PRIVATE CONTRACT

16-Week Delivery Plan with Weekly Acceptance Criteria

ProficientHub Platform — Commercial Launch Sprint

Client:	Santaliestra Limited (ProficientHub)
Represented by:	Jesús Santaliestra Grau, CEO
Contractor:	Brian Musonza
Date:	March 2026
Classification:	**CONFIDENTIAL**

1. Purpose of This Agreement

This Private Contract establishes the binding weekly deliverable schedule, acceptance criteria, and payment terms for the 16-week ProficientHub platform commercial launch sprint. It supplements and forms part of the engagement agreement between the parties and any terms agreed via the Toptal platform.

The weekly acceptance criteria listed herein consolidate all commitments made during the negotiation phase (10–17 March 2026), including items from the formal proposal, the QA response document, and all written communications between the parties.

Both parties agree that weekly payment is conditional upon verified completion of the acceptance criteria specified for the corresponding week, as reviewed in the Thursday weekly review meeting.

2. Commercial Terms

2.1 Agreed Rate and Payment Method

Both parties agree, by mutual consent, to a rate of **$120 USD per hour** on a **40-hour-per-week** billing basis, for a total engagement of **16 weeks**.

Item	Amount
Agreed rate	$120 USD / hour
Weekly billing basis	40 hours / week
Weekly payment amount	$4,800 USD / week
Engagement duration	16 weeks
TOTAL ENGAGEMENT	**$76,800 USD**

Payment method: Weekly payments of $4,800 USD shall be processed through the Toptal platform each week, subject to verified completion of the corresponding week's acceptance criteria as confirmed in the Thursday review meeting. If any acceptance criteria remain incomplete at the Thursday review, both parties shall agree on a resolution path; payment for that week may be adjusted or deferred until the outstanding items are resolved.

Working hours: Brian has committed to working materially above 40 hours per week (often 60+ hours including weekends) at no additional charge to meet the launch deadline. Hubstaff tracking may not fully reflect actual hours worked. The binding control mechanism for payment is the verified completion of weekly acceptance criteria defined in this contract, not time tracking.

2.2 Early Delivery Bonus

Santaliestra Limited commits to paying Brian Musonza an **early delivery bonus** in the event that the complete platform launch (as defined by all Week 16 acceptance criteria being fully met) is achieved before the end of Week 16.

Bonus calculation: $120 USD × 40 hours × number of full weeks by which delivery is advanced. This represents one full week's billing value per week saved. For example: if all launch criteria are verified at the end of Week 14 (2 weeks early), the bonus is $120 × 40 × 2 = $9,600 USD. If delivery is completed 3 weeks early, the bonus is $120 × 40 × 3 = $14,400 USD.

Payment method for bonus: The early delivery bonus shall be paid by Santaliestra Limited via bank transfer within 15 business days of Brian presenting a corresponding invoice. This payment is separate from and additional to the weekly Toptal payments.

2.3 Designer

The visual designer for the project shall be hired and managed directly by Santaliestra Limited. The designer is not part of Brian's scope of work or budget. Brian shall collaborate with the Client's designer from Week 1 and provide technical direction and design requirements as needed.

2.4 Scope Control

The scope defined in this contract is locked from Week 4 onwards. Any new major items requested after Week 4 will be evaluated jointly and may affect timeline or cost. Minor refinements and clarifications within the agreed scope do not constitute scope changes.

3. Delivery Structure

The build is sequenced in dependency order: foundations and environments; superadmin controls; institution setup and white-label activation; teacher workspace; student entry and dashboard; science and learning loop; commerce and usage loop; AI agents; retention, analytics, and launch hardening.

Throughout the build, ProficientHub will maintain an internal test institution for end-to-end validation. Commercial flows will be validated first in sandbox mode, then through controlled real-transaction checks.

Weekly reviews take place every Thursday. Brian presents working product against the acceptance criteria. Both parties sign off or document exceptions.

3.1 Team Structure

- Brian Musonza: Technical lead, architect, AI/ML engineer, infrastructure, delivery management, design direction
- Visual designer: Hired and managed by Santaliestra Limited, starting Week 1
- ProficientHub teachers (3): Exam format/UI/UX fidelity validation, feedback quality assurance, agent QA, golden-standard creation, Genome calibration, and any other function Brian requires

3.2 Teacher Review App

Brian shall build a standalone internal Teacher Review App to systematise teacher contributions. The app presents cards showing system or student input on the left and the teacher's expected output on the right. One teacher creates or edits the expected output; a second teacher verifies it. Approved records become the ground-truth dataset used for system evaluation, regression testing, and agent calibration.

Content types covered: study plans, exams, rubrics, flashcards, lesson materials, targeted exercises, agent responses, scoring outputs, and any additional areas identified during the engagement.

3.3 B2B Content Model and Teacher Role

ProficientHub is a B2B platform serving language academies, hospitals, universities, recruiters, and companies. Many client institutions already have their own training materials. The platform must support:

- **Academy content upload:** Institutions upload their own exams, study guides, exercises, and lesson materials. The platform ingests, indexes, and serves them within the branded environment.
- **New material from existing:** The system generates new materials (exercises, flashcards, study guides, practice tests) derived from academy-uploaded content, using AI generation pipelines.
- **ProficientHub native content:** AI-generated exam content and study materials as the default layer for institutions without their own content.
- **Blended mode:** Any combination of academy content and native content.

Primary teacher responsibilities in this engagement:

- **Exam format fidelity:** Ensuring every supported exam type replicates the official format, UI/UX, question types, timing, and scoring. The student must feel they are taking the real exam.

- **Feedback quality:** Validating that AI-generated feedback (scoring, explanations, recommendations) matches official exam standards.

- **Agent quality assurance:** Ensuring AI Agents (Tutor, Mock Exam Coach, Study Planner) provide premium support: accurate guidance, appropriate difficulty, contextually relevant advice, pedagogically sound recommendations.

- **Golden-standard creation:** Providing verified input-output pairs and grading rubrics. One teacher creates/edits, a second verifies. Approved records become ground-truth data for system evaluation.

- **Genome and science tools:** Supporting calibration of the Proficiency Genome, Exam Readiness Score, and any other scientific or pedagogical function where teacher expertise is required.

- **General availability:** Teachers are available to support Brian in any capacity needed throughout the engagement, including but not limited to the areas listed above.

4. Weekly Acceptance Criteria

Each week defines focus, measurable acceptance criteria, and unlocks. All items must be demonstrable in the Thursday review. Weekly payment of $4,800 USD through Toptal is released upon verified completion.

MILESTONE 1 — Platform Reset & Foundation (Weeks 1–4)

WEEK 1	
Focus	Platform foundations, environments, config safety, shared application shells.
Acceptance Criteria	1. Working staging baseline deployed 2. CI/CD pipeline with environment separation (staging + production) operational 3. Config validation and fail-fast startup behaviour in place 4. Structured logs, trace IDs, and health endpoints active 5. Shared web application shell running 6. Shared mobile application shell running (Expo — iOS, Android, web unified) 7. Public auth shell ready for role-based work 8. Infrastructure defined as code (IaC skeleton) 9. **Feature flags:** Global feature-flag system operational at superadmin level 10. **Product event collection:** Baseline event capture instrumented from Day 1 11. **Design collaboration:** Brian and Client's designer aligned on UX direction for key flows 12. **AI Core architecture:** AI Core designed as an independent service from the SaaS layer, communicating via internal API. Both can evolve separately
Unlocks	*Role-specific access and admin flows without reworking infrastructure.*

WEEK 2	
Focus	Superadmin as the root control layer.
Acceptance Criteria	1. Superadmin login and protected root access working 2. Superadmin dashboard shell functional 3. Global exam catalogue / admin view in place 4. Global feature-flag controls operational from superadmin 5. AI provider configuration surfaces available 6. Pricing configuration surfaces available 7. Demo-management and institution-provisioning scaffolding in place 8. **Multi-region:** IaC supports multi-region deployment via configuration
Unlocks	*The root layer that institution setup depends on.*

WEEK 3

Focus	Institution provisioning and white-label onboarding.
Acceptance Criteria	1. Superadmin-to-institution provisioning flow complete 2. Institution admin activation and login working 3. **Onboarding wizard:** Guided setup for academy: branding, exam config, billing config, agent config, student/teacher import 4. White-label setup: name, colours, fonts, logo fully configurable 5. Exam activation and feature activation at institution level 6. Institution dashboard shell functional 7. Brand/domain preview for the academy portal working 8. **Container App model:** Single app binary that rebrands with academy identity on student login
Unlocks	*A real academy context for later teacher and student flows.*

WEEK 4	
Focus	Institution operations, billing foundations, communication providers.
Acceptance Criteria	1. Team invite and role assignment flow working 2. Teacher invite and student invite/import flow working 3. Institution package and token configuration operational 4. **Payment gateway:** Provider-abstracted billing layer. Stripe default. Razorpay, Paystack, regional gateways via thin adapters 5. **Multi-currency:** Auto-conversion logic in billing core 6. **Communication providers:** Adapter layer for email/SMS/WhatsApp configurable per academy. Transactional email operational 7. Branded student portal bootstrap running 8. Institution operational settings complete 9. **SSO foundations:** OAuth2/SAML setup surfaces available 10. **Scope lock:** Launch scope formally locked as of end of Week 4
Unlocks	*An institution that can operate. Billing and comms foundations in place.*

MILESTONE 2 — Core Engines & Workflows (Weeks 5–8)

WEEK 5	
Focus	Teacher workspace, academy content management.
Acceptance Criteria	1. Teacher login functional 2. Teacher dashboard shell operational 3. Assigned-student view working 4. Review queue foundation in place 5. **Academy content upload:** Academies upload their own materials (exams, study guides, exercises). Upload, review, publish flow operational 6. **Material generation from existing:** System generates new materials (exercises, flashcards, study guides) derived from academy-uploaded content 7. Teacher messaging base operational 8. **Teacher Review App:** Standalone internal QA tool operational with card-based create/verify workflow for golden-standard data 9. **Class scheduling:** Calendar with Zoom/Meet/Teams integration for live classes. Per-exam and per-group scheduling
Unlocks	*Teacher-led review, content management, and academy content workflows.*

WEEK 6	
Focus	Student entry, first-run journey, LLM Gateway.
Acceptance Criteria	1. Student login through branded academy portal working 2. Forced password change flow operational 3. Student dashboard functional 4. Student counters and assigned exam view in place 5. Materials and history/schedule summary cards displayed 6. **Student calendar:** Study calendar with classes, Zoom/Meet/Teams links, AI reminders, and study plan integration 7. Placement entry flow started 8. Mobile login and mobile home/dashboard with full feature parity 9. **LLM Gateway:** Central routing with model selection by task, cost tracking per tenant with real-time dashboards, rate limiting, failover, semantic cache via Redis (target 20-30% API call savings), prompt registry with versioning and A/B testing 10. **PII anonymisation:** Automated PII stripping before LLM API calls inside Gateway (GDPR/HIPAA)
Unlocks	*First real end-user journey. LLM Gateway routing live.*

WEEK 7	
Focus	Placement, Genome, ERS, study-plan loop, AI validation.
Acceptance Criteria	1. Completed placement flow end-to-end 2. Placement result screen displayed 3. **Proficiency Genome:** 4-level ontology (Skill > Sub-skill > Attribute > Exam-mapping). Full 300–500 attribute registry seeded and

	versioned. Approximately 250–300 attributes operationally wired into scoring, diagnostics, ERS, sequencing, and prescriptions. Each student sees the 30–50 attributes most relevant to their exam and target band
	4. **L1 Transfer Intelligence:** L1 tracking data model in Genome. 80 L1 languages × 9 exams structure in place
	5. **Exam Readiness Score:** First ERS calculation functional
	6. First generated study plan displayed to student
	7. Updated student dashboard with readiness and next-step state
	8. **AI validation layer:** Automated checks for structure, length, type correctness. Anti-hallucination system (fact-checking, no false data/invented references). Anti-repetition system across exam catalogue (100 exams of same type must not repeat topics beyond official exam patterns). Quality gate: auto-publish, teacher review queue, or rejection with regeneration
Unlocks	*Scientific core and personalised study state.*

WEEK 8	
Focus	Daily practice, mock exam flows, scoring, exam governance.
Acceptance Criteria	1. Daily practice loop functional
	2. Smart review and targeted practice flow (spaced repetition by subskill)
	3. Metacognition/reflection capture in student flow
	4. Student exam dashboard operational
	5. **Exam simulator:** Official-feeling exam flows for priority exams. UI/UX identical to official format per exam type
	6. Result/report scaffolding and review hooks in place
	7. **Official feedback PDF:** Downloadable PDF report per exam mimicking the official exam's score report format. Criterion-by-criterion scoring
	8. Scoring and feedback connected to Genome
	9. Ontology schema locked, Genome/ERS foundations solid
	10. **Exam content governance:** Version control, validation, guardrails, review workflow, model-routing discipline, production monitoring
	11. **Exercises from Genome:** System detects student gaps, generates targeted exercises in student's exam format. Results feed back into Genome
	12. **Partial exam consumption:** A mock exam can be completed by sections across multiple sessions. Pending sections remain available without additional consumption. Fully consumed only when all sections completed
Unlocks	*Real exam-preparation behaviour across daily learning and mock execution.*

MILESTONE 3 — Product Differentiation & Mobile (Weeks 9–12)

WEEK 9	
Focus	Credits, usage, upsells, purchase loop, billing validation.
Acceptance Criteria	1. Credit and token usage logic live in student flows 2. No-credit redirect behaviour working 3. **Student upsell marketplace:** Academy configures and sells any product/service (ProficientHub + custom: courses, 1:1 classes, books, group sessions). Fully configurable pricing 4. Purchase flow end-to-end operational 5. Payment success and entitlement update flow working 6. Return-to-action loop after purchase functional 7. Controlled internal-institution billing validation established 8. **Billing 4-dimension model:** Exam selection, mock volume, AI tutor toggle, AI credits pack with volume discounts and multi-gateway support 9. **Independent quotas:** Writing Tests and Speaking Tests have independent counters separate from mock exam credits. Purchasing Writing Tests does not consume mock exam quota and vice versa 10. **AI Credit dual pool:** Institutional pool (academy's total inventory purchased from ProficientHub) + per-student pool (allocated from institutional pool or purchased as upgrades). Consumption tracked at both levels. Student upgrade purchases deducted from institutional volume
Unlocks	*Commercial operating loop for institutions and students.*

WEEK 10	
Focus	AI Agents: Tutor, Mock Exam Coach, Study Planner.
Acceptance Criteria	1. AI Tutor live, connected to Proficiency Genome, context-aware 2. Mock Exam Coach live, context-aware 3. Study Planner revision flow working 4. Agent session continuity maintained 5. Usage and credit deduction tied to agent actions 6. Low-credit disable state operational 7. **Agent context harness:** Efficient retrieval of relevant student data per interaction 8. **Placement Test as sales funnel:** Full placement-to-study-plan flow as main commercial entry for academies
Unlocks	*Premium AI layer inside the real product loop.*

WEEK 11	
Focus	Teacher intelligence, communications hub, didactic materials.
Acceptance Criteria	1. **Teacher Intelligence Dashboard:** At-risk students, weakness groups, intervention suggestions 2. Automated notifications and communications across channels

	3. **Communications Hub:** Transactional email, in-app notifications, exam results. Configurable providers per academy
	4. Teacher communication rules and delivery-state visibility
	5. **Didactic material generation:** PDF study guides, flashcard decks, audio (ElevenLabs TTS) per subskill from Genome
	6. **Branded material workflow:** Teacher approval for generated materials, including materials from academy uploads
Unlocks	*Teacher leverage and institution-level visibility at scale.*

WEEK 12

Focus	Enterprise readiness, CRM/ERP, ROI, demo, exam coverage.
Acceptance Criteria	1. **SSO:** OAuth2 and SAML complete
	2. **External CRM/ERP:** Bidirectional API connector layer. Sync leads, students, invoices, payments, progress with HubSpot, Salesforce, Xero, etc.
	3. Integration and mapping/admin controls operational
	4. **ROI Calculator:** For institutions and ProficientHub internal use
	5. **Demo/Sandbox:** Internal test institution as demo with pre-populated realistic data (students with progress, completed exams, dashboards with metrics, active AI agents). Functional demo, not mock-ups
	6. Superadmin sales/analytics tooling
	7. **Exam coverage:** Min 50 validated exams per type. Target 100. Official conditions and UI/UX replicated
	8. **Automated exam launch:** New exam types via Exam Adapter config file (scoring criteria, time limits, sections, question types, skill ontology mappings, AI evaluation prompts, dashboard templates, feedback format). Adding a new exam = new adapter, not code changes
	9. **Per-exam cost monitoring:** Real-time tracking of AI cost per exam generated. Cost per mock must remain under target threshold. Alerts on cost spikes
	10. Product analytics and operational signals wired
Unlocks	*Enterprise and sales readiness. Exam catalogue at commercial scale.*

MILESTONE 4 — Hardening & Commercial Launch (Weeks 13–16)

WEEK 13	
Focus	Retention and network-effect layers, 9 science tools completion.
Acceptance Criteria	1. **Gamification:** XP, streaks, badges, leaderboards live 2. **Per-exam community / forum:** Exam-scoped, moderation, notifications 3. **Academy Benchmarking:** Cross-academy comparison, switching cost, network effect 4. Reputation foundations tied to academy performance visibility 5. **9 science tools:** All 9 (Genome, ERS, Smart Review, Adaptive Sequencer, Diagnostic Feedback, Cognitive Diagnostic, AI Practice Generator, Self-Regulated Learning, Prescription & SLA) in v1 form connected to live flows
Unlocks	*Retention and differentiation on working core.*

WEEK 14	
Focus	Corporate English, avatar lip-sync, mobile hardening.
Acceptance Criteria	1. **Corporate English:** Multi-sector engine. Config-driven sectors, professions, competency maps, generation rules. Min 1 sector validated. Any sector supported from launch 2. Employee and manager/HR views with progress and analytics 3. Material generation with teacher approval 4. **Avatar lip-sync:** Audio → RunPod GPU worker → lip-synced video. Fallback to audio-only. Active for speaking exams with examiner simulation 5. **Mobile full parity:** All web functionalities on mobile. Admin flows optimised for mobile 6. **Offline mode:** Offline study, cached content, queued sync for key flows
Unlocks	*B2B corporate use cases, premium speaking, complete mobile.*

WEEK 15	
Focus	Security, compliance, launch operations, migration.
Acceptance Criteria	1. Security pass complete 2. Load testing with documented results 3. Backup and restore drills successful 4. Rollback drills successful 5. **GDPR/Privacy:** Consent management, data export, PII anonymisation, multi-jurisdictional controls, audit visibility 6. **Feature flags:** Controlled rollout per institution 7. **Migration tooling:** Users, credits, packs, content, historical data migration ready 8. Internal-institution validation for production-critical flows

	9. **Synthetic monitoring:** Active monitoring, audit trails, operational readiness 10. Push and notification hardening complete
Unlocks	*Operational launch readiness.*

WEEK 16	
Focus	Final launch candidate and controlled release.
Acceptance Criteria	1. Full launch candidate across all roles (superadmin, institution admin, teacher, student) 2. Final cross-role walkthrough completed 3. Monitoring and alerting readiness confirmed 4. Controlled release readiness verified 5. Go/no-go confidence for commercial launch 6. Runbooks and operational documentation delivered 7. **App Store:** iOS and Android release candidates submitted or ready 8. **Handover:** Architecture, deployment, and maintenance documentation complete
Unlocks	*Commercial launch.*

5. Exam UI/UX Standards

For every exam at launch, the user experience must be identical to the official exam format. Brian commits to:

- Defining the official-format checklist for each launch exam
- Building or refining the exam shell against that checklist
- Running automated interaction tests per exam type
- Obtaining teacher sign-off on format fidelity
- Minimum 50 validated exams per type at launch; target 100 where content supports it

Supported exams: OET, IELTS, CELPIP, TOEFL, TOEIC, PTE, Cambridge, Trinity.

6. Intellectual Property and Ownership

All code, architecture, documentation, designs, systems, configurations, and any other intellectual property produced during this engagement are 100% owned by Santaliestra Limited.

IP assignment clauses and a non-compete for EdTech / language testing (12–24 months) shall be included in the formal engagement agreement. Brian has signed the ProficientHub NDA.

7. Review and Communication Protocol

- **Weekly review:** Every Thursday at 9:00 AM UK time. Brian presents working product against acceptance criteria
- Both parties document sign-off or exceptions
- Weekly payment of $4,800 USD released through Toptal upon sign-off
- **Task tracking:** Linear. All milestone work broken into tickets. A weekly task list in Linear, agreed by both parties, serves as the reference to determine when a weekly milestone is completed and ready for payment
- Hubstaff runs where possible; binding control is verified deliverables
- Commercial decisions and priority calls must turn around quickly during sprint
- Teachers available for format/UI/UX validation, agent QA, Genome calibration, and any other function Brian requires

7.1 Definition of Done

For each weekly milestone, "Done" means:

- The scoped features for that week are implemented against the agreed acceptance criteria
- Role permissions, tenant rules, and security/privacy requirements for that scope are applied
- The feature is functional end-to-end in staging
- Tests and smoke checks for that scope are passing

- There are no open critical defects blocking normal use of that milestone scope
- The work is demoable in the Thursday review meeting
- Any required setup notes or handover notes are updated

"Done" does not mean code exists. It means the feature is working, testable, reviewable, and ready for the next dependent milestone.

8. Access Requirements (Week 1)

Client provides no later than start of Week 1:

- GitHub: Organisation member access (ability to create repositories and CI/CD workflows)
- AWS: Admin IAM role with full permissions to view and manage all resources
- Cloudflare: Admin access
- Domain/DNS management
- Payment sandbox (Stripe and others)
- Design assets and brand guidelines
- Toptal job status: reactivated and accepted at updated rate

8.1 Developer Handover

A handover meeting shall be arranged between Brian and Anamul (current lead developer) for Anamul to explain the repositories, infrastructure, and provide access to all currently configured systems.

8.2 Existing Tools and Providers

The following tools and providers are already in place and available to Brian:

- Task tracking: Linear
- Email provider: SendGrid
- SMS and WhatsApp: Twilio
- Payment provider: Stripe

Note: The B2B model requires each academy to configure their own email, SMS, WhatsApp, and payment providers. Platform architecture must support both ProficientHub's own provider accounts (B2B billing to academies) and academy-configured provider accounts (B2C billing from academies to students).

9. Technology Stack

Brian has confirmed the following technology stack in his QA Response document:

- **Backend API:** FastAPI for the main platform API and AI orchestration layer. Async-native for AI generation, scoring, and agent workflows
- **Database:** PostgreSQL as the primary database, with schema-level tenant isolation

- **Cache and queues:** Redis for caching, rate limiting, and circuit breaker state. AWS SQS + Lambda / ECS workers for async job queues, generation, scoring, and long-running AI tasks
- **Storage:** AWS S3 for generated exams, audio files, study guides, and reports
- **Web frontend:** React + Shadcn/ui, consistent with the current front-end direction
- **Mobile:** Expo (JavaScript/TypeScript) for unified iOS, Android, and web delivery
- **LLM Gateway:** Custom build on LiteLLM + proxy for model routing, cost controls, semantic caching, and prompt versioning
- **AI Core:** Independent service from the SaaS layer, communicating via internal API
- **Existing Django platforms:** Wrapped behind the FastAPI service boundary and migrated incrementally as each section is rebuilt. Full migration of remaining Django surfaces is a Phase 2/3 decision
- **Infrastructure:** Defined as code (IaC), CI/CD automated, staging + production environments
- **Repository consolidation:** Current multiple codebases consolidated into 1 backend repository, 1 frontend repository, and if needed 1 additional repository

10. Post-Launch Phases

Estimated but not contracted:

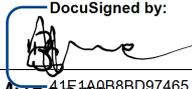| Phase	Duration	Est. Cost
Phase 2: CRM, ERP, Marketplace, Corporate English sectors, connectors	4–5 weeks	$57,600–$63,000
Phase 3: More exams/languages, predictive models, autonomy, voice/avatar, marketplace	4–5 weeks	$57,600–$63,000

Formal proposals for Phases 2–3 after launch. Maintenance arrangement to be negotiated separately.

11. Governing Law and Jurisdiction

This contract shall be governed by and construed in accordance with the laws of Ireland. Any dispute arising out of or in connection with this contract, including any question regarding its existence, validity, or termination, shall be submitted to the exclusive jurisdiction of the courts of Ireland.

12. Signatures

By signing below, both parties confirm agreement to all terms, weekly acceptance criteria, payment conditions, and the early delivery bonus defined in this contract.

CLIENT	**CONTRACTOR**
Santaliestra Limited (ProficientHub)	Brian Musonza
Firmado por:	DocuSigned by:
Jesus Santaliestra Grau	*[signature]*
2A86C9B73A93488...	A1F1A0B8BD97465...
Jesús Santaliestra Grau, CEO	Brian Musonza
Date: _____3/17/2026_____	Date: _____3/17/2026_____