Service level agreement.

The availability, support, maintenance and recovery we commit to on Enterprise plans, and the credits if we miss them.

In force since 13 November 2026

Arcolabs OÜ, registry code 17587452, Sepapaja tn 6, 15551 Tallinn, Estonia · hello@testee.ai

This agreement sets the service levels Arcolabs OÜ ("we") commits to for organizations on an Enterprise plan whose order refers to it ("you"). It is part of your agreement with us, together with the Terms of service and the Data processing agreement.

We wrote it for what our setup actually delivers today: testee runs in one Google Cloud region, on one PostgreSQL database with continuous backups, with a small team that answers support on business days. Where we can't promise more, it says so.

Availability target.

Each of these is available at least 99.5 % of every calendar month:

- **Candidate interviews.** Candidates can open their invitation link, start or resume an interview, and send their answers.
- **The dashboard and the public API.** Your team can sign in and work in the dashboard, and your systems can use the public API.

99.5 % allows about 3 hours 40 minutes of downtime in a 30-day month. Interviews that are running when a server restarts resume from where they were, so short interruptions rarely end an interview.

How we measure it.

- **Uptime checks.** Every 5 minutes, checks from the United States, Europe and Asia-Pacific open the candidate app, the dashboard and each of our services through the same address you use.
- **Downtime.** A 5-minute interval counts as down for a service level above when checks from at least two of the three locations fail for it, or when our incident records show it was unusable for most users, whichever shows more.
- **The result.** Availability is the share of the month's 5-minute intervals that weren't down, excluding the intervals below. We send you the month's figures on request.

What doesn't count as downtime.

- **Scheduled maintenance** announced as described below.
- **AI provider outages.** Outages of the AI providers the interviewer and analysis depend on (OpenAI, and Anam for the video avatar) that are beyond our control. We tell you about them as we do about our own incidents, and results catch up once the provider is back: the analysis waits and retries instead of failing.
- **Causes on your side.** Your or your candidates' networks, devices, browsers or browser extensions; your configuration, such as SSO, IP allowlists or webhooks pointed at systems that don't answer; use against the Terms or our documentation.
- **Limits working as designed.** Plan limits, rate limits (429) and the fair-use ceiling, and a suspension under the Terms, for example for unpaid invoices.
- **Other services.** Features marked beta or preview, the trial, and plans other than Enterprise.

Outages of Google Cloud, which runs testee, are not excluded: they count as our downtime.

Support and response times.

Write to hello@testee.ai from an address in your organization, with "Sev 1", "Sev 2" or "Sev 3" in the subject. Business hours are 09:00 to 17:00 Estonian time (EET/EEST), Monday to

Friday, except Estonian public holidays. The times below are for our first answer by a person, counted in business hours.

- **Sev 1, critical:** candidates can't take interviews, nobody can use the dashboard, or you suspect your data is exposed. First answer within 2 business hours, then updates at least every 4 business hours until it's resolved. Outside business hours our alerts reach the team and we act as soon as we can, but we don't promise a time.
- **Sev 2, major:** an important feature fails, such as results, invitations or an integration, and there is no reasonable workaround. First answer within 1 business day.
- **Sev 3, minor:** a question, a request, or a problem with a workaround. First answer within 2 business days.

We may change the severity you chose if it doesn't match these definitions, and tell you why.

Incident communication.

- **When it happens.** For an incident that affects you at Sev 1, we email your organization's owners and the contacts you named within 2 business hours after we confirm it, whether or not you reported it.
- **While it lasts.** Updates at least every 4 business hours, and a message when it's resolved.
- **Afterwards.** A written post-incident report within 5 business days of any Sev 1 incident and of any month below the target: what happened, its impact on you, the cause, and what we change so it doesn't happen again.
- **Personal data breaches** follow the Data processing agreement: we tell you within 48 hours.

Maintenance.

- **Window.** Planned maintenance that may interrupt testee happens on Sundays between 02:00 and 06:00 UTC, and at most 4 hours a month.
- **Notice.** We email your organization's owners at least 5 business days ahead, with the expected impact.

- **During it.** Interviews already running are allowed to finish; new interviews and changes in the dashboard pause, and a notice says when we'll be back.
- **Urgent fixes.** If a security issue can't wait, we may do maintenance at shorter notice. That time counts as downtime.
- **Most updates** roll out without any interruption and need no notice.

Service credits.

If a service level's monthly availability falls below the target, you get a credit on the monthly fee for the month it happened:

- **Below 99.5 %:** 10 % of the monthly fee.
- **Below 99.0 %:** 25 % of the monthly fee.
- **Below 95.0 %:** 50 % of the monthly fee.

Credits for one month are capped at 50 % of that month's fee, even if both service levels were missed. The monthly fee is what you pay for the Enterprise plan for that month, before tax and without extra usage.

To claim a credit, email hello@testee.ai within 30 days after the end of the month, with your organization's name and the dates and times affected. We confirm or explain our figures within 10 business days and deduct the credit from your next invoice. Credits can't be exchanged for cash.

Credits are your sole remedy for missing the availability target. If availability falls below 99.0 % in three months in a row, you may also end the Enterprise order without penalty, with 30 days' notice.

Backups and recovery.

Your data is kept in a PostgreSQL database in one zone of the region, backed up every day, with a continuous log of changes for point-in-time recovery. Recovery point (RPO) is how much recent data could be lost; recovery time (RTO) is how long it takes to get testee working again with the restored data.

- **A zone fails:** we restore the database into another zone of the same region from the backups and the change log. RPO 5 minutes, RTO 8 hours.
- **Data damaged by a fault or an error on our side:** we restore the database to a moment before the damage, any time in the last 7 days. RPO 5 minutes, RTO 24 hours.
- **The whole region is lost:** we rebuild testee in another region of the same area (EU for EU organizations) from the daily backups, which are stored across several regions of that area. RPO 24 hours, RTO 72 hours.
- **Recordings, audio and whiteboards** are stored in Google Cloud Storage, replicated within the region. They aren't part of the database backups.
- **Backups are kept 30 days** and then deleted. Data you or your candidates erased stays erased after a restore.

Relation to the Terms and the DPA.

- This agreement adds to the Terms of service; where the two differ on service levels, this agreement applies, and your Enterprise order applies over both.
- The Data processing agreement governs personal data, including security measures, sub-processors and breach notification.
- We may improve this agreement at any time. We won't reduce the service levels during your order's committed term, and we announce any other change at least 30 days ahead, as for the Terms.