Service Level Agreement (SLA)

Last Updated: 2026-11-05 Effective Date: 2026-11-05

1. Scope

1.1. This Service Level Agreement (**"SLA"**) of Genz IT Solutions GmbH (**"Provider"**) applies to the production infrastructure of the paid HotCodePush Service. It covers two surfaces, each measured separately:

- **Update Delivery**: everything an end-user device reads to receive an update, in particular the channel index, the bundle manifest, files, and packs.
- **Control Plane**: the API through which every release passes, as used by the command-line tool, the console and the MCP server.

1.2. The following are not covered:

- SDK runtime behaviour on end-user devices (no uptime concept),
- Features expressly designated as Beta,
- The user interface of the console,
- The endpoint that receives device events,
- The website, the documentation, and the status page itself.

1.3. This SLA applies only to the latest major release of each SDK. Downtime resulting from the use of unsupported major versions is excluded.

1.4. For the Free plan, the Availability Target is a target without commitment. Service Credits are not granted on the Free plan (see Section 5.1).

2. Availability Target

2.1. The Availability Target is **99.9%** per calendar month for each surface under Section 1.1. It applies to every paid plan.

2.2. **Fail-Open**: Downtime of either surface never breaks an installed application. End-user devices keep running the update they already have and only stop receiving new updates until the Downtime ends. The Provider designs the HotCodePush Service so that a device never depends on the availability of the Control Plane to run.

3. Measurement and Calculation

3.1. Availability is calculated per surface and per calendar month (UTC):

Availability (%) = (Total Minutes - Downtime Minutes) / Total Minutes x 100

3.2. Downtime is a period during which the affected surface is not accessible, as recorded by the public checks of the status page under Section 3.4. The checks run at one-minute intervals from several regions and measure at the edge of the content delivery network, which is what an end-user device experiences. An incident begins when a region fails three consecutive checks and ends when every region has recovered. A response served from the cache while the origin is unreachable is not Downtime. The exclusions listed in Section 4 shall not count as Downtime.

3.3. An incident the Provider declares on the status page counts as Downtime for its declared window, even where the checks passed. The Customer may submit verifiable measurement data from an independent monitoring system as counter-evidence. The Provider shall review such data within a reasonable period.

3.4. The current system status and the record of incidents are available at https://status.hotcodepush.com.

4. Exclusions from Downtime

4.1. The following periods shall not be taken into account when calculating Downtime:

- **Scheduled Maintenance**: Maintenance of the Control Plane announced at least 24 hours in advance. For major maintenance windows, the notice period is at least 48 hours. There is no maintenance window for Update Delivery.
- **Security Retraction**: Periods during which SDKs, APIs, or endpoints must be withdrawn, blocked, or patched without notice for security reasons (e.g. CVE-driven deprecations, abuse mitigation).
- **Force Majeure**: Events in accordance with the force majeure provision in GTC §16.
- **Targeted Security Incidents**: Targeted attacks on the HotCodePush infrastructure not attributable to the Provider (e.g. DDoS spikes against individual customer resources, zero-day exploits).
- **Third-Party Outages**: Outages at the Provider's upstream infrastructure providers or at internet service providers, unless attributable to the Provider.
- **Customer Fault**: Outages caused by faulty integration, configuration, or use by the Customer.
- **Customer-Side Waiting Time**: Periods during which the Provider is waiting for information, access credentials, or cooperation requested from the Customer.
- **Customer-Side Third-Party Integrations**: Outages of integrations configured by the Customer (e.g. webhook endpoints, the Customer's own identity provider).
- **Spending Cap**: Periods during which devices of the Customer receive no updates because the Spending Cap under GTC §6.3 is reached.
- **Abuse and Throttling**: Performance restrictions resulting from abuse, from exhausted rate limits, or from limitations under GTC §2.2 for use exceeding Fair Use.

5. Service Credits

5.1. Service Credits shall be granted exclusively to customers on a paid plan. No Service Credits shall be granted on the Free plan.

5.2. If the monthly availability of a surface falls below the Availability Target, the Customer shall be entitled to Service Credits in accordance with the following tiers:

Monthly Availability of the affected surface	Service Credit (% of Monthly Fee)
Below Availability Target to >= 99.0%	10%
Below 99.0% to >= 95.0%	25%
Below 95.0%	50%

5.3. The Monthly Fee is the usage charge for the affected calendar month under the Pay-as-you-go plan or the fee attributable to the affected calendar month under an Enterprise agreement. Service Credits are capped at **50% of the Monthly Fee** per calendar month. Where both surfaces fall below the Availability Target in the same month, the Service Credits do not stack beyond this cap.

5.4. Service Credits are the Customer's exclusive remedy for SLA breaches. This shall not apply to:

- statutory liability claims that cannot be limited (intent, gross negligence (Vorsatz und grobe Fahrlässigkeit), breach of material contractual obligations, the fulfilment of which is essential for the proper execution of the contract and on the observance of which the Customer may regularly rely (Kardinalpflichten), injury to life, body, or health, liability under the German Product Liability Act (Produkthaftungsgesetz), fraudulent concealment (Arglist), Art. 82 GDPR),
- termination rights expressly granted in this SLA or in the GTC.

6. Claim Process

6.1. Service Credits must be claimed within **30 days** after the end of the affected calendar month.

6.2. The claim shall be submitted in text form to support@hotcodepush.com and must contain the following information:

- Affected surface,
- Time and duration of the unavailability,
- Description of the impact.

6.3. The Provider shall review the claim and notify the Customer of the result within 15 business days (as defined in Section 3.2 of the Support Policy). Approved Service Credits under the Pay-as-you-go plan are paid out as a partial refund of the affected month's charge through the Merchant of Record. Approved Service Credits under an Enterprise agreement shall be offset against the next invoice.

7. Termination for Cause

7.1. A Customer under an Enterprise agreement shall have a right to terminate for cause if:

- the monthly availability of a surface falls below the Availability Target in **three consecutive calendar months**, or
- the monthly availability of a surface falls **below 95.0%** in any calendar month.

7.2. The termination right shall be exercised in text form within 30 days after the termination ground arises. The Customer shall receive a pro-rata refund for the unused period.

7.3. Customers on the Pay-as-you-go plan may cancel at any time in accordance with the GTC. They have no termination right under this Section.

8. Support

8.1. Response times and support channels are governed by the separate Support Policy. This SLA governs exclusively the availability of the infrastructure, not support response behaviour.

8.2. The authoritative language version of this SLA is the English version. The German version is for informational purposes only.