Flowpoint
E-invoicing & NRS compliance · Nigeria

SERVICE LEVEL AGREEMENT

Service Level Agreement (SLA)

Availability, support, incident-response, backup and recovery commitments for the Flowpoint e-invoicing service — written in plain language so that a technical or compliance reviewer can verify exactly what is promised and why.

Prepared for	**Nigeria Revenue Service (NRS) — SI/APP Accreditation**
Version	1.0 (Draft)
Date	June 2026
Classification	Confidential

1. Parties & scope

This Service Level Agreement ("SLA") is entered into between [INSERT: Provider legal entity name, e.g. Longconsole Consulting Limited, trading as "Flowpoint"] (the "Provider"), with registered office at [INSERT: Provider registered address] and RC number [INSERT: Provider RC number], and [INSERT: Customer / taxpayer organisation legal name] (the "Customer"). It takes effect on [INSERT: effective date] and remains in force for the duration of the underlying service agreement between the parties.

The SLA governs the availability, support, incident response, data handling and recovery commitments of the **Flowpoint** e-invoicing platform in its System Integrator (SI) capacity — the creation, signing and reporting of invoices to the Nigeria Revenue Service (NRS) with IRN issuance, verifiable QR/cryptographic stamp and a complete audit trail. It applies to the production environment and, where expressly stated, to the sandbox environment. It does not extend the Provider's responsibility to systems, networks or platforms operated by third parties, including the NRS / Merchant Buyer System ("MBS") platform, which is addressed in Section 10.

Why an e-invoicing SLA carries unusual weight. An invoice on this platform is not a convenience record — it is a tax-compliance document. When the platform "signs" an invoice, it obtains an Invoice Reference Number (IRN) and cryptographic stamp from NRS that legally certify the transaction. A taxpayer that cannot sign an invoice when it needs to may be unable to trade compliantly; a taxpayer whose audit record is lost may be unable to prove that it ever did. For that reason the commitments below treat *availability* and *auditability* — being up, and being able to prove what happened — as first-order obligations, not best-effort niceties.

> This SLA describes service commitments only. It does not alter the data-protection, security or compliance obligations set out in the wider service agreement or in **document 05 — Security & Compliance**; where any term conflicts, the signed master agreement prevails.

1.1 How to read this document (plain-language orientation)

This SLA is intended to be read by reviewers who are expert in security and operations but who are not specialists in tax or accounting. Each technical and commercial term is defined in Section 2 and then explained again, in plain English, at the point where it is used. As a quick orientation:

- **Availability** — how often the service is up and reachable, expressed as a percentage of the month (Section 3).
- **Downtime budget** — the small amount of "allowed" outage that follows mathematically from the availability target; everything beyond it is a breach (Section 3).
- **Severity (S1–S4) and Response/Resolution times** — how fast the Provider promises to react, graded by how badly an incident hurts the Customer's ability to stay compliant (Section 4).
- **RPO and RTO** — how much recent data could be lost, and how long restoration could take, if a disaster struck (Section 7).
- **Maintenance window and change management** — when planned work happens and how changes are controlled so they do not become outages (Section 6).
- **Service credits** — the financial remedy if availability targets are missed (Section 9).

2. Definitions

The terms below are the precise, contractual meanings used throughout this SLA. Each is written so that a reader without a finance or accounting background can follow it; the sections that use them add further plain-language explanation at the point of use.

Term	Meaning
Availability	The percentage of time in a calendar month during which the platform's core service functions — authentication, invoice creation, and submission of an invoice to NRS for signing — are operational and reachable by the Customer, calculated as defined in Section 3.
Downtime	Any period during which the core service functions are unavailable to the Customer for reasons within the Provider's control, excluding the periods listed in Section 3.3.
Incident	An unplanned interruption to, or reduction in the quality of, the service; or an event that has not yet impacted the service but is likely to do so if unaddressed.
Severity	The classification (S1–S4) assigned to an Incident based on its operational and compliance impact, as defined in the table in Section 4.
Response Time	The elapsed time between the Provider's receipt of a correctly reported Incident through a designated channel and the Provider's first substantive acknowledgement, which includes an Incident reference and a named owner.
Resolution Time	The elapsed time between receipt of an Incident and either a permanent fix or an accepted workaround that restores service to within normal operating parameters. "Resolution" of an Incident whose root cause is external (e.g. NRS-side) means restoration of the Provider's portion of the service, not the third party's.
Maintenance Window	A pre-announced period reserved for planned maintenance, upgrades or change deployment, during which the service may be wholly or partly unavailable. Time spent within an announced Maintenance Window is not counted as Downtime.
Business Hours	**[INSERT: business hours, e.g. 08:00–18:00 WAT, Monday–Friday, excluding Nigerian public holidays]** . Hours outside this window are "Outside Business Hours".

3. Service availability commitment

What "availability" means in plain terms. Availability is simply the share of time the service is up and doing its core job — letting a user log in, build an invoice, and submit it to NRS for signing. It is expressed as a percentage of the calendar month. A target such as "99.5%" looks close to perfect, but the gap between 99.5% and 100% is the part that matters: it is the *downtime budget*, the amount of outage the Provider is permitted before the commitment is breached. Stating the budget in minutes — rather than hiding it behind a percentage — is how a serious operator makes the promise auditable.

3.1 Target

The Provider commits to a monthly Availability target of **99.5%** for the core service functions, measured over each calendar month. At 99.5%, the permitted Downtime is approximately 3 hours 39 minutes per 30-day month — that figure *is* the monthly downtime budget. In practical terms it means that, across an entire month of continuous operation, total outage attributable to the Provider must stay under roughly three and a half hours for the target to be met. A higher target may be agreed in writing for an individual Customer: **[INSERT: agreed uptime target if different from 99.5%]**.

3.2 Measurement method

A commitment is only as honest as the way it is measured. Availability here is calculated from independent, externally observed monitoring of the live service endpoints — not from the Provider's own internal opinion of its health — and a transient network blip is deliberately not counted as an outage:

Formula	Definition
Availability % = (Total minutes in month − Downtime minutes) ÷ Total minutes in month × 100	Downtime is measured from the Provider's external monitoring of the core service endpoints, sampled at no greater than one-minute intervals. An endpoint is considered down only after two consecutive failed checks, to avoid counting transient network blips. Monitoring records and the public status page are the agreed sources of truth for an Availability calculation.

3.3 Exclusions

No availability target can hold the Provider responsible for outages it cannot control. The exclusions below define that boundary honestly: they carve out time lost to planned maintenance the Customer was warned about, to the Customer's own systems, to genuine force majeure, and — importantly for this platform — to the NRS / MBS side of the connection, which the Provider calls but does not operate. The following periods do not count as Downtime and are excluded from the Availability calculation:

- Planned maintenance carried out within an announced Maintenance Window (Section 6).
- Unavailability, degradation, throttling, rejection or non-response of the **NRS / MBS platform** or any NRS-side authorization, registration or rate-limiting condition (see Section 10) — including the intermittent NRS-side `HTTP 403 "access is denied"` condition under resolution with NRS.
- Outages caused by the Customer's own systems, configuration, credentials, network or misuse of the service.
- Force majeure: events beyond the Provider's reasonable control, including utility or carrier failures, natural disasters, civil disruption, or acts of government.
- Suspension of service required by law, by NRS, or to prevent a security incident from causing wider harm.

4. Incident severity & response

What "severity" and the response/resolution targets mean. When something breaks, not every fault is equally urgent — a total signing outage is a crisis, a typo in a report is not. *Severity* is the label (S1 = most severe, down to S4 = cosmetic) that decides how hard and how fast the Provider must react. Two clocks then start: the *Response Time* is how quickly a human acknowledges the incident, assigns it a reference number and names an owner (so the Customer knows it is being worked, not sitting in a queue); the *Resolution Time* is how quickly normal service is restored, either by a permanent fix or an accepted workaround. Grading these targets by compliance impact is what separates a world-class operator from a best-effort one — the most severe band is reserved for anything that stops a taxpayer signing invoices or that threatens the audit record those invoices depend on.

Every Incident is assigned a severity on first contact and may be re-classified as understanding improves. The first substantive acknowledgement always includes an Incident reference and a named owner. Targets below are measured from receipt of a correctly reported Incident; targets outside Business Hours apply to the 24/7 on-call rotation for S1 and S2 only.

Severity	Definition	Target response	Target resolution / workaround
S1 — Critical	Total service outage, or NRS invoice signing/reporting is unavailable, such that no Customer can sign or report invoices to NRS; or confirmed loss of, or inability to write, the mandatory submission audit log. Material compliance and revenue impact.	≤ 30 minutes, 24/7	≤ 4 hours to a fix or accepted workaround; continuous effort until resolved.
S2 — High	Severe degradation: a core function is materially impaired (e.g. intermittent signing failures attributable to the Provider, authentication failures, or backlog growth in submission processing) but a partial path remains, or the impact is limited to a subset of tenants.	≤ 1 hour, 24/7	≤ 8 business hours to a fix or accepted workaround.
S3 — Medium	A non-core feature is unavailable or behaving incorrectly (e.g. a report, an ERP import job, or the portal UI in a limited area) with a reasonable workaround available; no impact on the ability to sign and report invoices.	≤ 1 business day	≤ 5 business days, or scheduled into the next maintenance release.
S4 — Low	Cosmetic issues, documentation errors, minor UI inconsistencies, or feature requests with no operational impact.	≤ 2 business days	Scheduled into a future release at the Provider's reasonable discretion.

> **Audit-log integrity is treated as S1.** Because every NRS round-trip must write a complete row to the submission audit log, a confirmed failure of that audit write — even if signing itself appears to succeed — is classified as Critical, since it threatens the compliance record on which certified receivables depend.

5. Support

Support tiers, in plain terms. "Support" here means the people and channels a Customer reaches when something is wrong. The service is offered in tiers: routine questions and lower-severity issues are handled during published support hours, while the most damaging incidents (S1 and S2) are covered by a round-the-clock on-call rotation so that a critical signing outage is never left waiting for office hours to begin. The channels below tell the Customer exactly how to reach each tier, and the status page gives an at-a-glance, self-service view of whether an outage is already known and being worked.

5.1 Hours & channels

Item	Detail
Standard support hours	**[INSERT: support hours, e.g. 08:00–20:00 WAT, Monday–Friday]** for all severities.
On-call coverage	24/7 on-call rotation for **S1 and S2** Incidents only.
Email	**[INSERT: support email, e.g. support@…]** — primary channel for S2–S4 and for raising any Incident.
Support portal / in-app	**[INSERT: portal or in-app support URL]** — ticket creation, status tracking and history.
Urgent / S1 hotline	**[INSERT: emergency contact number for S1/S2 outside business hours]**.
Status page	**[INSERT: public status page URL]** — live service status and incident notices.

5.2 Escalation path

What "escalation" means. Escalation is the defined route by which an unresolved or worsening incident is handed up to progressively more senior people, each with more authority to mobilise resources. A mature escalation path does two things: it gives the Customer a clear lever to pull when targets are slipping, and — crucially — it commits the Provider to escalate *itself*, without being asked, for the most serious incidents. If an Incident is not progressing within the targets above, the Customer may escalate through the following tiers. The Provider will also escalate proactively for S1 and prolonged S2 Incidents.

Tier	Owner	Engaged when	Contact
Tier 1	Support engineer / on-call responder	On receipt of any Incident.	**[INSERT: Tier 1 contact]**
Tier 2	Engineering lead / platform owner	S1 immediately; S2 if Response Time is missed; or on Customer request.	**[INSERT: Tier 2 contact]**
Tier 3	Head of Engineering / Service Delivery Manager	S1 not resolved within target, or repeated breaches.	**[INSERT: Tier 3 contact]**

Executive	Accountable executive sponsor	Major outage, regulatory exposure, or unresolved escalation.	**[INSERT: executive sponsor contact]**

6. Maintenance & change management

Why planned maintenance and change control matter. Most serious outages in any platform are not caused by attackers or hardware — they are self-inflicted, introduced by a change that went wrong. "Change management" is the discipline of making changes deliberately rather than accidentally: scheduling disruptive work into pre-announced *maintenance windows* (so the Customer is never surprised), giving advance notice, version-controlling every change, and testing data migrations before they touch live records. A *maintenance window* is simply a published time slot reserved for upgrades, during which a short, expected interruption may occur; because it is announced in advance, time spent inside it is not counted against the availability budget. The Provider performs routine maintenance — patching, dependency upgrades, schema migrations and capacity changes — within scheduled Maintenance Windows to minimise disruption.

Item	Commitment
Standard maintenance window	**[INSERT: e.g. 01:00–05:00 WAT on weekends]**, scheduled to avoid peak invoicing periods.
Notice for planned maintenance	At least **72 hours'** advance notice by email and on the status page for any window expected to cause service interruption.
Maximum planned downtime	**[INSERT: e.g. ≤ 4 hours per calendar month]**; time within an announced window is excluded from Downtime.
Emergency change	Where a critical security patch or stability fix cannot wait, the Provider may deploy with shortened or immediate notice, recording the reason and notifying affected Customers as soon as practicable.
Change records	Material changes are version-controlled and recorded; migrations affecting compliance data follow the Provider's idempotent, per-record-transactional, dry-run-tested migration process.

> Database migrations that touch compliance tables (submission audit log, invoices, receivables) are never destructive by default and are tested in a dry-run mode before production deployment, consistent with the Provider's data-safety controls in **document 05**.

7. Data backup, retention & recovery

RPO and RTO explained without jargon. These are the two numbers that describe how well a service survives a disaster. *RPO — Recovery Point Objective* answers "how much recent data could we lose?"; it is measured *backwards* from the moment of failure. An RPO of fifteen minutes means that, in the worst case, you might lose the last fifteen minutes of work but nothing older, because backups are captured at least that often. *RTO — Recovery Time Objective* answers "how long until we are back up?"; it is measured *forwards* from the moment a disaster is declared. An RTO of four hours means service should be restored within four hours of the failure. For an e-invoicing platform these two numbers are compliance-critical: a large RPO could mean signed invoices and their audit records simply vanish, and a large RTO could mean a taxpayer is unable to invoice for the duration. A world-class operator therefore not only sets tight targets but periodically *tests its restores* — proving the backups can actually be brought back, rather than merely assuming they exist.

Item	Commitment
Backup frequency	**[INSERT: e.g. automated daily full backup + continuous transaction-log/point-in-time capture]**, with backups stored encrypted at rest.
Recovery Point Objective (RPO)	**[INSERT: target RPO, e.g. ≤ 15 minutes]** — the maximum acceptable data loss measured backwards from a disruptive event.
Recovery Time Objective (RTO)	**[INSERT: target RTO, e.g. ≤ 4 hours]** — the target time to restore core service after a declared disaster.
Backup verification	Restores are tested periodically (**[INSERT: test cadence, e.g. quarterly]**) to confirm backups are usable, not merely present.
Submission audit-log retention	NRS submission audit records (request, response, HTTP status, IRN, QR data, outcome, duration and attempt number) are retained for **7 years** to meet tax record-keeping requirements, in line with the immutable audit-trail commitment.
Data export	The Customer may export its invoice data (e.g. JSON, CSV, UBL XML) through the portal or API at any time. On termination, data is retained in a read-only, exportable form for the statutory window: **[INSERT: post-termination export availability, e.g. at least 60 days]**.

8. Security incident & breach notification

What this section commits to. Separately from ordinary outages, this section governs what happens if data is exposed or tampered with, or if the keys and credentials used to sign invoices are compromised. The core promise is timely, honest disclosure: the Customer is told what happened, what data was affected and what is being done, quickly enough to meet its own legal obligations. This matters acutely here because the signing keys are what make an invoice legally trustworthy — a compromise of that material is not just a data issue, it strikes at the integrity of the tax record itself, which is why it is treated at the highest severity.

A "Security Incident" is any actual or reasonably suspected unauthorised access to, disclosure of, alteration of, or loss of Customer or tenant data, or any compromise of the cryptographic material or credentials used to sign invoices.

Item	Commitment
Initial notification	The Provider will notify affected Customers without undue delay and in any event within **[INSERT: notification window, e.g. 72 hours]** of confirming a Security Incident that is likely to affect the Customer's data.
Content of notice	Nature of the incident, categories and approximate volume of data affected, likely consequences, measures taken or proposed, and a point of contact.
Regulatory alignment	Notification and handling are designed to align with the Nigeria Data Protection Act / NDPR principles and with any reporting obligation to the relevant data-protection authority. **[INSERT: DPO name & contact if appointed]**.
Containment & remediation	The Provider will take prompt steps to contain the incident, preserve evidence, and provide a post-incident review on request.
Cooperation	The Provider will cooperate with the Customer's own breach-notification obligations to NRS and to data subjects where applicable.

> A Security Incident affecting signing keys, authentication, or the integrity of the submission audit log is treated as an **S1** Incident under Section 4 and escalated immediately.

9. Service credits & remedies

What a "service credit" is. A service credit is the financial consequence the Provider accepts when it misses an availability target — typically a percentage of the monthly fee, refunded or applied against a future invoice. It puts the Provider's own revenue at stake for the reliability it promises, which is precisely why credits are a hallmark of a serious SLA rather than a marketing one. The credit is the Customer's defined remedy for missed availability; it does not cap the Provider's other obligations under the master agreement. If the monthly Availability target in Section 3 is not met for reasons within the Provider's control, the Customer's sole and exclusive remedy is a service credit calculated against the monthly fee for the affected service, requested within 30 days of the end of the affected month and applied to a subsequent invoice. The placeholder schedule below is to be confirmed in the commercial agreement.

Monthly Availability achieved	Service credit (% of monthly fee)
Below 99.5% but ≥ 99.0%	[INSERT: e.g. 10%]
Below 99.0% but ≥ 95.0%	[INSERT: e.g. 25%]
Below 95.0%	[INSERT: e.g. 50%]

Service credits do not apply to Downtime arising from any exclusion in Section 3.3 (including NRS/MBS-side conditions and force majeure). Credits are capped at [INSERT: cap, e.g. 50% of the monthly fee] for any single month and do not constitute a penalty.

10. Exclusions & dependencies

The commitments in this SLA apply only to components operated and controlled by the Provider. The following are explicitly outside the scope of the Availability and Resolution targets:

- **Availability of the NRS / MBS platform and NRS-side authorization.** Flowpoint signs and reports invoices by calling the NRS platform. If NRS is unavailable, slow, rate-limited, or rejects or fails to authorize a request — including the NRS-side `HTTP 403 "access is denied"` condition currently under resolution with NRS — the resulting impact is an external dependency outside the Provider's control and is excluded from Downtime. The Provider remains responsible for correctly forming, transmitting, retrying (per its idempotency rules) and auditing each request, and for accurately reflecting the NRS outcome.
- The Customer's own systems, ERP/accounting software, network connectivity, browsers, and credentials.
- Third-party infrastructure and subprocessors (e.g. managed database/cache, object storage, email/SMS delivery) where the failure originates with that provider; the Provider will nonetheless pursue mitigation and failover where designed.
- Capability explicitly disclosed as not yet built: **APP transmission to the buyer's Access Point (NRS Step 6)** and **inbound buyer acknowledgements** are roadmap items and are not covered by any availability commitment in this SLA. For B2B/B2G, an invoice reaches `irn_received` after signing; `transmitted_at` is set only for B2C.

> **Honest dependency statement.** A successful NRS round-trip requires NRS-side authorization that the Provider does not control. Production go-live is pending resolution of the sandbox authorization (403) matter with NRS. Until then, availability of end-to-end signing depends on NRS, and this SLA does not warrant the behaviour of the NRS platform itself.

11. Standards & operational excellence

The commitments in this SLA are credible because they sit on top of engineering practices that are already built into the platform, not aspirations bolted on for accreditation. The list below describes the controls that make the availability, recovery and integrity promises achievable, in plain terms a security reviewer can verify against the architecture in **document 05 — Security & Compliance**.

Practice	What it is, in plain language	Why it underpins this SLA
Immutable audit trail	Every round-trip to NRS — the request sent, the response received, the HTTP status, the IRN and QR data, the outcome, the duration and the attempt number — is written to an append-only submission log (`fp_submission_log`) that is not edited after the fact.	It is the source of truth behind the Resolution and recovery commitments: an outage can be reconstructed minute-by-minute, and a "certified" invoice can always be proven against a real, recorded NRS success.
Fail-closed audit write	If that audit record cannot be written, the operation deliberately fails rather than proceeding silently — so the system never records an invoice as signed without the matching evidence.	This is why audit-log failure is an S1 (Section 4): the platform is engineered to refuse to create a "phantom" certified invoice, protecting the compliance record the Customer relies on.
Idempotent retries	Retried submissions carry an idempotency key, so re-sending a request after a timeout cannot accidentally sign or report the same invoice twice.	It lets the Provider safely retry across transient NRS or network faults (Section 10) without creating duplicate tax documents — reliability without double-counting.
Fail-closed configuration	In production (`NODE_ENV=production`) the service guards its required secrets and refuses to start if any are missing, rather than booting in an insecure or half-configured state.	It prevents a misconfigured deployment from silently serving traffic — a common, avoidable cause of both outages and security incidents.
Encryption in transit and at rest	Data is protected with TLS while moving over the network, and cryptographic key material is encrypted at rest using AES-256-GCM.	It backs the breach-notification and data-protection commitments in Section 8 and protects the very keys that make a signed invoice trustworthy.
Monitored, containerised NRS connector	The platform runs in containers (Docker), is deployable to AWS (Elastic Beanstalk / ECS Fargate) or DigitalOcean, uses a MySQL 8 datastore, and offloads genuinely large bulk work to Redis + BullMQ worker queues; the NRS connector and core endpoints are externally monitored.	Container portability and queued bulk processing support the recovery (Section 7) and availability (Section 3) targets; external monitoring is the agreed, independent basis for measuring availability.

Recognised standards & controls	The platform implements UBL 2.1, ISO 20022 and NRS schema v1.1 for document formats; and JWT + OTP authentication, RBAC, per-tenant isolation, bcrypt password hashing, OWASP-aligned controls and NDPR alignment for security and privacy.	Conformance to published standards is what lets a reviewer check the platform against an external yardstick rather than the Provider's own word.

Standards & operational excellence. The "world-class" claim in this SLA rests on controls that are genuinely in place today — an immutable, append-only NRS audit trail; idempotency keys that make retries safe; fail-closed configuration and a fail-closed audit write that refuse to proceed without their evidence; TLS in transit with AES-256-GCM-protected key material at rest; and externally monitored core endpoints. These are the mechanisms that turn the promises in Sections 3–9 from assertions into things a reviewer can verify.

Honest status — not yet complete. In the interest of full disclosure, the following are Partial / roadmap and are *not* warranted by this SLA: **APP transmission to the buyer's Access Point (NRS Step 6)** and **inbound buyer acknowledgements** are Not yet (roadmap); **ISO 27001 certification** is on the roadmap; an **independent penetration test (VAPT)** is being commissioned; and **production go-live** remains pending resolution of an NRS-side `HTTP 403` authorization matter (Section 10). None of these is presented as already achieved.

12. Review & governance

Item	Commitment
Service review cadence	The parties review SLA performance [INSERT: cadence, e.g. quarterly], covering Availability achieved, Incidents raised, response/resolution performance and any service credits.
Reporting	The Provider makes Availability and Incident summaries available on request and via the status page; major-incident post-mortems are shared for S1 events.
SLA amendment	The Provider may update this SLA on [INSERT: notice period, e.g. 30 days'] written notice. Changes do not reduce a commitment already in force for the current period.
Continuous improvement	Recurring Incidents and missed targets feed corrective and preventive actions tracked through the Provider's change process.
Governance contacts	Provider Service Delivery Manager: [INSERT: name & contact]. Customer relationship owner: [INSERT: name & contact].

Acceptance

For the Provider	For the Customer
Name: [INSERT: signatory name] Title: [INSERT: title] Entity: [INSERT: Provider legal name] Date: [INSERT: date] Signature: _____	Name: [INSERT: signatory name] Title: [INSERT: title] Entity: [INSERT: Customer legal name] Date: [INSERT: date] Signature: _____