowledgment goes into the [SLA reports](#) alongside availability. Measure the call, confirm the outage, separate out the blindness, keep the evidence and alert the right person: that is what turns monitoring into a response.

Frequently asked questions

Are ping or SIP OPTIONS enough to monitor a SIP trunk?

No. Ping says the SBC's or the PBX's address responds on the network, and SIP OPTIONS says the SIP stack on the other side responds. Neither places a call: they don't show whether the carrier routes your number to you, whether audio flows in both directions or whether the right IVR answers. For that you need a real test call.

How often should a voice circuit be tested?

At an interval well below the downtime the target allows. A 99.9% target in a 30-day month tolerates 43 minutes; testing once an hour can miss an entire outage. At Monitor's default pace, the test call runs every 15 minutes, with retries every 2 minutes after a failure.

Why not alert on the very first failed call?

Because an isolated failure can be a route that blipped once, and a false alarm at 2:00 a.m. teaches the team to ignore the next ones. Repeating the call separates a blip from an outage: in Monitor, four failures in a row, by default (the first one and three retries), open an incident, and only a confirmed incident alerts anyone by email.

What is a no-visibility test?

It is a test that could not run because of a problem on the tester's side, for example the monitoring's calling platform not responding. It says nothing about the circuit. In Monitor, that period doesn't count against the circuit, stays out of the SLA and doesn't wake the customer's team; it is 2N5 that gets alerted.

What should I send the carrier when a SIP trunk goes down?

The number called, the origin, the time of each call with seconds and time zone, the SIP response code, the Q.850 cause when one was sent and the Call-ID. With that data the carrier can find each call in its own records. Include the last call that worked as well, to mark when the outage began.

Sources

1. RFC 3261: SIP, including the OPTIONS and REGISTER methods, response codes and the Call-ID
2. RFC 3264: the offer/answer model for media negotiation in SIP
3. RFC 3550: RTP and RTCP, including loss and jitter reports
4. RFC 3960: early media, the audio before the answer
5. RFC 6432: Q.850 causes in the Reason header of SIP responses
6. RFC 6314: NAT traversal practices for SIP, including keep-alive
7. ITU-T P.800.1: MOS terminology, including MOS-LQO and MOS-CQE
8. ITU-T G.107: the E-model and the R-factor

GUIDE

Telephony SLA: how to calculate availability by year and by month

The SLA for a SIP trunk or an 0800 (toll-free) number is a simple calculation, but full of details that change the result: the length of the period, scheduled maintenance, what counts as an outage and the resolution of the measurement. This guide walks through the calculation and each of those details.

By [2N5 team](#) Published October 7, 2026 Updated October 7, 2026

Summary

- Availability = (period – maintenance – downtime) ÷ (period – maintenance) × 100.
- 99.9% allows about 8 h 46 min of downtime in a 365-day year, and about 43 min in a 30-day month.
- In Brazil, telephony SLAs are usually measured over the year; the month serves as a reference and an early warning.
- An outage only enters the calculation if someone measured it: the more frequent the test, the more accurate the number.
- Losing visibility because of a problem in the monitoring itself is not a circuit outage and should stay out of the calculation.

Downtime each availability target allows

Period	99%	99.5%	99.9%	99.95%	99.99%
365-day year	3 d 15 h 36 min	1 d 19 h 48 min	8 h 45 min 36 s	4 h 22 min 48 s	52 min 34 s
Leap year (366 days)	3 d 15 h 50 min	1 d 19 h 55 min	8 h 47 min 2 s	4 h 23 min 31 s	52 min 42 s
31-day month	7 h 26 min 24 s	3 h 43 min 12 s	44 min 38 s	22 min 19 s	4 min 28 s
30-day month	7 h 12 min	3 h 36 min	43 min 12 s	21 min 36 s	4 min 19 s
Leap-year February (29 days)	6 h 57 min 36 s	3 h 28 min 48 s	41 min 46 s	20 min 53 s	4 min 11 s
February (28 days)	6 h 43 min 12 s	3 h 21 min 36 s	40 min 19 s	20 min 10 s	4 min 2 s
Week	1 h 40 min 48 s	50 min 24 s	10 min 5 s	5 min 2 s	1 min
Day	14 min 24 s	7 min 12 s	1 min 26 s	43 s	9 s

How do I calculate an availability SLA?

Divide the time the circuit worked by the time it was supposed to work. As a formula: **availability = (period – scheduled maintenance – downtime) ÷ (period – scheduled maintenance) × 100**. The period is counted in minutes, using the real calendar: a year has 525,600 minutes (527,040 in leap years), and a month has between 40,320 and 44,640.

Running it backwards answers the most common question in a contract negotiation: how much downtime a target allows. **Allowed downtime = (period – scheduled maintenance) × (100% – target)**. For 99.9% over a 365-day year, that is $525,600 \times 0.001 = 525.6$ minutes, or 8 hours, 45 minutes and 36 seconds.

The calculator and the table above use exactly this formula. The table assumes no scheduled maintenance; with maintenance, use the calculator.

How much downtime fits in each target?

The figures that show up most often in telephony contracts, over a 365-day year:

Target	Per year	Per 30-day month
99%	3 d 15 h 36 min	7 h 12 min
99.5%	1 d 19 h 48 min	3 h 36 min
99.9%	8 h 45 min 36 s	43 min 12 s
99.95%	4 h 22 min 48 s	21 min 36 s
99.99%	52 min 34 s	4 min 19 s

Each extra “9” divides the allowed downtime by ten. The difference between 99.9% and 99.99% is the difference between a whole afternoon down and less than an hour in the year. Before signing up to a high target, ask how it will be measured: a target nobody measures is just a number in the contract.

What counts as downtime?

This is the part the contract needs to spell out, and often doesn't. For a voice circuit, the situations usually treated as unavailability are:

- **the call doesn't connect:** the carrier rejects it with a SIP failure code, such as [SIP 503](#) or [SIP 404](#), or doesn't respond at all;
- **the call connects with no audio**, in one direction or both: to the caller, a silent 0800 is down;
- **the wrong message answers:** a swapped IVR, or the carrier's recording saying the number does not exist;
- **voice quality drops below the agreed level:** some contracts treat this as partial unavailability, others as a separate target.

And what normally stays out of the calculation:

- **scheduled maintenance**, announced and within the agreed window;
- **failures caused by the customer itself**, such as a PBX switched off at a branch, when the contract is with the carrier and the fault is not on the carrier's side;
- **time when the monitoring lost visibility** because of a problem of its own. A test that could not run says nothing about the circuit, and cannot be turned into unavailability.

Whoever measures has to separate these situations, and the report has to show them. In Monitor, a period in which the test could not run because of a problem on 2N5's side comes out of the calculation and is shown separately.

Yearly or monthly SLA: which one applies?

Whatever the contract says. In Brazilian corporate telephony, the availability target is commonly yearly, with each month's figures used as a reference. The two measures tell different stories:

	Yearly target	Monthly target
What it tolerates	A bad month, if the year closes within the target	No month outside the target

	Yearly target	Monthly target
Allowed downtime at 99.9%	8 h 45 min 36 s in the year	43 min 12 s in each 30-day month
When the penalty is calculated	At year end	Every month
What the monthly figure is for	Acting before the year blows through the target	It is the SLA itself

An example: a trunk is down for 6 hours in March and not at all for the rest of the year. Against a monthly target of 99.9%, March missed (6 hours against the 44 minutes and 38 seconds allowed in a 31-day month). Against a yearly target of 99.9%, the year closed at 99.932%, within the target.

That is why Monitor's reports show the **last 12 months** as the figure that decides the SLA, with the month and the last 30 days for reference.

How do I subtract scheduled maintenance?

Take the maintenance out of the period before the calculation, on both sides of the fraction. A 30-day month has 43,200 minutes; with a 4-hour maintenance window (240 minutes), the period that counts drops to 42,960. If the circuit was down for another 30 minutes outside the window, availability is $(42,960 - 30) \div 42,960 = 99.930\%$.

Two things to watch:

1. **The window has to be recorded beforehand**, with a start and an end. A "maintenance" declared after the outage is just an outage.
2. **Whatever runs past the window is an outage**. If maintenance scheduled to end at 4:00 a.m. ended at 5:30 a.m., the extra hour and a half goes into downtime.

In Monitor, a maintenance window covers the chosen circuits at a scheduled time, and failures on those circuits during the window do not open an incident.

How do I measure real downtime?

With a test that sees what the customer sees, at a frequency that fits the target. Two common mistakes make the SLA figure worthless:

- **Measuring the equipment, not the call**. A PBX that is powered on and a trunk that is registered do not prove that a call connects with audio. Downtime is only real if the test is a real call.
- **Measuring too rarely**. If the test runs once an hour, a 40-minute outage can go unseen, and the SLA comes out better than it was. For a monthly 99.9% target, which tolerates 43 minutes, the interval between tests has to be much shorter than that.

At Monitor's default pace, each circuit gets a test call every 15 minutes. When one fails, the call is repeated three times, every 2 minutes; four failures in a row open an incident, and two good calls in a row confirm recovery. That way the start and end of every outage are pinned down to within a few minutes, and an isolated blip does not enter the SLA.

What should I ask for in the carrier contract?

Before accepting a target, agree in writing on:

1. **the measurement period** (year or month) and the time zone;
2. **what counts as unavailability**: only calls that fail to connect, or also silent calls and quality below the threshold;

3. **how maintenance is announced** and the maximum window;
4. **who measures, and how:** the carrier's measurement, yours, or an independent one;
5. **what evidence is accepted** to open a dispute: the time, the number called, the SIP code and the Q.850 cause of each failure.

The last item carries the most weight when it is time to claim. An outage recorded with the SIP code and Q.850 cause of every test call is hard to dispute; an outage noticed because a customer complained is not.

What about customer service?

For SACs (the customer-service channels Brazilian law regulates) in regulated sectors, availability is not just a contract matter. Decree No. 11,034/2022 (Decreto nº 11.034/2022), until November 1, 2026, requires access to the SAC to be available 24 hours a day, 7 days a week, through at least one channel, and makes phone support mandatory; from November 2, 2026, Decree No. 13,135/2026 keeps phone support mandatory and requires at least one electronic channel available 24 hours. For telecom providers, Anatel's General Consumer Rights Regulation, Anatel Resolution No. 765/2023 (RGC), requires phone support to be available 24 hours a day. Measuring the availability of the 0800 is how you show those rules were met.

Frequently asked questions

What is the formula for an availability SLA?

Availability is the time the service worked divided by the time it was supposed to work, as a percentage: $(\text{period} - \text{maintenance} - \text{downtime}) \div (\text{period} - \text{maintenance}) \times 100$. A trunk that was down for 60 minutes in a 30-day month, with no maintenance, had 99.861% availability.

How much downtime is 99.9% availability?

In a 365-day year, 8 hours, 45 minutes and 36 seconds. In a 30-day month, 43 minutes and 12 seconds. In a week, 10 minutes and 5 seconds. In a day, 1 minute and 26 seconds.

Yearly or monthly SLA: which is better?

It depends on the contract. A yearly target tolerates a bad month as long as the year closes within it; a monthly one penalizes every bad month. In Brazilian corporate telephony the target is commonly yearly, with each month tracked for reference and to act before the year blows through it.

Does scheduled maintenance count as downtime?

Usually not, if it was announced with the notice the contract requires and stayed within the agreed window. It comes out of the period before the calculation. Maintenance that overruns its window counts as an outage from the moment the window ended.

How do I measure the downtime of an 0800 number?

By calling it at short intervals and recording when calls stop connecting and when they start connecting again. Precision depends on the interval: with one call every 15 minutes and retries every 2 minutes after a failure, the start and end of an outage are pinned down to within a few minutes.

Sources

1. ITU-T E.800: definitions of quality of service terms, including availability
2. Decree No. 11,034/2022 (Decreto nº 11.034/2022), art. 4: SAC access 24 hours a day, 7 days a week (applies until November 1, 2026)
3. Decree No. 13,135/2026, arts. 12 and 14: 24-hour electronic channel and mandatory phone support (applies from November 2, 2026)
4. Anatel, Resolution No. 765/2023 (RGC), art. 15: phone support available 24 hours

GUIDE

Voice quality in VoIP: MOS and MOS-LQO in practice

A call can connect, have audio in both directions and still sound choppy. The MOS score is the standard way to say how much, but there is more than one MOS, and they don't measure the same thing. This guide explains the scale, the ITU-T acronyms, what jitter, loss and delay do to voice, and what to do when the score drops.

By 2N5 team Published October 7, 2026 Updated October 7, 2026

Summary

- The ITU-T P.800 MOS scale runs from 5 (excellent) to 1 (bad) and was born from tests with human listeners.
- Under ITU-T P.800.1, MOS-LQO is listening quality calculated by an objective model; the output of the G.107 E-model is MOS-CQE, an estimate.
- A measured MOS compares the received audio with a reference recording, as P.863 (POLQA) does; P.862 (PESQ) was withdrawn on January 5, 2024.
- In the E-model, the default values give $R \approx 93.2$; $R\ 80$ corresponds to MOS 4.03, $R\ 60$ to 3.10, and the maximum MOS-CQE is 4.5.
- G.114 recommends at most 400 ms of one-way delay for planning; below 150 ms, delay is essentially transparent.
- In Monitor, a circuit is healthy with MOS-LQO from 3.8 up and degraded below 3.0, thresholds 2N5 adjusts per circuit.

What is MOS in VoIP?

MOS in VoIP is the *Mean Opinion Score* of a call's quality, on a scale from 1 to 5. The score was born from tests in which listeners rated audio samples, using the methods of ITU-T P.800. Today almost every MOS score is calculated by a model that tries to predict what those listeners would say.

Searching for "MOS score" gets you the same thing: the S already stands for *score*, and the phrase is simply how the acronym gets used in practice. What changes from one tool to another is what goes into the calculation. A PBX can show 4.2 for a call that a listening model rates at 3.1, and both numbers can be right, because they measure different things: one estimates quality from the network, the other listens to the audio that arrived. Before comparing two scores, you need to know what kind each one is.

What does the ITU-T P.800 MOS scale look like?

The ITU-T P.800 MOS scale has five categories, from 5 (excellent) to 1 (bad). In the absolute category rating method, each listener hears a sample and picks a category; the average of all the listeners' scores is the MOS for that sample. It is an average of opinions, not a physical measurement of the signal.

Score	Category	Label in P.800
5	Excellent	Excellent
4	Good	Good
3	Fair	Fair
2	Poor	Poor

Score	Category	Label in P.800
1	Bad	Bad

Two practical consequences. First, because it is an average, the score rarely reaches the extremes: it only takes some listeners not giving a 5 for the average to fall below that, even with flawless audio. Second, each model that predicts MOS has its own useful range. The E-model's estimated MOS, for example, tops out at 4.5. That is why a "4.3" from one tool and a "4.3" from another are not automatically the same quality.

What do MOS-LQO, MOS-CQE and the other acronyms mean?

MOS-LQO, MOS-CQE and the rest come from ITU-T P.800.1, which names each type of MOS by two questions: what was assessed, listening or conversation, and how the score was obtained, with listeners, with a model working on the signal or with a model working on network parameters. Each combination has its own acronym.

	Subjective: listeners	Objective: model on the audio	Estimated: model on parameters
Listening quality (LQ)	MOS-LQS	MOS-LQO	MOS-LQE
Conversational quality (CQ)	MOS-CQS	MOS-CQO	MOS-CQE

The difference between listening and conversation matters more than it seems. Listening quality assesses audio in one direction: whether the voice arrived clearly. It doesn't see delay, because a listener who only listens can't tell whether the reply was slow. Conversational quality includes what gets in the way of a real conversation, such as delay and echo. A high [MOS-LQO](#) with 500 ms of delay is possible: the voice sounds clean and the conversation is poor.

The output of the ITU-T G.107 E-model is a MOS-CQE: an estimate of conversational quality, made from parameters. A measurement such as P.863's is a MOS-LQO: listening quality, calculated on the audio. Calling both of them just "MOS" is the source of most arguments about voice scores.

What is the difference between measured MOS and estimated MOS?

The difference between measured MOS and estimated MOS is what goes into the calculation. Measured MOS, a MOS-LQO, comes from the audio the call received, compared with the expected recording. Estimated MOS, a MOS-CQE, comes from network figures, such as packet loss, jitter and delay, fed into a model that never hears the call.

Measured MOS: the audio against the reference

Reference-based models take two signals, the reference and the degraded one, and calculate how far the second has drifted from the first in a way a listener would notice. The most cited:

- **P.863 (POLQA)**: the ITU-T recommendation for objective listening quality, which returns a MOS-LQO.
- **P.862 (PESQ)**: the previous recommendation, withdrawn by the ITU-T on January 5, 2024 and replaced by P.863. It still shows up in older equipment and reports. See [POLQA and PESQ](#).

Because they listen to the audio, these models see what the network doesn't report: voice clipped at the source, noise, botched codec transcoding, the real effect of packet loss after the phone tried to conceal it, and even a wrong message in place of the greeting, because the comparison with the reference fails.

Estimated MOS: the E-model and the R-factor

The ITU-T G.107 E-model is a transmission planning tool. It adds up a route's impairments into an index, the R-factor:

$$R = R_o - I_s - I_d - I_{e,eff} + A$$

R_o is the basic signal-to-noise ratio, I_s the impairments simultaneous with the voice signal, I_d the impairments caused by delay and echo, $I_{e,eff}$ the equipment impairment, which includes the codec and packet loss, and A the advantage factor, the user's tolerance in exchange for mobility, for example. With all of G.107's default values, $R \approx 93.2$.

G.107 converts R into MOS-CQE with a formula valid for R between 0 and 100: **MOS = $1 + 0.035 \times R + R \times (R - 60) \times (100 - R) \times 7 \times 10^{-6}$** . For R of 100 and above, MOS-CQE is 4.5, the maximum. For the default R : $1 + 0.035 \times 93.2 + 93.2 \times 33.2 \times 6.8 \times 7 \times 10^{-6} = 1 + 3.262 + 0.147 \approx 4.41$.

R-factor	MOS-CQE
100 or more	4.50
93.2 (default values)	4.41
90	4.34
80	4.03
70	3.60
60	3.10
50	2.58

Note that the E-model's top score for a perfect route sits at 4.41 with the default values. Equipment that shows a MOS of 4.4 for every call on a lossless network is showing the model's ceiling, not a measured quality.

When to use each

	Measured MOS-LQO	Estimated MOS-CQE
Input	The received audio and the reference recording	The route's loss, jitter, delay and codec
Needs a reference	Yes	No
Sees delay	No: it is listening quality	Yes, through the I_d term
Sees audio clipped at the source or a wrong message	Yes	No
Good for	Knowing whether the voice reached the caller clearly	Planning a route and tracking the network

Both have their place. A network estimate is cheap and works for looking at every call on the system. A call that is actually listened to tells you whether the voice got through.

What MOS score is good in VoIP?

A good MOS score in VoIP, on a measured MOS-LQO, is around 4: the voice sounds clear and nobody has to strain to understand. Below 3, people start to complain. These figures apply to measured listening quality; an estimated MOS-CQE uses a different calculation and shouldn't be compared with them directly.

The table describes, roughly, what you hear in each MOS-LQO band, and how Monitor reads each band with the default thresholds:

MOS-LQO	What you hear	How Monitor reads it
4.0 or more	Clear voice, effortless; differences only show side by side with the reference	Healthy
3.8 to 4.0	Clear, with some loss of crispness noticeable if you pay attention	Healthy
3.0 to 3.8	Intelligible, with short dropouts, hiss or a metallic voice; the conversation takes attention	In service, with the score flagged as below ideal
2.0 to 3.0	Lost words and requests to repeat; this is the complaint zone	Degraded
Below 2.0	Conversation is hard or impossible	Degraded

The right threshold depends on the use. An IVR that only plays a greeting tolerates less degradation than it seems, because the customer can't ask it to repeat itself; an internal trunk between branches can live with a slightly lower score. That is why Monitor's thresholds are set per circuit.

How much jitter is acceptable?

Acceptable jitter is whatever the phone's or the PBX's jitter buffer can absorb without discarding packets and without adding too much delay. The recommendations cited here don't set a single number: the same jitter that a large buffer absorbs without loss turns into choppiness in a small buffer, and the large buffer charges its price in delay.

Jitter is the variation in delay from one packet to the next. Voice packets leave at regular intervals, for example one every 20 ms, and arrive at irregular ones. The jitter buffer holds packets for a few milliseconds and plays them out at the right pace. A packet that arrives after the buffer's deadline is discarded, and to the listener that is the same as a lost packet.

RFC 3550 defines how RTP calculates interarrival jitter: with each packet, the estimate moves 1/16 of the difference between the measured variation and the previous estimate, $J = J + (|D| - J) \div 16$. This is the value RTCP reports carry and many devices display. It is a smoothed average: short jitter spikes, which are the ones that overflow the buffer, barely show up in it.

The math that matters is total delay. An example, with illustrative values: 40 ms of network, 20 ms of packetization and a 60 ms jitter buffer add up to $40 + 20 + 60 = 120$ ms one way. Doubling the buffer to absorb more jitter takes the total to 180 ms, above G.114's essentially transparent range. A measured MOS-LQO shows the result of that trade-off in the audio; the jitter figure alone doesn't.

In Monitor, the jitter of each site's internet link is measured by the [data circuit](#) agent, along with latency, packet loss, DNS time and speed.

How much packet loss can voice take?

The packet loss voice can take depends on the codec, on the phone's ability to conceal the missing packet and on how the loss happens. Scattered loss, a packet here and a packet there, is concealed better than burst loss, when several packets in a row vanish and a whole syllable disappears.

The G.107 E-model captures this in the $I_{e,eff}$ term, the equipment impairment with loss: $I_{e,eff} = I_e + (95 - I_e) \times Ppl \div (Ppl \div BurstR + Bpl)$. I_e is the codec's impairment without loss, Ppl the percentage of packets lost, Bpl the codec's robustness to loss and $BurstR$ how concentrated the loss is in bursts. The higher $BurstR$, the greater the impairment for the same percentage of **packet loss**. The formula makes it clear that the same loss weighs differently on different codecs, which is why a single "acceptable loss" rule is misleading.

A reference-based measurement sidesteps the argument: it listens to the result after concealment, the buffer and the codec, and scores what arrived. The choice of **voice codec** also counts: each codec transcoding along the path degrades the audio a little, and a low-bitrate codec starts with a lower score before any loss.

How much delay is acceptable on a VoIP call?

Acceptable delay on a VoIP call, under ITU-T G.114, is up to 400 ms one way for network planning purposes. Below 150 ms, delay is essentially transparent for most applications. Between the two values, the conversation works, but people start talking over each other and waiting for replies.

Delay doesn't spoil the sound, it spoils the conversation. The voice arrives clean and late. That is why a MOS-LQO, which is listening quality, doesn't see delay, and the E-model, which estimates conversational quality, does, through the I_d term. Delay also makes echo more annoying: an echo that arrives quickly blends with your own voice, and one that arrives late is heard as a second voice.

The components of one-way delay, so you can do the math for your route:

- **packetization**: the amount of audio in each packet, for example 20 ms;
- **codec processing** on both sides;
- **network**: transit time, which the data circuit agent measures as latency on the site's link;
- **jitter buffer**: the time the receiver holds packets;
- **transcoding along the path**: each SBC or PBX that decodes and re-encodes adds its own time.

Why does the reference recording matter?

The reference recording matters because a reference-based model compares two signals: what should have been heard and what was heard. Without the reference, there is nothing to compare against, and the score becomes an estimate. With it, the same comparison that measures quality also shows when the wrong message answered.

The reference has to be what should actually play: the IVR greeting, recorded clean, or a known test tone. If the IVR changes its greeting on purpose, the reference has to change too, or every call starts to look wrong. In Monitor, 2N5 records each number's reference during setup. The audio kept is that of the recording or the IVR, never a person's voice, and it is kept for the period 2N5 sets.

How do I measure the voice quality of a SIP trunk, 0800 or IVR?

You measure the voice quality of a SIP trunk, an 0800 (toll-free) number or an IVR by calling the number and listening to the audio that arrives, compared with the reference. A real call sees the whole path, carrier, trunk, PBX and IVR, which no single piece of equipment sees.

That is what Monitor does on every test call for each **voice circuit**, the monitored SIP trunk, 0800 or support number. 2N5 records the audio the call received and compares it with the expected recording using a reference-based perceptual listening-quality model, which returns a MOS-LQO (ITU-T P.800.1). Quality is measured on the audio; Monitor does not estimate the score from the network.

With the default thresholds, a circuit is healthy with MOS-LQO from 3.8 up and degraded below 3.0; 2N5 sets the thresholds with you, per circuit, during setup. A score below the threshold marks the test call as degraded, and it goes through the same retries as any failure. If it persists, a quality incident opens, the escalation ladder is alerted, and the customer's contacts also hear about the problem and its resolution.

Each call's score appears on the circuit's page, with the threshold lines, on the dashboard next to the circuit's status and in the [SLA reports](#). See the details in [measured voice quality](#).

What should I do when the MOS score drops?

When the MOS score drops, the first step is to confirm that the drop repeats and to see its reach: a single circuit, several at the same site or several on the same carrier. The reach tells you where to look, and the evidence from each call is what gets the problem to whoever can fix it.

What you see	Likely cause	What to check
Low score on every circuit at the same site	The site's internet link: loss, jitter or congestion	Latency, loss and jitter on the data circuit at the same time
Low score only on one carrier's circuits	The carrier's route or codec	The negotiated codec and the evidence from each call
Low score at peak hours	Saturated outbound link, with no priority for voice	Link usage and voice traffic prioritization
Constantly low score, with no network loss	Low-bitrate codec, codec transcoding, level or noise at the source	The codecs along the path and the IVR recording itself
Suddenly very low score, with the call connecting	A message different from the reference, or one-way audio	The call's audio and the silent calls guide

In practice, the sequence:

1. **Confirm.** A single low score is not a trend. In Monitor, a degraded call is retried before any incident.
2. **Check the reach.** Compare the circuits at the same site and those on the same carrier at the same time.
3. **Look at the link.** If the site has a monitored data circuit, the charts for the last 100 hours show latency, loss and jitter at the moment of the drop.
4. **Check codecs and transcoding.** Every decode and re-encode along the path costs score. See whether the call goes through transcoding that could be avoided.
5. **Check voice priority.** On a shared link, voice needs priority over the rest of the traffic, especially in the upload direction, which is usually the narrowest.

6. **Take it to the carrier with evidence.** The time of each call, the number called, each call's score and the last healthy call.

What should I ask for in a voice quality contract?

In a voice quality contract, ask for the score to be defined by its full name: which MOS, in P.800.1 terminology, which model and version, against which reference, measured where and how often. Without that, "a minimum MOS of 4" applies just as well to a network estimate as to a listening measurement, and they give different numbers.

The items that make the target verifiable:

1. **The type of MOS:** measured MOS-LQO or estimated MOS-CQE.
2. **The model:** P.863 or another, and the version; P.862 has been withdrawn and shouldn't be the basis of a new contract.
3. **The reference:** which recording or tone is compared.
4. **The measurement point:** a call from outside, like a customer's, or equipment inside the network.
5. **The violation rule:** how many below-threshold scores in a row count as degradation, and whether degradation enters the SLA calculation as partial unavailability or as a separate target. The [telephony SLA](#) guide shows how each choice changes the math.

Frequently asked questions

What is a good MOS score for VoIP?

It depends on which MOS. On a measured MOS-LQO, voice sounds clear at around 4, and below 3 people start to complain. In Monitor, the default is healthy from 3.8 up and degraded below 3.0. On an E-model estimated MOS, the maximum is 4.5, and the values are not comparable with a measured score.

What is MOS-LQO?

It is the listening quality score, from 1 to 5, calculated by an objective model from the received audio. The acronym comes from ITU-T P.800.1: MOS, mean opinion score; LQ, listening quality; O, objective. Models such as P.863 (POLQA) compare the received audio with a reference and return a MOS-LQO.

How much jitter is acceptable on a VoIP call?

Whatever the jitter buffer can absorb without discarding packets and without pushing total delay past the budget. The recommendations cited here do not set a single jitter figure. G.114 gives the ceiling for one-way delay, 400 ms for planning, and the buffer's delay counts toward it.

Which measurement model should I use to compare scores?

Any reference-based model that returns a listening quality score, such as P.863 (POLQA), the ITU-T recommendation that replaced P.862 (PESQ). What matters is stating which model was used when comparing scores, because each one has its own scale and its own ceiling.

Why is my PBX's MOS different from the measured MOS?

Because the PBX usually estimates the score from network figures, such as packet loss, jitter and delay, using the G.107 E-model. That estimate doesn't listen to the audio. It can't see audio clipped at the source, poor codec transcoding or a wrong message, all of which a reference-based measurement sees.

What is the difference between MOS and MOS-LQO / POLQA?

MOS is the 1 to 5 scale of voice quality; MOS-LQO is a score on that scale measured on the audio that arrived, compared with a reference recording. POLQA is one of the models that compute MOS-LQO, alongside PESQ, which the ITU has withdrawn. Scores from different models cannot be compared, and a MOS estimated from the network is not a MOS-LQO.

Sources

1. ITU-T P.800: subjective assessment methods and the MOS scale
2. ITU-T P.800.1: MOS terminology (LQS, LQO, LQE, CQS, CQO, CQE)
3. ITU-T G.107: the E-model and the R-factor
4. ITU-T G.114: one-way transmission time
5. ITU-T P.863: POLQA, objective listening quality
6. ITU-T P.862: PESQ, withdrawn on January 5, 2024
7. RFC 3550: RTP and RTCP, including the interarrival jitter calculation

GUIDE

What is synthetic call monitoring

Synthetic call monitoring is a robot that calls your 0800 number, your IVR or your SIP trunk on a schedule, the way a customer would, and records what happened: whether the call connected, what answered and how the audio sounded. It answers the question no equipment dashboard answers: can a customer reach the company right now?

By [2N5 team](#) Published October 14, 2026 Updated October 20, 2026

Summary

- A synthetic call is a test call placed by a system, not a person, on a schedule and always the same way.
- It measures what the customer experiences: the call connects, the right IVR answers, audio flows both ways and the voice quality is good.
- Passive monitoring analyzes the real calls that already happened; synthetic monitoring also sees the hour when nobody called.
- To count as a measurement the call has to be real: a PBX that is powered on and a registered trunk do not prove that a call connects.
- In Monitor by 2N5, each voice circuit gets a test call every 15 minutes by default, and only four failures in a row open an incident.

What is synthetic call monitoring?

Synthetic call monitoring is the practice of deliberately calling, from a system and on a schedule, the numbers customers dial, and recording the result of each call. The call is "synthetic" because a robot places it, not a customer; it is real because it travels through the carrier's network for real, the same way a customer's call does.

Each test call follows a fixed script: dial the number, wait for the answer, listen to what answers, compare it with what should answer and hang up. What gets recorded is what matters to the people who run the telephony: whether the call connected, how quickly, with which SIP code, and how the voice sounded.

What is a test call for?

Every hour, or every fifteen minutes, it answers one question: **could a customer reach the company right now?** The PBX and trunk dashboards answer another question, whether the equipment is up, and the two answers can differ.

What happened	What the PBX dashboard shows	What the test call shows
The carrier stopped routing the 0800 to the trunk	All green, no calls	The call does not connect, with the SIP code returned
The IVR was swapped or the platform published a faulty flow	All green, calls answered	The greeting heard does not match the reference recording
The carrier answers with a "number does not exist" message	No calls	The call connects, but the message is not the IVR's

What happened	What the PBX dashboard shows	What the test call shows
The call connects, but audio goes one way only	Normal call	Missing or weak audio in the measurement
Voice quality dropped on the route	No alarm	MOS-LQO below the limit

An 0800 that spends the night down without anyone noticing is the classic case. The guide [Your 0800 went down and nobody told you](#) walks through a night hour by hour.

How does a synthetic call work?

In four steps, repeated on each circuit's schedule:

1. **Call.** The system dials the monitored number (0800, 0300, 4004, a landline number or a SIP trunk number) from outside the customer's network.
2. **Listen.** It records whether the call connected, the SIP code and Q.850 cause returned, and the audio received.
3. **Compare.** It checks what answered against the [IVR](#) reference recording and measures the voice quality of the audio received, the [MOS-LQO](#).
4. **Decide.** An isolated failure is retried; only consecutive failures become an incident and an alert.

The last step is what separates useful monitoring from a false-alarm generator. At Monitor by 2N5's default pace, each circuit gets a call every 15 minutes. When one fails, the call is repeated three times, 2 minutes apart; four failures in a row, the first and three retests, open the incident, and two good calls in a row confirm the recovery. The guide [Line-down alerts, without false alarms](#) describes that path.

How is it different from passive monitoring?

Passive monitoring watches the real calls that cross the network, by capturing traffic, and computes indicators about them. Synthetic monitoring places calls on its own. The practical consequence is simple: passive monitoring only sees what happened when there was a call, and at night, on a weekend or on a rarely dialed number there may be none.

The two complement each other, and neither replaces the other. The guide [Passive monitoring vs active test calls](#) compares them and says when each one fits.

What does a synthetic call not see?

It sees what happens from the origin it calls from. A failure that only affects customers of one mobile carrier, or one region, can go unnoticed if the test call leaves from another path; the guide [Why an 0800 won't connect from a mobile phone](#) covers that case. It also does not replace analyzing real calls to find the cause inside your own network.

Which terms show up with the topic?

People who search for it find "test call", "active monitoring" and "availability test" in Portuguese, and "synthetic calls" and "synthetic voice monitoring" in English. They are the same concept. What changes from tool to tool is what it calls (just the number, the whole IVR, the SIP trunk), where it calls from and what it does with the result.

How does Monitor by 2N5 do it?

Monitor by 2N5 is a managed service: 2N5 registers each circuit, records the IVR reference greeting, sets the schedule and places calls from Brazil to 0800 numbers, IVRs and SIP trunks. Each call becomes a record with time, number called, SIP code, Q.850 cause and MOS-LQO, and the set becomes the [SLA report](#) for the month and the last 12 months. See [how it works](#) and [0800 and IVR monitoring](#).

Frequently asked questions

What is a synthetic call?

It is a test call started by a system, on a schedule, that follows the same script every time: dial a number, listen to who answers, sometimes press an IVR option and hang up. The result of each call is recorded with its time, SIP code, Q.850 cause and audio assessment, and becomes the basis of the measured availability.

Is synthetic monitoring the same as active monitoring?

In telephony, in practice, yes. Both terms describe tests the monitoring system starts itself, as opposed to passive monitoring, which only observes the calls customers already made. Synthetic is the name used in website and network monitoring, and active is the name more common among people who work with VoIP.

Why use synthetic calls if I already monitor the PBX?

PBX monitoring shows the equipment from the inside. A synthetic call shows the whole path from the outside, from the carrier to the answer. A broken carrier route, a faulty number portability change or a swapped IVR leave the PBX all green, because no call reaches it. Only a call from outside sees that.

Do synthetic calls get in the way of my trunk channels?

Hardly at all. Each test call takes up a channel for a few seconds, and Monitor works within an agreed call limit so it does not compete with your customers for channels. The default frequency is one call every 15 minutes per circuit, which comes to 96 short calls a day.

How often should synthetic calls run?

The shorter the interval, the less time an outage goes unseen. At 15 minutes, an outage is seen within 15 minutes and confirmed 6 minutes later by the retests. For a 99.9% monthly target, which tolerates 43 minutes, a test once an hour is too little.

Sources

1. [RFC 3261: SIP, the protocol of the test calls and its response codes](#)
2. [RFC 6432: Q.850 causes in the Reason header of SIP responses](#)
3. [ITU-T P.800.1: MOS terminology, including MOS-LQO](#)
4. [ITU-T E.800: quality of service definitions, including availability](#)

GUIDE

Why an 0800 number won't connect from a mobile phone when it works from a landline

When the company's 0800 (toll-free) number answers from the office landline but won't connect from a mobile phone, the problem is almost never the answering side. It sits somewhere between the caller's carrier and the number's configuration. This article lists the causes in the order worth checking them.

By [2N5 team](#) Published October 7, 2026 Updated October 14, 2026

Summary

- An 0800 call is free for the caller; the subscriber pays, and a call from a mobile phone usually costs the subscriber more.
- Many 0800 services let the subscriber choose which origins the number accepts calls from; a forgotten restriction knocks out mobile calls.
- Each origin follows its own route to the 0800: one route can fail while the others work.
- Test from at least two mobile phones on different carriers, and write down the time, the calling number and what you heard.
- For telecom providers, Anatel requires the phone support line to take calls from both landlines and mobile phones.

Why does the 0800 work from a landline but not from a mobile phone?

Because a call from each origin takes a different path to the 0800, and it only takes one of those paths failing. A call from a mobile phone goes through the caller's mobile carrier, which has to hand it to the carrier that owns the 0800, which translates the 0800 into the real answering number and only then delivers the call to the company's trunk and IVR. An origin restriction, a misconfigured route or a translation set up for only some origins knocks out mobile calls without touching landline calls.

Are “not working from a mobile phone” and “not connecting from a mobile phone” the same problem?

In practice, yes. Someone who says the 0800 “does not work from a mobile phone” is usually describing a call that does not connect: silence, a busy signal, a drop after a few seconds or a carrier announcement. If the whole 0800 is down, the failure also applies to the landline; see the guide [Your 0800 went down and nobody told you](#). If only mobile phones fail, the causes below are the starting point.

What are the most common causes?

Cause	What a mobile caller hears	Where to check
Origin restriction on the 0800 service	A carrier announcement saying the call cannot be completed, or a busy signal	The 0800 contract and configuration with the carrier that owns the number
0800 translation set up only for some origins or areas	A number-does-not-exist announcement, or silence	The carrier that owns the 0800

Cause	What a mobile caller hears	Where to check
A faulty route between the mobile carrier and the 0800 carrier	Silence, or the call drops after a few seconds	The caller's mobile carrier and the 0800 carrier
The trunk or IVR rejects the format of the calling number	A busy signal, or an error message from the PBX itself	Inbound rules on the PBX or SBC (A-number format)
Trunk capacity exhausted at peak hours	Busy at certain times	Contracted concurrent channels and trunk occupancy

Origin restriction

An 0800 call is free for the caller: under Anatel's numbering regulation, the subscribing organization pays for the call. Since a mobile call usually costs the subscriber more, many 0800 services let it choose which origins the number accepts calls from. A restriction set up to save money, or inherited from an old contract, makes the number work from a landline and fail from a mobile phone. It is the first thing to ask the carrier.

Translation and routing

Behind every 0800 there is a real number. The carrier that owns the 0800 holds that translation and can configure it by origin or by region, for example to send each state's calls to a different support center. If the translation for mobile calls, or for one region, is wrong or missing, only those calls fail. The same happens when the route between the caller's mobile carrier and the 0800 carrier has an interconnection problem.

The PBX or SBC itself

Calls reach the company's trunk carrying the caller's number, and that number comes in different formats depending on the origin: with or without the leading 0 and the carrier selection code, with the area code, sometimes in international format. An inbound rule that accepts only one format, or that uses the calling number to pick the destination, can reject mobile calls. In the PBX log they show up as rejected, for example with [SIP 403](#) or [SIP 484](#).

How do I test before opening a ticket?

Place test calls from several origins at the same time and write down the result of each one.

1. **Call from a landline and from mobile phones on at least two different carriers**, within the same 10 to 15 minute window.
2. **For each call, write down** the calling number, the carrier, the time with its time zone and what you heard: a normal answer, silence, a busy signal, a carrier announcement (copy the wording).
3. **Check the PBX or SBC log** for the calls that should have arrived. If there is no record, the call never reached your trunk and the problem is upstream, at the carrier. If there is a record with a rejection, the problem is in your inbound rule.
4. **Open the ticket with the carrier that owns the 0800**, attaching the table of calls. Ask it to check the origin restriction and the number translation.

What do the rules require of the support line?

For SACs (the customer-service channels Brazilian law regulates) in regulated sectors, Decree No. 11,034/2022 (Decreto nº 11.034/2022), until November 1, 2026, and Decree No. 13,135/2026, from November 2, 2026, guarantee free access to the SAC and make phone support mandatory. For telecom providers, Anatel's General Consumer Rights Regulation, Anatel Resolution No. 765/2023 (RGC), is explicit: phone support must be free, available 24 hours a day and take calls placed from landlines and from mobile phones. An 0800 that won't connect from mobile phones leaves out exactly the customer who only has a mobile phone.

What about monitoring?

Monitoring with real calls shows whether the 0800 connects and whether the IVR answers with the right message, from the origin it calls from. Monitor calls from the 2N5 platform: it alerts you when the number stops connecting or the IVR changes, but it does not test each mobile carrier. For a restriction that only affects mobile phones, the manual test above is still the way to go. See how [0800 and IVR monitoring](#) works.

Frequently asked questions

Is calling an 0800 number from a mobile phone free?

Yes. Under Anatel's numbering regulation, on the 0800 code the subscribing organization takes responsibility for payment, and the call is free for the caller, from a landline or a mobile phone. That is why the cost of a mobile call falls on the company that owns the number.

Why do some companies block mobile calls to their 0800?

Because a mobile call to an 0800 costs more for whoever pays, which is the company. Some 0800 contracts let the subscriber choose which origins to accept. A restriction set up to save money, or inherited from an old contract, keeps customers who only have a mobile phone from getting through.

How do I prove to the carrier that the 0800 doesn't connect from mobile phones?

Place test calls from mobile phones on at least two carriers and from a landline, at the same time. For each one, write down the calling number, the time with its time zone and what you heard: silence, a carrier announcement, a busy signal. With that, the carrier can find the route that failed.

Does monitoring that calls the 0800 catch this problem?

It does when the failure affects the origin it calls from. Monitor calls from the 2N5 platform, so it shows whether the 0800 connects and whether the IVR answers with the right message from that origin. It does not replace testing from each mobile carrier.

My 0800 doesn't work from a mobile phone but works from a landline. What does that mean?

It means the failure applies only to calls arriving through the mobile network, not that the number is down. The most common causes are an origin restriction on the 0800 service, a translation set up only for some origins, a faulty route between the mobile carrier and the 0800 carrier, or a PBX inbound rule that rejects the number format.

Why won't my 0800 connect calls from mobile phones?

Almost always because the fault is in the route or in the number rules, not because the 0800 is switched off. The most common causes are an origin restriction on the 0800 service that blocks the mobile network, a translation set up only for some origins, a faulty route between the mobile carrier and the 0800 carrier, or a PBX inbound rule. Test from mobile phones on two carriers and from a landline.

Sources

1. Anatel, Resolution No. 749/2022 (Regulamento de Numeração, the numbering regulation), art. 18: the 0800 code is free for the caller
2. Anatel, Resolution No. 765/2023 (RGC), art. 15: free phone support, 24 hours, taking calls from fixed-line (STFC) and mobile (SMP) networks
3. Decree No. 11,034/2022 (Decreto nº 11.034/2022), arts. 3 and 4: free access to the SAC and mandatory phone support (applies until November 1, 2026)
4. Decree No. 13,135/2026, arts. 8 and 14: free access to the SAC and mandatory phone support (applies from November 2, 2026)
5. RFC 3261: SIP response codes

Legal

Version 1, in effect from 2026-10-06 [Version 1](#) [Version 2](#) [Version 3](#)

Terms of Use

Voice Circuits is a service of **2N5 Global Enabler**. This page has two parts:

1. **2N5 Global Enabler's Terms of Use**, which apply to every 2N5 service.
2. **Voice Circuits**, the conditions specific to this service.

Part 2 adds to Part 1. Where the two differ, Part 2 applies to Voice Circuits. A signed contract or commercial proposal between 2N5 and the Client prevails over both.

Part 1 — 2N5 Global Enabler Terms of Use (LGPD – Brazil)

Valid for operations in Brazil – applicable legislation: LGPD (Law 13,709/2018), Marco Civil da Internet (Brazilian Internet Civil Framework), and Anatel Regulations.

Welcome to 2N5 Global Enabler ("2N5", "we", "our"). These Terms of Use govern the use of our services and platforms, including telecommunications solutions (STFC/VoIP), Cloud PBX, IVR, Omnichannel Integrations, APIs, and Artificial Intelligence.

By using our services, the User declares to have read, understood, and fully agreed to these Terms, as well as to our [Privacy Policy](#), which is part of this document.

1. Definitions (LGPD)

For the purposes of these Terms:

- **Personal Data:** information related to an identified or identifiable natural person.
- **Sensitive Data:** special categories provided for in Art. 5, II, of the LGPD.
- **Controller:** natural or legal person responsible for decisions regarding the processing of personal data. (The Client is the Controller.)
- **Processor:** person who processes personal data on behalf of the Controller. (2N5 acts as the Processor.)
- **Data Protection Officer (DPO):** person appointed to act as a communication channel between the Controller, Data Subjects, and the ANPD.
- **User:** any individual or legal entity that uses 2N5's services.
- **Client:** User who contracts a plan or service.
- **Platform:** 2N5's web environment, APIs, apps, and digital systems.

2. Acceptance and Term

Use of the services implies full and unrestricted acceptance of these Terms. The contract takes effect upon first access, registration, or subscription.

3. Eligibility

Only the following may use the services:

- duly registered companies;

- individuals over 18 years of age with full legal capacity;
- authorized legal representatives.

4. Services Provided

2N5 offers:

4.1. Telecom Services (STFC/VoIP)

In accordance with Anatel regulations, including:

- Phone numbers (DID), portability, and interconnection;
- Inbound and outbound traffic;
- Anti-fraud;
- Per-minute rate, subscription, or allowance.

4.2. Cloud PBX

- Extensions, IVR, queues;
- call recording in accordance with the law (the Client is responsible for obtaining consent);
- reports and dashboards.

4.3. Smart IVR

- advanced flows;
- dynamic routing;
- DTMF/voice collection.

4.4. Artificial Intelligence Services

- virtual agents;
- transcription, semantic analysis, classification;
- service automation;
- natural-language productivity tools.

4.5. APIs and Omnichannel Integrations

- telephony, SMS, WhatsApp (when authorized), email, webhooks;
- native or custom integrations.

5. User Responsibilities

The User agrees to:

1. Comply with Brazilian law, including Anatel rules, the LGPD, and the Marco Civil da Internet.
2. Not engage in spam, abuse, or fraudulent or illegal calls.
3. Keep registration data accurate and up to date.
4. Protect passwords, API tokens, and credentials.
5. Obtain consent from their own customers for call recording, use of AI, data collection, and any purpose required by the LGPD.
6. Be fully responsible for data entered into the platform.

6. Specific Conditions — Telecom, Cloud PBX, IVR, and AI

6.1. Use of Telecom

The following is prohibited:

- unlawful masking of Caller ID;
- fraud attempts, abusive robocalls, spam, and illegal telemarketing;
- use for criminal purposes or contrary to Anatel rules.

2N5 may immediately suspend the service in case of abuse.

6.2. Call Recording

- The feature is optional.
- The Client is solely responsible for informing and obtaining consent from the Data Subject.
- 2N5 only stores and processes data according to the Controller's instructions.

6.3. Artificial Intelligence

The User acknowledges that:

- AI operates on statistical models and may contain inaccuracies;
- Responsibility for decisions made based on AI results lies exclusively with the Client;
- The Client must have adequate legal bases to process personal data via AI (Legitimate Interest, Consent, Performance of Contract, etc.);
- 2N5 is not responsible for biases, misinterpretations, or errors arising from improper use of AI.

6.4. WhatsApp and Messaging

- Subject to Meta/WhatsApp policies;
- Irregular accounts may be blocked without refund;
- 2N5 does not control Meta's rules and is only an integrator.

6.5. Security and Logs

2N5 keeps logs for the legal period:

- 6 months under the Marco Civil da Internet (Art. 15);
- additional periods upon court order.

7. Data Protection (LGPD)

7.1. Processing Roles

- Client = Controller
- 2N5 = Data Processor

7.2. Data Processing by 2N5

2N5 processes data only in accordance with:

- the Client's instructions;
- the contract;
- applicable law.

7.3. Security Measures

2N5 adopts:

- encryption;
- access control;
- logical segregation;
- backups;
- internal security policies;
- certified data centers (ISO, SOC, etc.).

7.4. Storage and Retention

Data is retained for the period necessary for the purpose, regulatory compliance, and contract performance.

7.5. Sharing

Data may be shared with:

- telephone carriers;
- technology providers;
- authorities upon legal order.

8. Payments, Billing, and Credits

- Postpaid, prepaid, or hybrid plans;
- Credits are non-refundable;
- Late payment may result in immediate suspension;
- Overage calls and additional services are charged by usage.

9. Intellectual Property

All code, infrastructure, trademarks, and content are the exclusive property of 2N5. Copying, redistributing, reverse engineering, or reproducing parts of the systems is prohibited.

10. Limitation of Liability

To the extent permitted by law, 2N5 is not responsible for:

- loss of profits, revenue, or data;
- indirect damages;
- unavailability caused by third parties (carriers, providers, Meta, etc.);
- misuse of the services;
- decisions made by AI.

2N5's total liability is limited to the amount paid by the Client in the last 90 days.

11. Suspension and Termination

The service may be suspended for:

- breach of these Terms;
- abusive use;

- fraud;
- late payment;
- court order.

The Client may close the account at any time. Amounts paid are non-refundable.

12. Technical Support

Provided via tickets, email, or the admin panel. Plans with an SLA may be contracted separately.

13. Changes to these Terms

2N5 may update this document at any time. Continued use implies automatic agreement.

14. DPO Contact

To exercise your rights, clarify questions, or request additional information, contact:

- **Data Protection Officer (DPO)**, 2N5 Global Enabler
- **Email:** privacidade@2n5.com.br
- **Phone:** +55 (11) 2050-1096
- **Website:** www.2n5.com.br

Part 2 — Voice Circuits

1. Acceptance

The Client accepts these Terms when it signs the contract or proposal. Each User accepts them, together with the Privacy Policy, the first time they sign in, and 2N5 records the version accepted and when. Users act on behalf of the Client.

2. The service

Voice Circuits is a **managed monitoring service**. Depending on what the Client contracts, it:

- **places test calls** to the Client's telephone numbers and judges whether each line connects and carries audio of acceptable quality;
- **measures the Client's data circuits**, with an agent installed at the Client's site or with checks made from outside;
- **checks the Client's websites and email delivery**;
- **alerts the people the Client chooses**, step by step, when a failure is confirmed, and tells them when it is resolved;
- **produces service reports**, including availability figures.

2N5 sets up the monitoring from the information the Client provides. The Client's Users may also see it and, where allowed, change it.

The service **measures and reports**. It does not repair lines, carriers' networks or websites, and it is not an emergency service. A test can miss a problem, or report one that a later test does not confirm. Decisions the Client takes on the strength of a result are the Client's.

3. Limits of the plan

The Client's plan limits the number of voice circuits, data circuits, websites and email checks, and the number of test calls per hour. When a limit is reached, nothing more of that kind can be added until 2N5 raises it. Monitoring of what already exists continues.

4. Test calls

1. The Client may ask 2N5 to call only numbers it is **entitled to have called**, such as its own lines or lines it manages with the owner's agreement.
2. A test call that may be answered by a person first plays a recorded message saying it is an automated test.
3. **Voices are never kept.** Audio is kept only when the call is answered by the recording or automated menu (IVR) the Client told 2N5 to expect, and only to prove the line and measure its quality. Audio from a call answered by anyone or anything else is destroyed when the call ends.
4. The Client tells 2N5 when a number stops being its own, and plans maintenance windows for planned work, so that alerts and figures stay correct.
5. 2N5 may slow down or stop test calls that could harm a carrier or a third party, or that break telecommunications rules.

5. The monitoring agent

2N5 installs the data-circuit agent only on a computer the Client chooses, and **only with the Client's written authorization**. The agent sends network measurements and basic facts about the computer, as listed in the Privacy Policy. It opens no port and reads no files or traffic. The Client may ask 2N5 to remove it at any time.

6. The Client's responsibilities

In addition to Part 1, section 5, the Client must:

1. tell its staff and other contacts that their details are used to send them alerts and reports;
2. own, or have the right to use, the recordings it supplies for comparison;
3. tell 2N5 at once about any unauthorized access to its accounts.

7. Personal data

1. **Roles.** For the monitoring data the Client provides or the service produces for it, the Client is the controller and 2N5 the processor, as Part 1, section 7 says. For demo requests, sign-in accounts and access and security records, 2N5 is the controller.
2. **As processor, 2N5:**
 - uses only the kinds of providers listed in the Privacy Policy, and tells the Client before adding a new kind;
 - helps the Client answer requests from data subjects and from the ANPD;
 - tells the Client **without undue delay, and within 48 hours at most**, of a security incident affecting the Client's data;
 - transfers data outside Brazil only under a mechanism allowed by LGPD Art. 33.

8. Availability

2N5 aims to keep the service available at all times, but does not guarantee it. Planned maintenance is announced in advance where possible. Any availability commitment or service credit applies only if it is written in the contract.

9. The Client's data

The Client keeps all rights to its data and recordings, and allows 2N5 to use them only to provide the service. After the contract ends:

- the Client may ask for an export of its reports and results within **30 days**;
- 2N5 then deletes the Client's data within **90 days**, including backups as they expire;
- 2N5 keeps only what the law requires, such as access records for 6 months.

10. Changes

2N5 publishes any change to these Terms with its date and tells the Client and Users at least **30 days** before it takes effect. A Client that does not accept a material change may end the contract before it takes effect.

Privacy Policy

Voice Circuits is a service of **2N5 Global Enabler**. This page has two parts:

1. **2N5 Global Enabler's Privacy Policy**, which applies to every 2N5 service.
2. **Voice Circuits**, which says exactly what this service collects, why, and for how long.

Part 2 adds to Part 1. Where the two differ, Part 2 applies to Voice Circuits.

Part 1 — 2N5 Global Enabler Privacy Policy (LGPD, GDPR, PIPEDA and CCPA/CPRA)

1. Introduction

2N5 Global Enabler is committed to protecting personal data and to transparency in the processing of information. This Privacy Policy sets out how we collect, use, share, store, and protect personal data, in compliance with:

- Lei Geral de Proteção de Dados – LGPD (Brazil)
- General Data Protection Regulation – GDPR (European Union)
- Personal Information Protection and Electronic Documents Act – PIPEDA (Canada)
- California Consumer Privacy Act / California Privacy Rights Act – CCPA/CPRA (USA – California)

By accessing our services, the data subject declares to have read, understood, and agreed to this Policy.

2. Important Definitions

For the purposes of this Policy:

- **Personal Data:** any information related to an identified or identifiable natural person.
- **Sensitive Personal Data:** racial origin, religious belief, health, biometrics, etc.
- **Data Subject:** the person to whom the data refers.
- **Controller:** the party that decides on the processing (2N5 Global Enabler).
- **Operator/Processor:** the party that carries out the processing on behalf of the controller.
- **DPO (Data Protection Officer):** responsible for the communication channel with data subjects and authorities.
- **Consent:** free, informed, and unambiguous expression by the data subject.

3. Personal Data Collected

3.1. Information provided directly by the data subject

- First name, last name
- Email, phone, WhatsApp
- Company, job title, and business information
- Data entered in forms, chat, or email
- Support and customer service records

3.2. Information collected automatically

- IP address (anonymized when possible)
- Browsing data, approximate geolocation, and online identifiers
- Cookies, tags, and tracking pixels
- Device and system data
- Log records, site usage, and access behavior

3.3. Information from third parties

We may receive information from:

- Analytics platforms (Google Analytics, Meta, LinkedIn)
- Marketing and CRM tools
- Business partners and integrated platforms

4. Legal Bases for Processing

2N5 Global Enabler carries out processing based on legal grounds, including:

4.1. LGPD – Bases

- Consent
- Performance of contract
- Legitimate interest
- Compliance with legal/regulatory obligation
- Regular exercise of rights

4.2. GDPR – Bases

- Consent
- Contractual necessity
- Legitimate interests (when not overriding the data subject's rights)
- Legal obligation
- Protection of vital interests (exceptional situations)

4.3. PIPEDA – Principles

- Informed consent
- Limiting collection and retention
- Accountability and transparency
- Appropriate safeguards

4.4. CCPA/CPRA – Rights and Obligations

- Transparency
- Opt-out of sale/sharing of data
- Limiting use of sensitive data
- Prohibition of discrimination against data subjects who exercise their rights

5. Purposes of Processing

We process personal data to:

- Provide and improve services
- Customer service and technical support
- Manage accounts, integrations, and communications
- Perform contracts and business obligations
- Analytics, metrics, and operational intelligence
- Authorized marketing, remarketing, and advertising campaigns
- Prevent fraud, incidents, and suspicious activity
- Comply with legal, tax, and regulatory obligations
- Security, auditing, and internal control

We do not use personal data for exclusively automated decisions that produce significant legal effects without human intervention, except where permitted by law.

6. Sharing and International Transfer

Data may be shared with:

- Service providers (hosting, analytics, cloud computing, automation, telephony, CRM)
- 2N5 group companies
- Business partners under a confidentiality agreement
- Government authorities, when legally required

International Transfer

Data may be stored or processed on servers in Brazil, Canada, the USA, Europe, or other countries. We ensure adequate mechanisms, such as:

- Standard Contractual Clauses – SCC (EU)
- Specific contractual clauses
- Privacy policy compatible with GDPR, LGPD, PIPEDA, and CCPA
- Data Processing Agreements – DPA

7. Data Retention and Deletion

We keep data only for the period necessary for:

- Fulfillment of the purposes of this Policy
- Legal and regulatory obligations
- Defense in judicial or administrative proceedings

After this period, data is deleted or anonymized according to secure technical standards.

8. Data Subject Rights

8.1. Rights under the LGPD

- Confirmation of processing
- Access

- Correction of incomplete data
- Anonymization, blocking, or deletion
- Portability
- Information about sharing
- Revocation of consent
- Objection to processing

8.2. Rights under the GDPR

- Access
- Rectification
- Erasure ("right to be forgotten")
- Restriction of processing
- Data portability
- Objection
- Not to be subject to automated decisions

8.3. PIPEDA Rights

- Access
- Correction
- Revocation of consent
- Transparency

8.4. CCPA/CPRA Rights

- Right to know
- Right to delete
- Right to correct
- Right to opt-out of sale/sharing
- Right to limit sensitive data use
- Non-discrimination

Rights may be exercised by contacting the DPO directly.

9. Security Measures

We use technical, administrative, and organizational measures consistent with international standards:

- Encryption at rest and in transit
- Access controls and multi-factor authentication
- Firewalls, IDS/IPS, and 24/7 monitoring
- Data Loss Prevention (DLP)
- Network segmentation
- Secure backup and redundancy
- Internal policies and periodic audits
- Employee training

10. Cookies

We use essential, analytics, functional, and marketing cookies. For more information, see our Cookie Policy.

11. Updates

This Policy may be updated periodically. The most recent version will always be available on the website.

12. Contact – Data Protection Officer (DPO)

To exercise your rights, clarify questions, or request additional information, contact:

- **Data Protection Officer (DPO)**, 2N5 Global Enabler
- **Email:** privacidade@2n5.com.br
- **Phone:** +55 (11) 2050-1096
- **Website:** www.2n5global.com

Part 2 — Voice Circuits

Voice Circuits monitors a company's voice circuits, data circuits and online services, and alerts the right people when one fails. It is sold to **companies** as a managed service, and 2N5 sets up each customer's monitoring from the information the customer provides.

1. Who decides how data is used

Data	Controller	2N5's role
Demo requests from our website, the accounts people use to sign in, access and security records	2N5	Controller
Monitoring data a customer gives us or the service produces for it: alert and report contacts, numbers called, recordings, results and reports	The customer company	Processor (operador): we use it only on the customer's instructions

If your details were given to us by your employer or another company, for example as an alert contact, that company is the controller. We will help it answer your request.

2. What we collect, why, and for how long

2.1. Booking a demo on our website

- **What:** your name, work email, company (optional), language, what you want to monitor and its approximate size.
- **Not collected:** we do not store your IP address, browser details or any tracking identifier. Your IP address is held in memory for a few minutes only, to stop abuse of the form.
- **Why:** to reply to you and arrange the demo.
- **Legal basis:** steps taken at your request before a contract (LGPD Art. 7, V).
- **Kept for:** 12 months after our last contact with you, unless you become a customer. You can ask us to delete it sooner.

2.2. Your sign-in account

- **What:** your name, email and language; your password, stored only as a one-way hash and never in readable form; the parts of your organization you can see; when you last signed in.
- **Why:** to let you sign in, show you your organization's monitoring and keep the account secure.
- **Legal basis:** performance of the contract with your organization (Art. 7, V) and our legitimate interest in keeping the service secure (Art. 7, IX).
- **Kept for:** as long as you have access. The account is deleted with your organization or at your request.

2.3. Contacts who receive alerts and reports

- **What:** name, role, work email, phone and WhatsApp numbers, language and the hours you do not want to be called. We also keep the messages we sent you and whether they were delivered or acknowledged.
- **Why:** to tell the right people when a line or service fails, and to send the reports your organization asked for.
- **Legal basis:** the customer's instructions, as controller (Art. 39).
- **Kept for:** while your organization keeps you as a contact. Messages already sent are kept for 13 months, so that yearly service reports can show who was told and when.

2.4. Test calls

- **What we keep:** the number called, the call's technical record (signalling, timings and result codes), the verdict and the measured quality.
- **Before a person can hear it:** a call that may reach a person first plays a recorded message saying it is an automated test.
- **We never keep the voice of a person who answers.** Audio is kept only when the call is answered by the recording or automated menu (IVR) the customer told us to expect, and only to prove the line works and to measure its quality. If anything else answers, including a person, the audio is destroyed when the call ends and only the result is kept.
- **Kept audio:** stored encrypted and deleted after 90 days, or the shorter period the customer sets.
- **Legal basis:** the customer's instructions, as controller (Art. 39).
- **Kept for:** technical call records 90 days; verdicts and quality figures for the length of the contract, so that service reports and yearly figures are complete.

2.5. Data-circuit monitoring agent

At the customer's written request, 2N5 installs a small agent on a computer at the customer's site.

- **What it sends:** the computer's name, operating system and agent version; the computer's local IP address, the router's address and the public IP the measurements come from; network measurements such as latency, loss, jitter, DNS time and speed; processor and memory use.
- **What it never collects:** files, browsing, traffic contents, user names, programs, Wi-Fi names or other devices on the network.
- **Legal basis:** the customer's instructions (Art. 39).
- **Kept for:** 13 months.

2.6. Website and email checks

- **Website checks** visit only public addresses the customer gives us.
- **Email checks** send one test message an hour to a test mailbox the customer creates for the purpose. They read only the headers of our own test message, to see whether and where it arrived, then delete it. No other message in that mailbox is read.
- **Legal basis:** the customer's instructions (Art. 39).
- **Kept for:** results for the length of the contract.

2.7. Access and security records

- **What:** for each request, the IP address, date and time, and what was requested; for each sign-in, whether it succeeded; a record of changes made in the service: who did what, and when.
- **Legal basis:** legal obligation under Marco Civil da Internet Art. 15 (Art. 7, II) and legitimate interest in security (Art. 7, IX).
- **Kept for:** access records 6 months, or longer only when an authority or court orders it; the record of changes for the length of the contract.

3. Cookies and storage in your browser

Part 1's analytics, functional and marketing cookies are **not used** in Voice Circuits. It has no analytics, no advertising, no tracking pixels, no third-party cookies, and loads nothing from other websites, so it does not ask you to accept cookies. It stores only what it needs to work:

Cookies and browser storage

Name	Kind	Purpose	Lifetime	Flags
<code>vcia_session</code>	cookie	Keeps you signed in	Until you sign out or it expires	HttpOnly, Secure, SameSite=Lax
<code>vcia.language</code>	cookie + local storage	Remembers English or Portuguese	1 year	Secure, SameSite=Lax
<code>vcia.signed-in</code>	local storage	Shows Open the app on the front page	Until you sign out	—
<code>vcia.sidebar</code>	local storage	Remembers whether the sidebar is open	Until cleared	—
<code>vcia.theme</code>	local storage	Remembers light, dark or your computer's setting	Until cleared	—
<code>vcia.staleBuildReloadAt</code>	session storage	Stops a reload loop after an update	Until the tab closes	—
<code>vcia.toast</code>	session storage	Carries a confirmation across one page load	One page load	—
<code>vcia.mailbox-link</code>	session storage	Holds a test mailbox's connect link while it signs in	Until the tab closes	—

You can delete these in your browser settings. Without `vcia_session` you cannot stay signed in.

4. Who we share data with

We **do not sell** personal data and do not use it for advertising. We share it only with:

- **Service providers that run Voice Circuits for us:** hosting and cloud infrastructure, email delivery, text-message and voice-call alerts, and the telephone carriers that carry the test calls. Each is bound by contract to use the data only for this purpose and to protect it.
- **The customer's own email provider**, when the customer connects a mailbox to send email or to run an email check.
- **The customer company** that the monitoring data belongs to.
- **Public authorities**, when a law or a court order requires it.

Some of these providers store or process data outside Brazil, including in the United States. We make these transfers under LGPD Art. 33, using the **standard contractual clauses approved by the ANPD** or another mechanism the ANPD allows.

5. How we protect data

- All traffic to the service is encrypted, and browsers are told to use encryption every time.
- Call audio, recordings, stored passwords for connected services, and signing keys are each encrypted in the database. The key that encrypts them is never kept on the server's disk.
- Every organization's data is kept separate, and a manager sees only the parts of their organization they were given.
- 2N5's administrators sign in with two-step verification. Every manager can turn it on, and an organization can require it of all its managers.
- Repeated failed sign-ins pause sign-in for that address, and its owner is told by email. Changing a password signs out every other session.
- Administrative access to the servers is possible only over a private network. Daily backups are kept for 14 days.

If a security incident may cause relevant risk or harm, we will tell the ANPD and the people affected within 3 working days of learning that personal data was affected. For data we hold for a customer, we will tell that customer without delay.

6. Decisions made automatically

Voice Circuits decides automatically whether a **line, link, website or mailbox** is working. It does not make automated decisions about people. You may still ask for a decision based only on automated processing to be reviewed (LGPD Art. 20).

7. Children

Voice Circuits is a service for companies. It is not meant for anyone under 18, and we do not knowingly collect their data.

8. Your rights

You have the rights in Part 1, section 8. How to use them with Voice Circuits, and how long we take to answer, is on our [LGPD](#) page. You may also complain to the **Autoridade Nacional de Proteção de Dados (ANPD)** at gov.br/anpd.

9. Changes

When this page changes in a way that matters, we publish the new version here with its date and tell signed-in users at least 30 days before it takes effect. Earlier versions stay available at their own address.

10. Contact

Write to the **Encarregado / DPO** of 2N5 Global Enabler, **Cláudio Longo Dacome Jr**, at privacidade@2n5.com.br, or call **+55 (11) 2050-1096**.