

Ref	Requirement		
	HEADER		
	ICAEW Technical Accreditation Scheme Financial Products Software Evaluation		
	Crescit Pay		
			
	Sept-26		
	© ICAEW. Technical Accreditation Questionnaire		
	CONTENTS		
1	Introduction and Prologue		
2	Issues identified and evaluation conclusion		
	-- GLOBAL REQUIREMENTS:		
3	Access and security		
4	Data processing and reporting		
5	Usability		
6	SaaS hosting		
	-- SPECIFIC REQUIREMENTS: (to follow deep dive demonstration if deemed appropriate)		
7	FX Hedging & Payments		

Ref	Requirement	Response	Reviewer Comments
1.	INTRODUCTION AND PROLOGUE		
Introduction			
1.01	The suitability of software for each particular user will always be dependent upon that user's individual requirements. These requirements should therefore always be fully considered before software is acquired. The quality of the software developers or suppliers should also be considered at the onset.		
1.02	Fundamentally, good software should: 1. Be capable of supporting the functions for which it was designed. 2. Provide facilities to ensure the completeness, accuracy, confidentiality and continued integrity of these functions. 3. Be effectively supported and maintained. It is also desirable that good software should: 4. Be easy to learn, understand and operate. 5. Make best practical use of available resources. 6. Accommodate limited changes to reflect specific user requirements. It is essential, when software is implemented, for appropriate support and training to be available.		
Approach to evaluation			
1.03	The objective is to evaluate a product against a set of criteria developed by the ICAEW to ensure that the software meets the requirements of Good Accounting Software, as laid down in the summary.		
1.04	In order to effectively evaluate the software, a product specialist from the vendor completed the detailed questionnaire and provided it to the ICAEW to examine. The ICAEW's Scheme Technical Manager then reviewed the operation of the various aspects of the software assisted by a member of the vendor's technical staff and checked the answers to confirm their validity. The questions were individually reviewed and commented on and the majority of assessments were confirmed.		
1.05	The Technical Manager discussed the assessment with a member of the vendor's staff in order to clarify any points requiring further information. In the event of disagreement between the supplier and the Technical Manager, the Technical Manager's decision was taken as final and the response was changed accordingly.		
1.06	The latest version of the software was used throughout the evaluation.		
1.07	When the evaluation had been completed, a draft copy was sent to the ICAEW Scheme Manager for review before completion of the final report.		
Prologue: Matters to consider before purchase		NB: This text has been provided directly by the software vendor and does not form part of the ICAEW's & RSM's evaluation	
1.08	General Overview:	CrescitPay is a cloud-based (SaaS) foreign exchange risk management and treasury automation platform for mid-market businesses. It identifies the FX exposure arising from a business's forward obligations — receivables, payables, and purchase and sales orders — quantifies the resulting risk, and supports mitigating it through hedging programmes, forward deal management, hedge accounting and cross-border payments. CrescitPay is the workflow, analytics and automation layer. It does not hold client funds, take positions or act as principal in any transaction: all FX execution and payment settlement is performed by a regulated embedded payments provider operating under its own permissions. The platform is delivered entirely through a web browser, with no client software to install and no customer infrastructure to maintain.	
1.09	Supplier background:	Crescit Pay Ltd is registered in England and Wales, company number 17099549. Incorporated in 2026 following successful completion of the build by Matthew Harris.	
1.10	Product background and suitability for the user:	The product is aimed at finance and treasury teams in businesses with recurring multi-currency exposure that is currently managed on spreadsheets or ad hoc through a bank.	

Ref	Requirement	Response	Reviewer Comments
1.11	Add-on modules:	CrescitPay is delivered as a single, fully integrated application rather than as separately installed modules; there is no inter-module set-up effort. Capability areas are enabled or hidden per organisation through feature toggles, grouped as: FX (converter, live rates, forecast, payment calendar); Risk (deals, FX book, coverage, rollover, FX P&L, policies, alerts, budget rates, VaR, natural hedge, counterparties, group netting, stress test, what-if, intercompany loans); Compliance (IFRS 9, FRS 102 S12 and S30, EMIR, ISDA exposure, audit log); Money (payments, bank connections, reconciliation); Reports (reports, history, deal import, scheduled reports); and Settings (integrations, ERP write-back, consolidation).	
1.12	Typical implementation (size):	The product is targeted at mid-market businesses — CrescitPay's own materials cite a range of approximately £10m to £500m revenue — with recurring FX exposure across multiple currency pairs. Multi-entity groups are supported through consolidation, group netting and intercompany loan functionality. Implementation is configuration rather than deployment: there is no customer infrastructure to provision. Set-up comprises creating the organisation and base currency, adding users and roles, setting approval thresholds and trading limits, defining hedge policies and budget rates, and connecting any ERP/banking integrations, assisted by an in-app onboarding wizard.	
1.13	Vertical applications:	The platform is sector-agnostic: it is designed around the nature of a business's foreign currency exposure rather than around any particular industry, and there are no industry-specific editions or vertical variants. It is most relevant to importers and exporters, manufacturing and distribution businesses, and multi-entity groups with cross-border trade or intercompany flows.	
1.14	Server platform and database:	The service is hosted on Google Cloud Platform (Firebase) and is fully managed and serverless — customers provide no servers, database or infrastructure of any kind. Database: Google Cloud Firestore, a managed NoSQL document database. Application logic and integrations: Node.js Cloud Functions. Web delivery: Firebase Hosting (managed CDN with SSL). Authentication: Firebase Authentication, supporting email/password, Google sign-in and enterprise SSO via OIDC (Microsoft Entra ID, Okta). Secrets are held in Google Secret Manager and are never exposed to the browser. The Firestore database and Cloud Functions are provisioned in Google's europe-west2 (London) region, so customer data resides in the UK. Data is encrypted in transit (HTTPS/TLS 1.2+) and at rest (AES-256).	
1.15	Client specification required:	No client software, plug-ins or downloads are required. CrescitPay is a browser-based single-page application and runs entirely within a standard web browser.	
1.16	Partner network:	Technology and execution partners evidenced in the product are: the embedded payments provider for regulated FX execution and cross-border settlement (Ebury); open banking aggregators (TrueLayer, Plaid); ERP, accounting and accounts payable systems (Xero, Sage, NetSuite, Microsoft Dynamics 365, Coupa, Tipalti, Odoo, Unit4, Basware, Zip); market rate data (ExchangeRate-API); and AI-generated commentary (Anthropic).	

Ref	Requirement		
2.	<u>ISSUES AND CONCLUSION</u>		
Highlighted issues			
2.01	There are a number of limitations in the product, which while not adversely impacting upon this evaluation may be of importance to some organisations. It is important that any business contemplating the purchase of software reviews the functionality described and limitations therein against its detailed requirements. Attention is drawn in particular to the following areas where the product, on its own, may not be suitable for businesses with certain requirements:		
2.02	Findings for considerations by potential customers: (See vendor comments against the various questions)	Question reference	
	* The product's security framework benefits from Google Cloud controls and enterprise SSO, but some elements of CrescitPay's own security assurance remain under development. Native MFA for email/password users, configurable inactivity timeouts and enhanced intrusion monitoring are targeted for 2028, alongside CrescitPay's planned ISO 27001 certification and alignment with the NIST AI Risk Management Framework.	3.09, 3.17, 3.22, 6.03, 6.36 - 6.38	
	* Service delivery depends materially on Google Cloud and regulated execution partners, while data is hosted in a single London region with no user choice of jurisdiction. Users can export data but cannot perform a full self-service restore, and the service has a limited operating history against its 99.9% availability target. Users should therefore consider the contractual responsibilities, resilience, recovery and contingency arrangements across the supplier chain.	6.01 - 6.02, 6.27 - 6.29, 6.51 - 6.56, 7.30, 7.40 - 7.45	
	* Integration is available through APIs, connectors, open banking and scheduled or on-demand synchronisation, but manual import is limited to CSV and scheduled ad-hoc file transfers are not supported. Native Excel import is targeted for 2028, so users should confirm that the available formats and interfaces meet their operational requirements.	4.19, 7.59 - 7.60, 7.63 - 7.64	
	* The product offers limited scope for tailoring data capture and retaining supporting evidence. User-defined fields, transaction attachments and a document repository are not supported. Attachment capability may be added from 2028, but no delivery target is stated for user-defined fields or document storage.	4.17, 5.10, 5.46	
	* Reporting and presentation options are relatively standardised. Reporting is based on built-in outputs, saved searches are limited, dashboards have fixed layouts, and the English-only interface cannot be branded or white-labelled. Localisation, saved filters and configurable dashboards are roadmap items without firm delivery dates.	4.36 - 4.37, 5.01 - 5.02, 7.76	
	* Records are not explicitly locked, with editable fields operating on a last-write-wins basis. Approval workflows also lack formal out-of-office delegation, so temporary cover requires another approver or a manual permission change. Organisations should assess whether these alternative controls meet their concurrency and access requirements.	5.05 - 5.06, 5.12	
	* Users cannot test releases before production or receive advance notice before changes enter service. User-facing updates are communicated by email and new features may be disabled, but organisations with formal change-control requirements should assess whether these arrangements are sufficient.	6.68 - 6.69, 6.73	

Ref	Requirement		
	* Support is provided through an account manager, general email inbox or telephone line rather than a user-facing ticketing system. A ticketing capability is included in the Series A roadmap without a stated delivery date, so users should confirm how requests and incidents will be tracked and evidenced in the interim.	5.32	
	* The product supports spot and forward FX contracts but not more complex instruments such as options or swaps. Users requiring a broader range of hedging instruments should assess product suitability, as no delivery target has been stated for extending this coverage.	7.31	
	* The product is subject to automated vulnerability scanning using OWASP ZAP quarterly and before major releases, with any adverse findings stated to be remediated immediately. However, automated vulnerability scanning is not equivalent to penetration testing performed by an appropriately qualified security specialist. Organisations should therefore consider whether independent specialist penetration testing and evidence of remediation are required to meet their security assurance requirements.	6.46 - 6.48	
Evaluation conclusion			
2.03	For the specific use-cases in support of assisting accountancy firms to make effective use of Financial Products software, for which the product is designed, the solution appears to meet this criteria. A notable strength is the platform's integrated approach: exposure flow directly into deal booking, market valuation, hedge accounting and reporting. This is supported by policy controls, four-eyes approval, real time risk analyse, and stress testing against historical and hypothetical market scenarios. It continues to be actively developed and enhanced. Planned Series A developments from 2028 include native MFA, enhanced security monitoring and broader import capabilities, including native Excel files. Members should be aware of the limitations of the solution as above, and fully understand the role that it can play in helping manage their compliance needs. * NOTE THAT THE QUESTIONNAIRE RELATES TO THE SOFTWARE PRODUCT AND NOT ANY SUPPLEMENTARY SERVICES PROVIDED BY THE SUPPLIER TO THE ACCOUNTANCY FIRM USING THAT PRODUCT *		
Disclaimers			
2.04	Any organisation considering the purchase of this software should consider their requirements in the light of proposals from the software supplier or its dealers and potential suppliers of other similarly specified products. Whilst the contents of this document are presented in good faith, neither ICAEW, nor the ICAEW's Technical Manager (RSM UK Consulting LLP or any party nominated by the ICAEW to perform this role on the ICAEW's behalf) will accept liability for actions taken as a result of comments made herein. The decision to purchase software resides entirely with the organisation.		

Ref	Requirement	Response	Reviewer Comments
3.	<u>ACCESS AND SECURITY</u>		
Access control			
3.01	What security features are included to control access to the application?	Access is controlled by Firebase Authentication (email/password, Google, and enterprise SSO). All data access is additionally governed server-side by Cloud Firestore security rules, which enforce authentication, per-organisation data isolation (a user can only reach records for their own orgId), role checks, and an explicit deny-all catch-all so any collection not expressly permitted is blocked. The React application layers route guards on top of this (unauthenticated users are redirected to the login page; role-restricted areas are gated by dedicated guards). Suspended users and suspended organisations are blocked at the rules layer and force-signed-out. Application secrets (API keys, bank/ERP OAuth tokens) are held in Firebase Secret Manager and never exposed in the browser.	Noted.
3.02	Can access to functions be managed via a permissions matrix so users can only see (in menus and other links) and access those areas they are authorised to access?	Yes. Access is governed by a role-based permission model that functions as a permissions matrix. Each user holds an organisation role (org_admin, approver, trader, viewer, or a time-limited read-only auditor). A single permissions module derives capabilities (canWrite, canApprove, canAdmin, isAuditor) from that role. These are enforced both in the UI (menu visibility and route guards) and, critically, at the Cloud Firestore rules layer, so a user cannot reach data or actions outside their authority even outside the UI. Each organisation can also enable/disable individual feature areas via a per-feature visibility map.	Noted.
3.03	Is this access to the application managed by:- - Individual user profiles? - User groups or job roles?	Both. Each user has an individual profile document holding their identity, platform role, organisation membership and organisation role. Authority is then applied through user groups / job roles, i.e. the organisation role hierarchy: org_admin > approver > trader > viewer, plus a dedicated read-only auditor role.	Noted.
3.04	Can a report be produced detailing all current users, their user groups if relevant, and their authority levels and/or access rights?	Yes. The Organisation Admin screen lists all current members with their role, account status, MFA-enrolment flag, last login, per-user trading limit and (for auditors) access-expiry date, and administrators can review and change roles from there. The audit log (which records role changes and other user-management events) can additionally be exported to Excel for a point-in-time record.	Noted.
3.05	If menus can be tailored does the system limit the display of menu options to those for which permission has been granted for each user?	Yes. Menus are tailored to the user's role and to the organisation's enabled features. The navigation is filtered so that, for example, an auditor sees only the Compliance area, non-administrators do not see Settings/admin areas, and any feature an organisation has disabled is hidden. Menu options are only displayed where the corresponding permission has been granted.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
3.06	Does security allow for access to be limited to: - Read only? - Read/write? - Read/amend/delete?	Access levels map to roles: read-only (viewer and auditor), read/write (trader and above), and approval authority (approver and org_admin). Financial transaction records (hedge deals, payments) are deliberately NOT deletable by any client role — they are cancelled via a controlled status change rather than removed — so the strict 'read/amend/delete' level does not apply to core transactions by design. Certain configuration/reference records (e.g. beneficiaries, policies) can be amended and removed by authorised roles.	Confirmed.
3.07	If data can be accessed by separate reporting facilities, such as ODBC or an external report writer, is the user access security control applied?	Not applicable in the ODBC sense. The database (Cloud Firestore) is not exposed via ODBC or an external report writer; there is no separate reporting connection to secure. All data access — including in-app Excel/PDF exports — passes through the same authenticated, org-scoped, role-checked security rules, so reporting is subject to the identical access controls as the rest of the application.	Noted.
3.08	Does the system security integrate with Microsoft's Active Directory or other tools that provide a single sign-on?	Yes. The system supports single sign-on via Firebase Authentication OIDC providers, including Microsoft Entra ID (Azure AD / Active Directory), Okta, and Google Workspace. SSO is detected by the user's email domain at login; for Google Workspace the sign-in is restricted to the registered hosted domain. SSO provider secrets are held server-side and never exposed to the browser.	Confirmed.
3.09	Does the system provide multi-factor authentication (MFA)?	MFA is supported by the underlying Firebase Authentication platform and the data model already carries organisation ('MFA required') and per-user ('MFA enrolled') flags. However, an in-application enrolment/enforcement flow is not implemented in the current codebase; for SSO users, MFA is enforced by the corporate identity provider (Entra ID / Okta). In the interim, accounts without MFA are protected by the Firebase Authentication password policy (minimum length and complexity), short-lived session tokens (~1 hour) that are silently refreshed only while valid, strict per-organisation data isolation enforced server-side, forced sign-out of suspended users, and the immutable audit trail. Native in-app MFA enrolment and enforcement — including the ability for an organisation to make MFA compulsory for all its users is a planned enhancement as part of Series A, target 2028.	Confirmed.
Passwords and access logs			
3.10	Is access to the software controlled by password?	Yes. Email/password sign-in is provided via Firebase Authentication, alongside Google and enterprise SSO options.	Confirmed.
3.11	Does each user have a separate log on (user ID)?	Yes. Every user authenticates with their own credentials and is issued a unique user ID (Firebase Auth UID), with their own individual profile document. There are no shared logins.	Noted.
3.12	If there is no password facility please state how confidentiality and accessibility control is maintained within the software?	Not applicable — a password facility is provided (see 3.10).	N/A.

Ref	Requirement	Response	Reviewer Comments
3.13	Are passwords masked for any user logging in?	Yes. The password field on the login screen uses a masked (type="password") input, so characters are not displayed as they are typed.	Confirmed.
3.14	Is password complexity available and enforced?	Password strength is governed by the Firebase Authentication password policy configured for the project (minimum length and complexity rules are set at the identity-platform level). Custom complexity rules are not additionally enforced in the application code. For SSO users, password complexity is dictated by the corporate identity provider.	Noted.
3.15	How many previous passwords are retained / the password history?	Password history is not retained at the application level; the application never stores passwords. For email/password accounts this is governed by the Firebase Authentication platform, and for SSO accounts it is governed by the organisation's identity provider (Entra ID / Okta), where password-history policy is typically configured.	Noted.
3.16	Are passwords encrypted?	Yes. The application never stores or handles raw passwords itself — credentials are managed by Firebase Authentication, which stores them as salted hashes, and all traffic is over TLS/HTTPS. SSO credentials never reach the application at all.	Noted.
3.17	Are users automatically logged off after a pre-set idle time? - Can the time period be changed? - Can any information be viewed without being logged in, including after logging off, if so what information?	There is no configurable idle-timeout auto-logout implemented in the application at present. Sessions rely on the Firebase Authentication ID token, which is short-lived (approximately one hour) and silently refreshed while the session is valid; suspended users are force-signed-out. No protected data can be viewed without an active authenticated session — unauthenticated users are redirected to the login page and protected routes render nothing. A configurable inactivity timeout is a recommended future enhancement as part of series A - target 2028.	Noted.
Deletion of transactions			
3.18	Is it possible to delete a transaction?	Core financial transactions cannot be deleted. Hedge deals and payments have client deletion explicitly disabled at the Firestore rules layer; they are instead cancelled through a controlled status change. Certain configuration/reference records (e.g. beneficiaries) can be deleted by authorised users.	Confirmed.
3.19	If so, then how are deletions controlled by the system?	Cancellation is performed through a dedicated, controlled flow rather than a raw delete. Cancelling a deal transitions its status to 'cancelled' within an atomic batch (also releasing any linked pipeline items), writes an audit-log entry and notifies the arranger. The security rules only permit specific, valid status transitions (e.g. open→cancelled, pending_approval→cancelled), and approvals enforce four-eyes (the approver must differ from the arranger). The original record and its financial fields remain immutable throughout.	Noted.

Ref	Requirement	Response	Reviewer Comments
3.20	Are deleted transactions retained in the audit trail (see below) and denoted as such?	Yes. Nothing is physically removed — the transaction is retained with a 'cancelled' status and a corresponding 'deal_cancelled' entry is written to the append-only audit log, so cancellations remain fully visible and traceable.	Confirmed.
Audit trails			
3.21	Does the system have an audit trail (log) which records all changes to transactions in the system?	Yes. An append-only audit log records key changes and actions across the system, including deals booked, approved, cancelled and settled, pipeline additions/deletions, user-management actions, invites, and organisation changes. Each entry captures the acting user's ID, name and email, a target reference and a human-readable description.	Confirmed.
3.22	Does this log also record any system error messages and/or any security violations?	Partially. Business and security-relevant events (role changes, suspensions, cancellations, approvals, etc.) are captured in the application audit log. System error messages and access-rule denials are recorded separately in Google Cloud Logging by the Cloud Functions and the Firestore platform rather than in the application audit-log collection.	Noted.
3.23	Is it possible to turn off or delete the audit trail?	No. The audit log cannot be turned off or deleted. Entries are append-only and updates and deletes are explicitly forbidden at the security-rules layer, and the recorded actor is forced to match the authenticated caller so entries cannot be fabricated against another user.	Confirmed.
3.24	Does the software allocate a system generated sequential unique reference number to each transaction in the audit log, date and time stamp it and record the user ID?	Partially, as described. Each entry is assigned a unique reference (a system-generated document ID), is date/time-stamped, and records the user ID (plus name and email) of the actor, with the actor enforced to be the authenticated caller. The reference is unique but is not a running sequential number.	Noted.
3.25	Are all master file changes recorded in the audit trail?	Yes, for user master data. User-management changes — role changes, trading-limit changes, suspensions/activations, removals and invites — are written to the audit log. Some organisation-level configuration changes are managed through admin-only, rules-protected paths but are not all individually logged.	Noted.
Compliance			
3.26	Does the system operate in a way that is compliant with data protection legislation including GDPR? How does the system facilitate this?	The architecture is designed to support data-protection compliance: all data is hosted in the EU (Cloud Firestore and Cloud Functions are pinned to the europe-west2 / London region for data residency); access is authenticated, role-based and strictly isolated per organisation; personal data held is minimal (name and email); and an immutable audit trail evidences access and change. Deletions of user records are handled through controlled server-side (Cloud Function) paths. Formal data-subject-request tooling (automated export/erasure) is not yet a packaged feature and would be handled operationally.	Noted.

Ref	Requirement	Response	Reviewer Comments
3.27	Describe your use of sub-processors if any?	The platform relies on the following sub-processors. Core infrastructure: Google Firebase / Google Cloud (Authentication, Firestore, Cloud Functions, Secret Manager), EU-hosted. FX rates and execution: Visa, Ebury and ExchangeRate-API. AI-generated narratives: Anthropic (Claude API). Optional, opt-in integrations enabled per organisation include open-banking aggregators (TrueLayer, Plaid) and ERP/AP connectors (NetSuite, Microsoft Dynamics 365, Coupa, Tipalti, Odoo, Unit4, Basware, Zip). Optional integrations are only active where an organisation configures them.	Noted.
Backup and recovery			
3.28	Is there a clear indication in the software or manuals as to how the data is backed-up and recovered?	Backup and recovery are provided by the underlying managed platform (Google Cloud Firestore) rather than by a bespoke, separately documented mechanism within the application.	Noted.
3.29	How often are backups taken and to what point can restores be done?	Backups are a managed capability of Cloud Firestore, which provides scheduled backups and point-in-time recovery to any point within the preceding seven days.	Noted.
3.30	How does the software facilitate recovery procedures in the event of software failure? (E.g. roll back to the last completed transaction).	Firestore is a managed, highly-available database; writes are atomic and there is no partial-write corruption to roll back at the application level. Multi-record operations (for example, cancelling a deal and releasing its linked pipeline items) are performed in atomic batches that commit all-or-nothing. Recovery from data loss is via the platform's point-in-time recovery / backups (see 3.29).	Noted.
3.31	If software failure occurs part way through a batch or transaction, will the operator have to re-input the batch or only the transaction being input at the time of the failure?	Only the individual action. Multi-record operations use atomic batches/transactions that commit all-or-nothing, so a failure part-way through rolls the operation back rather than leaving a partial commit. The operator re-submits the single failed action (e.g. re-books the one deal) rather than re-inputting a whole batch of otherwise-successful work.	Noted.
3.32	What features are available within the software to help track down processing problems?	Cloud Functions emit structured logs (including error context) to Google Cloud Logging, and Firestore records security-rule denials and authentication failures at the platform level, giving a server-side trail for diagnosing processing problems. Client-side errors surface in the UI. A dedicated in-application error-monitoring dashboard (e.g. a third-party error tracker) is not currently integrated.	Noted.

Ref	Requirement	Response	Reviewer Comments
4.	<u>DATA PROCESSING AND REPORTING</u>		
Input and validation of transactions			
4.01	Is data input controlled by self-explanatory menu options?	Yes. Data entry is driven through a clear, labelled navigation menu and purpose-built screens/forms (e.g. book deal, add exposure, create payment, set budget rate), each with descriptive fields, help tooltips and an onboarding wizard, rather than through raw configuration.	Confirmed.
4.02	Are these menus user/role-specific?	Yes. Menus are role- and feature-specific. The navigation is filtered by the user's organisation role (for example an auditor sees only the Compliance area and non-administrators do not see admin/settings areas) and by the features each organisation has enabled, so users only see the input options relevant to their permissions.	Confirmed.
4.03	Can the creation or amendment of standing data (e.g. customer account details) be undertaken using menu options or dialogue boxes as opposed to requiring system configuration?	Yes. Standing/reference data — such as beneficiaries, budget rates, approval rules, hedge policies and organisation settings — is created and amended through ordinary in-app forms and dialog screens by authorised users, not through back-end configuration.	Confirmed.
4.04	Does the software provide input validation checks such as: - (account) code validation? - reasonableness limits? - validity checks?	Yes. Input is validated on entry — for example amounts must be positive, numeric and within bounds; currencies must be recognised ISO codes and must differ from the base currency; dates must be valid; and required fields must be present. Both the manual entry forms and the bulk import apply these code-validity, reasonableness and validity checks.	Confirmed.
4.05	What control features are within the software to ensure completeness and accuracy of data input?	Completeness and accuracy are supported by required-field enforcement, typed inputs, ISO currency-code validation, reasonableness bounds on amounts/rates, date validity checks, and — for financial records — immutability of key fields once written (notional, rate, amounts, counterparty cannot be altered afterwards, enforced at the database-rules layer). Bulk imports show a pre-commit preview flagging every invalid row before anything is saved.	Confirmed.
4.06	How does the software ensure uniqueness of the input transactions? (i.e. to avoid duplicate transactions)	Uniqueness is enforced by a deterministic, idempotent fingerprint on imported/interfaced exposures. Where an upstream reference exists (e.g. an ERP invoice or PO number) it is trusted as the key; otherwise the fingerprint is composed from the invariant fields (currency, direction, value date, amount to 2 d.p. and counterparty), scoped per organisation. Re-importing the same exposure produces the same fingerprint and is de-duplicated rather than duplicated. Each booked deal also receives its own unique system-generated identifier.	Noted.

Ref	Requirement	Response	Reviewer Comments
4.07	Is data input by users validated by scripts or routines in the browser, or other client software, before transmission to the server?	Yes. Input is validated client-side in the browser before submission — form-level checks on amounts, currencies, dates and required fields, and per-row validation on imports — giving immediate feedback before any data is transmitted to the server. The following are enforced server-side and cannot be bypassed from the browser: authentication and organisation membership (a user can only reach records for their own orgId); role/authority checks (write, approve, admin); immutability of key financial and identity fields once written (notional, rate, amounts, counterparty); permitted status-transition rules (e.g. open→cancelled, pending_approval→cancelled only); four-eyes approval (the approver must differ from the arranger); a default-deny catch-all blocking any access not explicitly permitted; and, in Cloud Functions, validation of required parameters and credentials before any action is taken. These are enforced at the Cloud Firestore security-rules layer and in server-side Cloud Functions, i.e. authoritatively at the server regardless of the client.	Noted.
4.07 (Cont.)		Field-level value validation (numeric ranges, date formats, ISO currency codes) is primarily performed client-side for immediate feedback; the server layer focuses on authorisation and data-integrity constraints.	
4.08	Is data input by users validated by routines running on the server before data files are updated?	Partially. Server-side Cloud Functions validate required parameters and credentials before acting, and the database security rules enforce authoritative server-side constraints on every write — organisation membership, role, immutability of financial/identity fields, and permitted status transitions — so invalid or unauthorised writes are rejected at the server regardless of the client. Field-level value validation (e.g. numeric ranges) is primarily performed client-side; the server layer focuses on authorisation and integrity constraints.	Noted.
4.09	Does the above validation ensure that data entered in all input boxes: - Cannot be longer than a maximum length? - Cannot contain unaccepted characters such as semi-colons etc?	Partially. Client-side forms and the import validator constrain field content — rejecting empty values, non-numeric amounts, unrecognised currency codes and invalid dates. Because the data store is a document database rather than SQL, classic injection characters such as semi-colons are not a query risk; values are stored as typed fields, not concatenated into queries. Explicit maximum-length limits are applied on some fields but are not uniformly enforced on every free-text input. The fields without an explicit maximum-length constraint are the free-text descriptive fields rather than the structured/typed financial fields. Examples are the free-text note/description fields on deals and payments and the counterparty and beneficiary name/reference fields.	Noted.
4.10	Are input errors highlighted?	Yes. Validation errors are highlighted inline on the entry forms, and on bulk import each problem row is clearly flagged in the preview with its specific error(s).	Confirmed.
4.11	If Yes are they: - Rejected and error report generated on-screen? - Rejected and error reports generated? - Accepted and posted to a temporary account/area?	Rejected with an on-screen error. Invalid manual entries are blocked at the point of entry with a clear message. On import, invalid rows are identified in the on-screen preview and are not committed; only valid rows are imported. There is no 'accept-and-post-to-a-suspense-area' behaviour — invalid data is rejected rather than quarantined.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
4.12	Are responses to erroneous data input clear so that they do not lead to inappropriate actions?	Yes. Error messages are specific and plain-language (e.g. 'notional must be a positive number', 'unknown currency', 'settlement_date must be YYYY-MM-DD'), so the user is told exactly what to correct rather than being led into an inappropriate action.	Confirmed.
4.13	Does the software have an automatic facility to correct/reverse/delete transactions?	There is no automatic correction facility. Financial transactions (deals, payments) cannot be edited in their key financial fields or deleted; instead a deal is reversed through a controlled cancellation flow that sets its status to 'cancelled'. Corrections are therefore made by cancelling and re-booking, preserving a complete record.	Noted.
4.14	If yes, are these logged in the audit trail?	Yes. Cancellations/reversals are written to the append-only audit trail (a 'deal_cancelled' entry recording the acting user, timestamp and target), and the original record is retained with its updated status.	Noted.
4.15	Are all data entries or file insertions and updates controlled to ensure that should part of a data entry fail the whole transaction fails?	Yes. Multi-record operations are performed as atomic batches/transactions that commit all-or-nothing, so if any part of the operation fails the whole operation is rolled back and no partial write is left behind.	Noted.
4.16	Are messages provided to users clearly explaining whether the data entry or file upload has been processed successfully or not?	Yes. The user receives clear success or failure feedback — confirmation on successful entry, and on bulk import a results banner summarising how many rows succeeded and how many failed, with the failed rows itemised.	Confirmed.
Import and export of data			
4.17	Can files/attachments be uploaded and stored against any transaction?	No. The system does not currently support uploading and storing arbitrary files/attachments against individual transactions. Transactions instead carry structured references — including an upstream source reference (ERP invoice/PO number) where interfaced — rather than attached documents. Document attachment is a potential future enhancement subject to demand - Series A with a target of 2028 is the earliest estimated time, we would need much more data on the use case to build it efficiently.	Noted.
4.18	Is there an additional charge made for storage of uploaded files? - If yes, please indicate the cost.	Not applicable — file/attachment storage against transactions is not offered (see 4.17), so no storage charge arises.	Noted.
4.19	Can data be imported into the system from multiple types of files, e.g. XLS, text, CSV?	Import is currently supported from CSV files (a downloadable CSV template is provided). Direct import of native Excel (XLS/XLSX) or free-text formats is not offered in the manual importer; however, transactions and exposures can also be interfaced automatically from connected ERP and open-banking systems. Broadening manual import to additional file types is a straightforward enhancement. Manual import is currently from CSV (with a downloadable template and per-row validation). Should native Excel (XLSX) import be added, it would reuse the existing import pipeline — the same row-by-row validation (required fields, recognised currency, base-currency difference, valid direction, positive numeric notional and rates, valid dates), the same pre-commit preview flagging each valid/invalid row, and the same fingerprint de-duplication — so the control environment would be identical to CSV import. It is not planned for any earlier than Series A, with a target date of 2028.	Noted.

Ref	Requirement	Response	Reviewer Comments
4.20	Explain how the system validates imports into the system and what happens to any import which fails?	Each row is validated before anything is written. The importer parses the file and checks every row (required fields, recognised currency, base-currency difference, valid buy/sell direction, positive numeric notional and rates, valid dates) and presents a pre-commit preview marking each row valid or invalid with its specific error(s). Only valid rows are imported; invalid rows are left out and reported back with a success/failure count so they can be corrected and re-submitted. Re-imports are de-duplicated by fingerprint.	Noted.
4.21	Are imported /interfaced transactions detailed in the audit trail? (See also 3.27)	Yes. Imported/interfaced transactions are recorded in the audit trail (an 'external_deal_imported' action is logged), consistent with the audit-trail controls described in Section 3.	Noted.
4.22	Can data be exported from all areas of the system to multiple formats e.g. XLS, CSV, PDF, text; if so specify which formats are supported?	Yes. Data can be exported broadly across the system. Excel (.xlsx) and CSV export are available from data grids and reports (via SheetJS), and PDF export is available for formatted reports (e.g. the Board Report and disclosure reports, via jsPDF). Supported export formats are therefore XLSX, CSV and PDF.	Confirmed.
Data processing			
4.23	Does the software ensure that menu options or programs are executed in the correct sequence (e.g. outstanding transactions are processed before month end is run)?	Correct sequencing is enforced through controlled status transitions rather than a batch-run order. A deal must pass through valid states (e.g. pending approval → open → settled/cancelled), approvals require a different user from the arranger (four-eyes), and these transitions are validated at the database-rules layer so steps cannot be performed out of order. The system is real-time rather than batch-sequenced, so there is no 'must run X before month-end' dependency.	Confirmed.
4.24	Does the software provide automatic recalculation, where appropriate, of data input? (e.g. VAT)	Yes, where applicable. FX conversions, forward rates, hedged costs, mark-to-market valuations and coverage figures are recalculated automatically from the underlying inputs and live or booked rates as data changes. VAT/tax recalculation is not applicable — this is an FX treasury and hedging system, not a tax ledger.	Confirmed.
4.25	Is a month/period-end routine required to be undertaken?	No mandatory period-end close routine is required — the system operates in real time. Period-end reporting (e.g. FRS 102 / IFRS 9 fair-value and P&L disclosures, month-end management figures) is available on demand for any chosen period without a locking close process.	Noted.
4.26	Is it possible to delete accounts if the balance is Nil but transactions have been recorded against the code?	Not applicable in the general-ledger sense — the system does not maintain a chart of nominal accounts that can be deleted. Financial transactions themselves (deals, payments) cannot be deleted at all; they are retained and cancelled by status, so there is no route to remove a record that has activity against it.	Noted.
4.27	What is the size and format of reference numbers and descriptions within: - Ledgers? - Stock? - Currencies?	Ledger/stock reference numbering in the traditional accounting sense does not apply. Transactions (deals, payments, exposures) are assigned unique system-generated identifiers and carry free-text descriptive fields (e.g. counterparty, notes). Currencies use standard ISO 4217 three-letter codes (21 major currencies supported). There is no stock/inventory module.	Noted.

Ref	Requirement	Response	Reviewer Comments
4.28	How does the software guard against/warn about duplicate account numbers on set up?	For organisation identifiers, a unique organisation code is generated and used, avoiding collisions. For imported/interfaced exposures, the deterministic fingerprint (see 4.06) prevents duplicate transactions on set-up and re-import. The system does not maintain a manually-numbered chart of accounts, so account-number duplication does not arise.	Noted.
4.29	How does the software enable the traceability (from, to and through the accounting records) of any source document or interfaced transaction?	Traceability is provided through linked records and the audit trail. Booked deals are linked to the specific pipeline exposures they hedge (and vice versa), interfaced items retain their upstream source reference (ERP invoice/PO number), and every material action is recorded in the append-only audit log with actor, timestamp and target. This allows an item to be traced from source document/interface through to the accounting and reporting outputs.	Confirmed.
4.30	What drill down/around functionality is available within the software?	Summary screens and dashboards drill through to the underlying detail — for example dashboard coverage figures resolve to the individual exposures and the deals hedging them, and the hedge-accounting views tie deals to their designations and to the linked exposures. Reports and grids can be filtered and exported for further analysis.	Confirmed.
4.31	If the software uses a lot of standing information which changes frequently or regularly, does the software allow for such changes to be effected through the use of parameters or tables?	Yes. Frequently-changing standing information is held as configurable parameters/tables maintained through the UI — for example budget rates, approval-threshold rules, hedge policies, enabled features and organisation settings — so such changes are effected through data configuration rather than code changes.	Confirmed.
Report writer			
4.32	Does the system have an in-built report generator or is a third-party solution used (if so please specify)?	Reporting is delivered by the system's own built-in report generation rather than a third-party report writer. This comprises a formatted Board Report (PDF), analytics summaries, and a broad set of purpose-built reports (deals, FX P&L, VaR, stress test, EMIR, FRS 102 / IFRS 9 disclosures, hedge coverage, etc.), each exportable to Excel/CSV/PDF.	Confirmed.
4.33	Is the report writer based on a standard SQL-type approach and is it flexible and easy to use?	No. There is no general-purpose, SQL-based ad-hoc report writer. Because the data store is a document database, reporting is provided through the curated built-in reports plus flexible export of any data grid, rather than a free-form SQL query tool. There is no customer-facing BI/ODBC API or ad-hoc query endpoint at present. External analysis is supported through in-application data extraction — Excel (XLSX), CSV and PDF export from data grids and reports — which customers can feed into their own BI or spreadsheet tools. The internal integration layer (Cloud Functions with partner APIs/webhooks) serves system-to-system connectors (ERP, banking, FX execution) rather than general BI extraction. On audit logging: the append-only audit log records key business events (deals booked/approved/cancelled/settled, user-management actions, etc.); user-initiated report/export actions are also logged.	Noted.
4.34	Can the report generator operate over the financial and operational aspects of the system, e.g. combining service metrics with financial information?	Yes, within the built-in reports. Several reports already combine financial and operational information — for example hedge-coverage reporting combines exposures and positions against policy targets, and the Board Report combines positions with risk scenarios. There is not, however, a user-driven query builder for arbitrary cross-domain combinations.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
4.35	Is a comprehensive data dictionary provided to aid field selection?	No. A comprehensive data dictionary for ad-hoc field selection is not provided, as there is no free-form report builder that would require one. Exported data grids carry clear, labelled column headers.	Noted.
4.36	Does the system provide a library of reports and templates which can be amended, saved and re-run?	Partially. A library of built-in reports and templates (Board Report, disclosure reports, and the selectable scheduled-report types) can be re-run at any time. User-authored, saveable custom report templates are not currently supported beyond configuring the scheduled-report selection.	Confirmed.
4.37	Can users create their own reports? If so, what are the controls on users doing this?	Not in the sense of building bespoke reports. Users cannot author arbitrary custom reports; they run the built-in reports (subject to their role and the organisation's enabled features), apply on-screen filters, and export the results. Access to reporting is governed by the same role-based permission model as the rest of the system.	Confirmed.
4.38	Can users create saved searches/filters/queries?	On-screen filtering of grids and reports is supported. Persistent, named saved searches/queries are limited: the main saved, re-usable configuration is the per-user scheduled-report selection. Broader saved-filter functionality is a candidate enhancement.	Confirmed.
4.39	Can regular reports be added to user menus in the appropriate area of the system?	Reporting areas are surfaced in the role- and feature-driven navigation menu, so the relevant reports appear in the appropriate section for each user. Pinning individual custom reports to a user's menu is not a distinct feature; scheduled reports are instead delivered automatically by	Confirmed.
4.40	Does the system support the production of on demand (interactive) and scheduled batch reports?	Yes. On-demand interactive reports are available throughout the application, and scheduled batch reports are supported via automated email delivery on a daily (08:00 UTC), weekly (Monday 08:00 UTC) or monthly (1st of the month, 08:00 UTC) cycle, with each user selecting which report types they receive.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
5.	USABILITY		
Ease of use			
5.01	Does the solution provide a multi-language user interface?	No. The user interface is currently provided in English only; no multi-language/localisation framework is present in the application. (The platform supports international currencies and formats, but the UI language itself is not switchable.) The target customers are UK-based firms initially, but future localisation functionality is in the Series A roadmap.	Confirmed.
5.02	Does the system allow for customisable branding and UI (e.g. corporate colour palate, upload company logo, etc)?	Not currently. The application uses CrescitPay's own branding. Per-customer branding — uploading a company logo or applying a corporate colour palette — is not offered in-app today. The interface is built on a consistent internal design-token system, and organisations can toggle which feature areas are visible, but customer-facing white-label branding is not a current feature.	Confirmed.
5.03	Does the system have a similar look and feel and overall consistency between screens and modules?	Yes. The application uses a shared design system — consistent design tokens (colours, typography, spacing), a common layout and navigation, and reusable UI components — giving a consistent look and feel across all screens and modules.	Confirmed.
5.04	Is data entry easily repeated if similar to previous entry?	Partially. Re-entry is reduced through defaults and reuse: the organisation base currency, saved beneficiaries and budget rates pre-populate where relevant; exposures imported or synced from ERP/banking systems pre-fill deal bookings; and the 'what-if' tools let users iterate on a scenario. There is not a one-click 'duplicate the previous transaction' action on every entry screen.	Confirmed.
5.05	Does the software prevent access to a record while it is being updated?	The platform uses a real-time database with optimistic concurrency rather than pessimistic record locking, so records are not locked during editing. Integrity is protected instead by making key financial fields immutable once written and by using atomic transactions for multi-record operations, so concurrent activity cannot corrupt data; for editable fields a last-write-wins model applies.	Noted.
5.06	Is there locking at file or record level?	There is no explicit file- or record-level locking. Concurrency is handled optimistically, with atomic transactions and field immutability preventing conflicting or partial updates rather than locks.	Noted.
5.07	Does the software allow for the running of reports whilst records are being updated?	Yes. Because there is no locking and data is read in real time, reports and views can be run at any time, including while other users are updating records.	Confirmed.
5.08	Can timestamps or user comments be added to transactions?	Timestamps: yes — every record is server-timestamped, and the immutable audit trail records who made each change and when. User comments: limited — records carry structured reference/description fields and the audit trail captures human-readable descriptions, but free-text user comment threads attached to individual transactions are not a distinct feature.	Noted.

Ref	Requirement	Response	Reviewer Comments
5.09	Is there the ability to store preferences and default values on a per-user basis e.g. department/team/user?	Partially. Per-user settings include help-hint visibility and scheduled-report preferences, and notifications are held per user. Organisation-level defaults (base currency, enabled features, budget rates, approval rules) apply across the organisation. Distinct department/team-level default profiles are not a separate feature.	Confirmed.
5.10	Does the system have the ability to provide user-defined fields with associated validation of data input?	No. The data model is fixed; user-defined custom fields with their own validation rules are not supported.	Noted.
5.11	Can the system provide users with reminders and notifications e.g. workflows?	Yes. The system provides in-app notifications, rate alerts / target orders (the platform monitors live rates and flags when a target is hit), an approval workflow for deals above configured thresholds, and scheduled email reports.	Confirmed.
5.12	If the system provides workflows, does it have functionality to substitute/delegate authorisations?	No. Approvals are governed by role (approver / org_admin) with four-eyes enforcement (the approver must differ from the arranger). Out-of-office delegation or substitution of a specific individual's authority to another named user is not currently supported — any user holding the approver role can act. In the case of approver absences, the org admin can change permissions of other users. These will need to be made manually and returned to their normal levels when the temporary authority is no longer needed.	Noted.
5.13	Describe the tools and features available for a power user to make configuration changes such as amending a workflow.	Configuration is performed through in-app administration screens rather than a free-form workflow designer. Power users/administrators can maintain approval-threshold rules (including multi-tier chains and required approver counts), hedge policies, budget rates, per-organisation feature toggles, SSO settings, integrations, and user roles/trading limits. These are data-driven settings applied immediately.	Confirmed.
5.14	Is there the ability for users to define and configure layouts of letters and forms?	Limited. Formatted outputs — such as the board report and the FRS 102 / IFRS 9 disclosure narratives — are generated by the system to a fixed layout; users cannot freely design their own letter or form templates. Data and reports export to Excel/PDF for any downstream formatting the user requires.	Noted.
5.15	Can users save the parameters of searches?	On-screen filters are available within each area, and scheduled-report selections are saved per user, but named, reusable saved searches spanning the application are not a general feature.	Noted.
5.16	Does the system have a "universal search" option, allowing a search to be undertaken over all modules of the system?	No. There is no single cross-module 'universal search'; each area provides its own filtering and search.	Confirmed.
5.17	Can the system store menu option 'favourites' on a per user basis?	No. The navigation menu is role- and feature-driven; per-user menu 'favourites' or pinning are not supported.	Confirmed.
5.18	Can a user open multiple windows accessing the same or different modules of the system?	Yes. As a browser-based application, users can open multiple browser tabs or windows onto the same or different areas of the system simultaneously, and the real-time backend keeps data consistent across them.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
5.19	Can more than one software function be performed concurrently?	Yes. The application is asynchronous and multi-user, so multiple functions can be performed concurrently across users and browser tabs, and scheduled background processes (rate-alert monitoring, hedge-policy evaluation, bank/ERP syncs) run alongside interactive use.	Confirmed.
User documentation and training			
5.20	Confirm whether a user manual / instructions is provided and how this is distributed?	Yes. An in-app Guide documents each area of the platform with usage guidance and can be exported to PDF. This is complemented by a step-by-step onboarding wizard and context-sensitive help. Distribution is in-app, with PDF export available. (Architecture and integration-specification documents also exist for technical audiences.)	Confirmed.
5.21	Does the user manual include: - An index or search facility? - A guide to basic functions of the software? - Pictures of screens and layouts? - Examples? - A tutorial section? - Details of any error messages and their meanings?	Partially. The in-app guide provides a per-module functional guide with practical 'when to use' examples and is text-searchable within the browser/PDF. It does not currently include annotated screenshots, a dedicated step-by-step tutorial section, or a glossary of error messages and their meanings.	Confirmed.
5.22	Is context-sensitive help available within the system?	Yes. Context-sensitive help is available throughout via inline help tooltips (which can be toggled on/off per user), alongside the onboarding wizard.	Confirmed.
5.23	Is the manual and/or help editable by the user (subject to the permissions matrix)?	No. Help and guide content is maintained by CrescitPay and is not editable by the customer.	Noted.
5.24	Will the Software House make the detailed program documentation (e.g. file definitions for third party links) available to the user, either directly or by deposit with a third party (ESCROW)?	Not beyond headline descriptions in line with data protection and privacy policies.	Noted.
5.25	Please detail the training options available?	The platform supports self-onboarding through its onboarding wizard, in-app guide and context-sensitive help. Every org-admin will go through comprehensive training with a member of CrescitPay staff upon commencement of service.	Noted.
5.26	Who provides training: - Software House? - VAR?	Training is provided by CrescitPay when a customer is onboarded.	Noted.
Support and maintenance			
5.27	How is the software sold: - Direct from the software house? - Via a Value Added Reseller (VAR) or Integrator?	CrescitPay operates a direct to customer sales engine as well as receiving introductions from commercial partners such as accountants who can share in revenue generated.	Noted.
5.28	How is the product supported: - Direct from the software house? - Via a Value Added Reseller (VAR) or Integrator?	Every customer will have a dedicated account manager who will be their point of contact for support throughout their journey.	Noted.
5.29	Do VARs have to go through an accreditation process?	N/A	N/A.
5.30	Is the software sold based upon number of named users or a number of concurrent users?	Neither. There is no subscription cost to being a customer as monetisation is realised through FX transactions undertaken via the platform. A small margin is applied by the liquidity provider such as Ebury who will share that revenue with CrescitPay.	Noted.
5.31	The supplier should detail the support cover options available, covering: - The hours provided? - Associated costs? - The global regions covered?	Support will be provided during business hours, and where capacity exists, beyond this. We are onboarding UK customers as a first wave and will conform to UK Business Hours - Mon-Fri 9am - 6pm excluding Bank Holidays.	Noted.

Ref	Requirement	Response	Reviewer Comments
5.32	Detail the process by which customers raise support requests and how these can be viewed/managed?	Customers can contact their account manager or email a generic help inbox or telephone line. Ticket system is on the roadmap for Series A. SLA response and resolution times are defined in policy.	Noted.
5.33	Please note the methods of support available: - Telephone? - Internet chat? - Remote access to customer workstation? - Other, please specify?	Telephone or Email support are available initially. These will all be recorded due to the financial nature of the product.	Noted.
5.34	Do you offer service credits for failure to meet performance around SLA and uptime (if applicable)	N/A	N/A.
5.35	What is your escalation path for tickets which have not been resolved within a reasonable time?	They can be raised via a complaint to the directors in writing via post or email.	Noted.
5.36	How often are general software enhancements provided?	As a single-version SaaS platform, enhancements are delivered on a continuous/rolling basis rather than in discrete annual releases.	Noted.
5.37	Will they be given free of charge?	There are currently no plans for any SaaS or update fee.	Noted.
5.38	How are enhancements and bug fixes provided to customers?	Enhancements and bug fixes are deployed centrally to all customers automatically (no customer-side install required), with zero-downtime deployment; per-organisation feature toggles allow phased enablement of new capabilities where appropriate.	Noted.
5.39	Is “hot line” support to assist with immediate problem solving available?	Yes	Noted.
5.40	If so, is there an additional cost involved?	No	Noted.
5.41	At what times will this support be available?	UK Business Hours 9am - 6pm Monday to Friday, excluding bank holidays.	Noted.
Integration and www facilities			
5.42	Are the different modules of the system fully integrated (i.e. no set-up effort required in order to use the various modules together)?	Yes. The modules are fully integrated out of the box: they share a single data model, so — for example — exposures flow into deal booking, deals flow into mark-to-market, hedge accounting and reporting, and payments link to settled obligations, with no inter-module set-up	Confirmed.
5.43	Are they integrated on real time basis or batch basis?	Real time. The modules share the live database and reflect changes immediately. Integrations with external systems (bank feeds, ERP/accounting) additionally run on schedules (e.g. daily syncs) as well as on demand.	Confirmed.
5.44	Can the software be linked to other packages e.g. word processing, graphics, financial modelling, to provide alternative display and reporting facilities?	Via export. Data can be exported to Excel, CSV and PDF for use in spreadsheets, word processing, graphics and financial-modelling tools. There is no direct embedded link that drives a Word or graphics package from within the application.	Confirmed.
5.45	Can definable links to spreadsheets be created?	Export to Excel and CSV is available across the application. Live, refreshable spreadsheet links (for example ODBC- or API-driven cells that update automatically) are not provided; exports are point-in-time files.	Confirmed.
5.46	Does the system provide a secure document storage capability: If so, please give examples of the document types saved and what transactions these might relate to.	No. The platform does not provide a document repository. It stores structured transaction data — including upstream source references (e.g. ERP invoice/PO numbers) where interfaced — rather than attached documents.	Noted.
5.47	Can documents be scanned into a secure repository?	No. Scanning documents into a secure repository is not supported.	Noted.

Ref	Requirement	Response	Reviewer Comments
5.48	Does the system provide data migration tools for transactional and master data sets (e.g. employees customers, suppliers, journals, invoices).	Yes, for the relevant data. A CSV deal-import tool (with a downloadable template and per-row validation) handles bulk/transactional loading, and automated connectors ingest exposures and transactions from ERP/accounting and open-banking systems. General master-data migration of unrelated entities (e.g. employees) is outside the scope of this FX/treasury platform.	Noted.
5.49	What connection mechanisms does the software have and what breadth of functionality in terms of: - operations (add, update, delete)? and - what transactions/data it can access? E.g. if webservices APIs available, then can customers connect to whatever software they wish?	Integration runs server-side through Cloud Functions using partner APIs and webhooks, with a set of pre-built connectors: open banking (TrueLayer, Plaid), ERP/accounting (Xero, Sage, NetSuite, Microsoft Dynamics 365, Coupa, Tipalti, Odoo, Unit4, Basware, Zip) and FX execution/settlement (Ebury). Depending on the connector these support read (sync) and write-back operations. There is not, at present, an open public API allowing customers to connect arbitrary third-party software of their own choosing; integration is via the supported connectors and webhooks.	Confirmed.
5.50	Does the system support mobile working?	Yes. The application is a responsive web application that can be used in tablet and mobile browsers as well as on desktop/laptop. There is no separate native mobile app; access is through the web browser.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.	<u>SaaS HOSTING</u>		
	This evaluation covers the system but not the method by which it is delivered and/or contracted for. Potential users need to satisfy themselves on the security and disaster recovery aspects and licensing of the online system and any data protection issues of their own and customer/supplier information, contained therein, being held on the system, as well as the return of the data when the contract expires or is terminated.		
Data centres and customer data			
6.01	Whose data centres are used and where are these located: - If hosted -- where data centre controlled by a third-party? - If SaaS -- where the software vendor will be in control?	CrescitPay is a SaaS platform hosted on Google Cloud Platform (Firebase). The physical data centres are owned and operated by Google Cloud; CrescitPay controls the application layer. The Firestore database and Cloud Functions are provisioned in Google's europe-west2 (London, UK) region, and the web application is served via Firebase Hosting's managed CDN. Note that FX execution, payment settlement and any holding of client funds are performed by a separate, regulated embedded-payments provider — CrescitPay is the workflow/analytics layer and does not hold client money.	Noted.
6.02	Does the customer get a choice of the jurisdiction in which their data resides?	The current deployment is single-region in europe-west2 (London), so all customer data resides in the UK/EU. A per-customer choice of hosting jurisdiction is not offered in the present configuration; alternative regions could be discussed commercially.	Noted.
6.03	What certification(s) do you or your platform operators hold relating to your data centres and your business operations?	The underlying hosting platform (Google Cloud) holds an extensive set of certifications, including ISO/IEC 27001, 27017 and 27018, SOC 1/2/3 and PCI DSS. We plan to seek ISO27001 certification for the platform itself and work towards the NIST AI RMF as part of our Series-A push in 2028.	Noted.
6.04	Do you or your platform operator have an SSAE16 (System and Organisation Controls) report available?	Google Cloud publishes SOC 1/2/3 reports covering the underlying infrastructure, available under NDA via Google.	Noted.
6.05	What are the physical controls over the:- - Premises? - Fileservers? - Communications equipment?	Physical security of premises, servers and communications equipment is provided by Google Cloud's data centres (24/7 staffed security, multi-factor biometric access, video surveillance, redundant power/cooling, etc.), as documented in Google's data-centre security materials. CrescitPay operates no physical infrastructure of its own; these controls are inherited from the cloud provider.	Noted.
6.06	Is the space in this/these data centre(s) shared with any other companies?	Yes — this is a public-cloud (Google Cloud) multi-tenant environment, so physical infrastructure is shared. Customer data is kept logically isolated as described in 6.07.	Noted.
6.07	Is data for different customers/companies kept:- - On separate servers? - In different databases? - In separate database tables? - In a database with data for other customers and companies using logical security to partition customers' data?	Data for different customers is held in a shared, managed database (Google Cloud Firestore) and partitioned by logical security: every record carries an organisation identifier (orgId) and access is enforced per organisation by server-side security rules. This corresponds to the last option — a shared database with logical security partitioning.	Noted.
6.08	How is it ensured that data for different customers and companies is reliably identifiable and only accessed by authorised users for each customer/company?	Every document is tagged with its owning organisation (orgId), and the Firestore security rules permit access only where the authenticated user is a member of that organisation, with a default-deny rule blocking anything not explicitly allowed. This ensures each customer's data is reliably identifiable and reachable only by that customer's authorised users.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.09	What controls are in place to prevent users from one customer/company accessing data from another customer/company by accident or by design?	Cross-customer access is prevented at the server-side rules layer: access requires matching organisation membership, the orgId on records is immutable once written, and a catch-all deny rule blocks all other access. These checks are enforced by the database regardless of the client, so isolation cannot be bypassed from the browser (see also Section 3).	Noted.
6.10	How is (Internet) communication traffic monitored to identify potential problems before they happen: - From a performance perspective? - From a security standpoint?	Network-level performance and security monitoring (including DDoS protection) is provided by Google Cloud. The application emits structured logs to Google Cloud Logging/Monitoring.	Noted.
6.11	What procedures are in place to prevent a break in Internet Connection (at the server, client or in between) from causing data corruption?	Database writes are atomic — multi-record operations use transactions/batches that commit all-or-nothing, so a dropped connection at the client, server or in between causes the operation to fail cleanly with no partial commit or corruption; the client can safely retry. Firestore's managed replication further protects durability.	Noted.
6.12	Are communications between the user's computer and the software service encrypted: - User log in data only? - All data exchanged between user client and software service?	All data exchanged between the user's browser and the service is encrypted in transit via HTTPS/TLS (Firebase Hosting managed SSL; Firestore and Cloud Functions use TLS 1.2+), not merely the login exchange.	Noted.
6.13	Is data on your servers encrypted at rest?	Yes. Google Cloud Firestore encrypts all stored data at rest by default.	Noted.
6.14	What level of encryption is used?	Encryption at rest uses Google-managed AES-256; data in transit uses TLS 1.2+.	Noted.
6.15	Is a staging environment provided that is an exact replica of production; which can be used for testing purposes?	CrescitPay operates a staging environment for testing purposes which does not contain client data.	Noted.
6.16	Is a test environment provided to test configuration changes? If so, is there an additional charge for this?	Not to clients.	Noted.
Access to customer data			
6.17	What are the implications of the Data Protection Act over information held by the hosting service provider, and how does the vendor mitigate these?	As a data processor, CrescitPay is subject to UK GDPR / the Data Protection Act. Mitigations evidenced in the platform include UK/EU data residency, authenticated role-based access with strict per-organisation isolation, encryption in transit and at rest, and an immutable audit trail.	Noted.
6.18	Are you subject to any legal or regulatory requirements obliging you to retain a copy of customer data?	There is a minimum period of seven years post-transaction for records to be kept on Forward contracts.	Noted.
6.19	Who will be able to access or see customer data?	End users can see only their own organisation's data, governed by role-based access. Google (as infrastructure provider) has no application-level access to customer data. CrescitPay support and operations staff will be able to see customer data when troubleshooting.	Noted.
6.20	Explain the procedures to prevent unauthorised access from staff, or contractors, working for the service provider or any other people with access to the service provider's internal systems.	Application secrets are held in Google Secret Manager (never in the browser bundle) and infrastructure access is governed by Google Cloud IAM.	Noted.
6.21	Explain the release management procedures in place and the associated segregation of duties ?	The release management policy includes a defined deploy sequence (automated tests → build → deploy). These will all be authorised by the CTO or CEO before deployment.	Noted.
6.22	Is there sufficient segregation of duties preventing system developers from accessing and changing live applications and data files?	Segregation between developers and live data/applications is enforced through Google Cloud IAM role separation.	Noted.
6.23	Explain the review and approval procedures covering system operations staff when emergency changes need to be made to live applications and data?	Customer data cannot be changed by anyone in operations staff. They can view transaction history but cannot edit or delete. Emergency changes can only be made by the CTO, or in their absence, the CEO. Such changes would be limited to suspending functionality in an emergency, no new functionality would be deployed without completing rigorous testing.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.24	Is an audit trail always maintained of these emergency changes?	The application maintains an immutable, append-only audit log of business events, and Google Cloud Audit Logs record infrastructure/deployment activity.	Noted.
6.25	What procedures are in place when members of staff leave to ensure that their system access is stopped?	When a staff member leaves, their authenticatoin account will be disabled, and their access to Google Cloud systems revoked immediately.	Noted.
Platform and service levels			
6.26	Which databases and servers are used to host the software?	The software is hosted on Google Cloud: Cloud Firestore (a managed NoSQL document database) for data, Cloud Functions (Node.js, serverless) for server-side logic and integrations, and Firebase Hosting for the web application — all in the europe-west2 (London) region.	Noted.
6.27	What forms of user authentication are supported e.g. user names, passwords certificates, tokens etc.?	Supported authentication methods are: email address and password; Google sign-in; and enterprise single sign-on via OIDC (Microsoft Entra ID / Azure AD and Okta). Sessions are carried by short-lived Firebase ID tokens (signed JWTs). Multi-factor authentication is supported by the underlying platform (see 3.09).	Confirmed.
6.28	What is the proposed product/service availability percentage?	99.9% uptime is the target SLA.	Noted.
6.29	What percentage availability has been achieved over the past 12 months?	Platform only went online in May. Uptime has been 100% since then. In the event of a downtime of five minutes or more, an email is sent to CEO/CTO inboxes. Customers do not have access to this data.	Noted.
6.30	Is a service level agreement ("SLA") offered regarding: - Service availability? - Data recovery?	99.9% service availability, and data recovery should be 100% upon completion of any entry as the program writes atomically. Execution functionality is only avaialable Mon-Fri 0630-2200, however all other functions are available 24/7.	Noted.
6.31	Is the service available 24x7 or are there downtime periods for maintenance?	The service is designed for continuous (24x7) availability on Google's managed, always-on infrastructure; application deployments are zero-downtime. FX transactions can only be undertaken when markets are trading, which is a 24x5 window.	Confirmed.
6.32	Is the customer made aware of maintenance periods in advance?	1 week notice will be given of any notice periods, however we do not expect downtime to be associated with these, the customer would have to simply log out and log back in again.	Noted.
6.33	Does the application software:- - Require any client software to be installed on the user's computer? - Work entirely within Internet Browser software on the user's computer?	The application works entirely within a standard web browser — it is a browser-based single-page application served over the web. No client software needs to be installed on the user's computer.	Confirmed.
6.34	Where the product/service relies upon downloading and running an executable program, has that program been secured with a digital certificate to verify the source and integrity of the program?	N/A	N/A.
6.35	Does the product/service currently use any technologies which are obsolescent / out of support / soon to be end of life? If so, describe how the user can mitigate this risk.	No	Noted.
Platform security			
6.36	What security steps are taken to prevent and detect intrusion attempts?	Intrusion prevention/detection is layered: Google Cloud provides network-level protections including DDoS mitigation; access requires authentication (Firebase Auth) and passes server-side authorisation rules with a default-deny posture; secrets are held in Secret Manager; and all traffic is over HTTPS/TLS.	Noted.
6.37	Is firewall hardware and software used to protect the live systems from unauthorised access?	The architecture is serverless on Google Cloud, so there are no customer-managed servers or firewalls; network perimeter protection is provided and managed by Google Cloud.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.38	Which monitoring software is used to create alerts when intrusion attempts are suspected?	Nothing is currently set up beyond systems implemented by Google Cloud. Additional monitoring and reporting are on the Series A roadmap, 2028 estimate. Such alerts would be received by the CTO/CEO as per the downtime ones.	Noted.
6.39	Are designated staff responsible for receiving and urgently responding to these alerts?	Yes. Under the Customer Support Standards Policy, a security incident affecting client data is classified as a Severity 1 issue, which carries a 30-minute response target (available 24/7 during extended support hours) and a 4-hour resolution target. The CTO is notified immediately on any Severity 1 issue, and the CEO is additionally notified where the issue affects client data, multiple clients, or carries potential regulatory implications — which a security incident of this kind would. The affected client receives proactive updates at least every 60 minutes until resolution.	Noted.
6.40	Have clear procedures been established for identifying and responding to security incidents?	Yes. The Information Security Policy sets out a five-stage incident response process — Detect, Contain, Investigate, Notify, and Recover & Review. All personnel are required to report a suspected security incident to the CTO immediately, regardless of time of day. Critical and High severity incidents are escalated to the CEO immediately, who is responsible for external communication (including with execution partners and, where applicable, regulators). Notification timelines are defined — including the ICO within 72 hours for personal data breaches under UK GDPR, and execution partners typically within 24 hours — and the process is tested at least annually through a tabletop exercise. The complementary escalation and client-communication path once an incident is live is set out in the Customer Support Standards Policy (see 6.39).	Noted.
6.41	Is all security sensitive software, such as operating systems and databases, kept up to date with the latest software patches? Please indicate how regularly updates are applied.	The managed platform components (Firestore, Cloud Functions runtime, Hosting) are patched automatically by Google. Application dependencies are maintained by theCrescitPay with all recommended updates installed immediately upon release.	Noted.
6.42	List the procedures and software tools in place to prevent or detect and eliminate interference from malicious code, such as viruses?	The platform is a serverless managed service with no customer-operated operating systems to run anti-virus on, and the application does not provide a file-upload/execution surface; infrastructure hardening against malicious code is handled by Google Cloud.	Noted.
6.43	Is a system log maintained by the service provider that details - User access? - User activity? - Error messages? - Security violations?	Collectively yes: the application audit log records user activity and key changes; Google Cloud Logging records function/error activity; Firebase Authentication records user access; and denied access attempts are logged at the rules/platform layer.	Noted.
6.44	Is this log available to the customer?	The in-application audit log is available to organisation administrators/members and can be exported to Excel. Underlying infrastructure logs (Google Cloud Logging) are not exposed directly to customers.	Noted.
6.45	Have there been any successful unauthorised access attempts been made during the last year? If Yes:- - What was the effect on the business and users? - What steps are in place to prevent this happening again?	No. Regarding unauthorised attempts - Failed authentication attempts are recorded by Firebase Authentication, and access-rule denials (attempts to reach data or actions outside a user's authority) are logged at the Firestore/platform layer in Google Cloud Logging. These platform-level logs provide a server-side trail of failed and denied access attempts, distinct from the in-application business audit log.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.46	Is penetration testing regularly carried out by (please indicate frequency of tests): - Staff specialising in this field? - External specialists?	We run OWASP ZAP locally and on-demand to ensure all endpoints are secure and check all critical vulnerability points. These tests will run quarterly and prior to release of any major update.	Noted.
6.47	Are procedures in place to ensure that any weaknesses found by penetration testing are addressed quickly?	Yes - any Pen test results with adverse findings will be immediately rectified.	Noted.
6.48	If penetration testing by a specialist is not performed regularly, please indicate the main procedures in place to identify weaknesses?	N/A	OWASP ZAP is an automated vulnerability scanning tool rather than penetration testing performed by a security specialist.
6.49	Are security procedures regularly reviewed? Please indicate frequency of reviews.	Reviewed quarterly with the CTO reporting findings and recommendations to the board.	Noted.
6.50	What security reporting is provided demonstrating compliance against certification(s) and policy(ies)?	We are working towards the NIST framework for continuous improvement.	Noted.
6.51	How are security breaches communicated to customers?	These would be confirmed via email in line with the information security policy.	Noted.
Backups by the service provider			
6.52	In relation to backups undertaken by the system provider please explain: - How is a customer's data backed up? - How often is this undertaken? - What is backed up? - What's the media used? - Where are backups stored? - How many copies are there? - How long are they retained for? - Who has access to them? - Is the data encrypted?	Data is stored in Google Cloud Firestore, which provides durable, automatically-replicated managed storage with encryption at rest, and supports scheduled backups and point-in-time recovery.	Noted.
6.53	How frequently is a test-restore of backups undertaken?	Quarterly as part of the security review.	Noted.
6.54	Can the provider restore from a backups that it has taken at a customer request?	In principle yes — Firestore backups, point-in-time recovery and data exports allow the provider to restore data.	Noted.
6.55	Does a customer have the ability to undertake their own backups?	Customers can export their own data from across the application (Excel/CSV/PDF), giving them an independent copy. This is a data export rather than a full database backup.	Confirmed.
6.56	If so, can a customer restore data a backup that they have taken?	Exported data serves the customer's own records and portability. A full self-service restore is not offered; re-import is limited to the structured CSV deal importer (see Section 4), not a complete data restore. They cannot restore their system of record on Crescitpay to an earlier point in time.	Noted.
Platform recovery			
6.57	What contingency plans are in place to enable a quick recovery from: - Database or application software corruption? - Hardware failure or theft? - Fire, flood and other disasters? - Communication failures?	Resilience is underpinned by Google Cloud's managed, multi-zone services: Firestore data is automatically replicated; serverless compute has no single hardware dependency to fail or be stolen; and Google's regional infrastructure provides protection against localised disasters. Atomic writes prevent application/DB corruption from partial operations.	Noted.
6.58	How often are these plans tested?	We will test quarterly that we can resurrect control of our systems from a standing start.	Noted.
6.59	How often are these plans reviewed and updated?	Reviewed following each quarterly test.	Noted.
6.60	What is the longest period of time envisaged that service may not be available?	Our partners, GCP, Firebase, Visa and Ebury all have a 99.9% uptime SLA. Whilst outages caused by third parties are inevitable, we expect them to last 3-6 hours at most for a catastrophic failure.	Noted.
6.61	What are your: - Recovery Point Object (RPO) standards? - Recovery Time Objective (RTO) minimum standards?	Firestore's point-in-time recovery supports a low Recovery Point Objective (to a recent point, within the retention window). Atomic writes should all be recoverable with zero loss. RTO to minimum standards is 24 hours.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.62	If transaction records are dated and time stamped are the times used local to the user or based on where the server is located?	Transaction records are date/time-stamped using server-side UTC timestamps (Firestore server timestamps), and are displayed in the user's local time in the interface. Storage is therefore server/UTC-based, not dependent on the user's local clock.	Confirmed.
6.63	What protection is in place to enable users to able to access their accounting and other data if the service provider should experience serious difficulties, cease trading or decide to stop providing the service?	Customers can export their accounting and other data from the application at any time (Excel/CSV/PDF), giving an independent copy. CrescitPay will continue to build liquidity partnerships so other providers can take over relationships if needed.	Confirmed.
6.64	Do these arrangements include: - Standby arrangements for another organisation to continue providing the full service? - Minimal arrangements to at least enable customers to access their data for a sufficient period of time to extract data copies, produce reports and make alternative arrangements?	N/A	N/A.
6.65	If the system is hosted are there arrangements in place for this third party to continue providing a hosting service in the short term to allow time for customers to negotiate their own arrangements? If so, how long does the arrangement allow?	Hosting is provided by Google Cloud, so short-term hosting continuity is underpinned by Google's platform rather than a smaller third party.	Noted.
6.66	Are there any individual members of the vendor's staff whose leaving or illness would significantly reduce, or even stop, the service provider's ability to provide a full and reliable service to customers?	CrescitPay has a small team, so any individual suffering incapacitation would be inconvenient, but not a point of total failure. Both the CEO (Matt) and CTO (Bav) have full access to all systems, and detailed recovery instructions exist should the CFO (Peter), or a third party, need them.	Noted.
Platform change management			
6.67	Describe your approach to upgrades including what option customers have not to take upgrades (if any)?	CrescitPay is a single-version SaaS platform: all customers run the current release, so there is generally no option to decline an upgrade.	Noted.
6.68	Are users able to test the application before new versions go into live use?	No. Testing is not a customer facing function. Customers can toggle on/off on any feature in the platform.	Noted.
6.69	Are users given notice before application changes are applied to the live system?	No, but they are able to toggle on/off new features and can download all their records to CSV at any time.	Noted.
6.70	Are changes delivered into the live environment "switched off" to enable users to test them before enabling them for their environment?	Yes — the platform includes a per-organisation feature-visibility mechanism (enabledFeatures), so new capabilities can be delivered but kept switched off for an organisation until enabled.	Confirmed.
6.71	Describe what testing and QA processes are undertaken before upgrades and other changes are made live/available to customers?	The repository includes an automated test suite and a defined pre-deploy checklist (run tests → build → deploy). There are no circumstances where a failed test would still lead to deployment.	Noted.
6.72	Explain the release management procedures in place and the associated segregation of duties?	Releases are expected to feature bundled rollouts of features and enhancements rather than with a specific regular cadence.	Noted.
6.73	Are users informed when they next login of the application changes that have gone into live use?	No, but all updates will be accompanied by an email to all users detailing any customer-facing changes.	Noted.
6.74	Do customer staff have to take any action (e.g. regression testing) when new editions, patches or upgrades are released?	As a single-version SaaS platform, upgrades are applied centrally and typically require no action from customer staff.	Confirmed.
Subscription options			
6.75	What is the minimum level of commitment must the customer sign up to, e.g. 36 months?	N/A. It's a no-cost model without a subscription contract.	N/A.
6.76	Where online payment is used, what type of security is used to protect sensitive information?	N/A	N/A.
6.77	Where online subscription / payment is used, is an invoice provided to the customer and, if so, in what format?	N/A	N/A.
6.78	When subscriptions need to be renewed, what advance notice is provided and what is the time limit for renewal?	N/A	N/A.
6.79	Is there a procedure for late renewal and is there a time limit after which subscriptions cannot be renewed?	N/A	N/A.
6.80	How soon after creating or renewing a subscription (if applicable) can the system / service be used?	N/A	N/A.

Ref	Requirement	Response	Reviewer Comments
6.81	What notifications / confirmations are provided to the customer regarding subscriptions and payments?	N/A	N/A.
6.82	To what extent are users able to access their accounting and other data if: - They miss one or two payments? - They cease being customers?	They have total access until they ask for their account to be deleted.	Noted.
6.83	At the end of the contract term, how long does a customer have to obtain a copy of their data from you?	We are obliged not to hold data for longer than is necessary, however some forward contracts require record keeping for seven years post-transaction. Records should be available through this time.	Noted.
6.84	At the end of the contract term, how is a customer's data destroyed (if appropriate) and will that destruction be certified?	Deletion of the user profile within our firebase repository should effectively delete all data associated with that customer	Noted.
6.85	What is your processes regarding disposal of end-of-life and failed hardware devices that were used to operate your service?	The service operates no physical hardware of its own; end-of-life media disposal is handled by Google Cloud under its documented secure data-destruction and media-sanitisation standards.	Confirmed.
SaaS/Hosted reporting			
6.86	Are reports produced from the same software as the financial applications or is separate reporting software used?	Reports are produced by the same application as the financial functionality — reporting (board report, disclosure reports, analytics, and data-grid exports) is built into the platform. No separate reporting product is used.	Confirmed.
6.87	Does any application software (i.e. other than a web browser or PDF reader) need to be installed on the user's computer in order to prepare or view the reports?	No. Reports are viewed in the browser or produced as PDF/Excel/CSV downloads; the user needs only a web browser (and a standard PDF reader or spreadsheet application to open those respective formats).	Confirmed.
6.88	What browser versions are support: - On desktop/laptop (PC, Mac, Linux)? - On Tablets? - On mobiles?	The application targets modern evergreen browsers — current versions of Chrome, Edge, Firefox and Safari — on desktop/laptop (Windows, macOS, Linux). The interface is responsive (a mobile viewport is configured) for use on tablets and mobiles.	Noted.
6.89	Is access to the reporting facilities and data controlled by the same procedures as access to the main application?	Yes. Reporting uses exactly the same authentication and role-based access controls, and the same per-organisation data isolation, as the rest of the application.	Confirmed.
6.90	If it's different, explain the user access control facilities available to ensure information is only viewed by users with appropriate authority?	Not applicable — access to reporting is governed by the same controls as the main application.	N/A.
6.91	In what electronic formats are reports produced:- - PDF? - XML? - MS Excel spreadsheet? - CSV file? - As html for viewing in a web browser? - Other, please specify?	Reports are produced as PDF, Microsoft Excel (XLSX) and CSV, and can be viewed as HTML in the browser. XML is not currently a standard output format.	Confirmed.
6.92	Are report documents stored on the web server or on the user's computer? If stored on the web server, are they secure to ensure only users with appropriate authority can get access?	Reports are generated on demand and delivered to the user (as a download) or rendered in the browser; the underlying data is held in Firestore and protected by the same authentication and authorisation rules. Report artifacts are not persisted in a shared server-side store, and access to the source data always requires appropriate authority.	Confirmed.
6.93	For documents viewable in a browser is any data stored on the user's computer in a web browser cache or temporary file? If Yes: - Is there any protection against other users viewing the report or data on which it is based? - Is it clear on the reports when they were produced and the date of the data on which they are based, so the user can tell whether they are viewing out of date information?	Reports are generated dynamically from live data. Any downloaded report file resides on the user's device under their control, subject to standard browser handling; protection of files on a shared device is the customer's responsibility. Generated reports are dated (for example the board report records its report date) so users can see the currency of the data.	Confirmed.
6.94	Are communications between the browser and the server encrypted for any report related communications?	Yes. All communication between the browser and the server, including report-related data, is encrypted over HTTPS/TLS.	Noted.

Ref	Requirement	Response	Reviewer Comments
6.95	If reports are produced dynamically each time the user views them can historical reports be reproduced at any time?	Reports are generated dynamically from stored data, so a report for a previous period can be regenerated at any time from the retained records (subject to the underlying data	Confirmed.
6.96	Can reports viewable in a browser be navigated dynamically by users? For example: - Enabling drill down to more detailed information? - Altering which columns and rows of data are displayed. - Choosing time periods? - Specifying selection criteria?	Interactive reports and dashboards support dynamic navigation in several areas — drill-down from summary figures to the underlying deals and exposures, filtering by selection criteria, and choice of time periods. Full pivot-style reconfiguration of arbitrary columns/rows is limited; data can be exported to Excel/CSV for further manipulation.	Confirmed.
6.97	Can report data be reliably copied and pasted direct from browser viewable reports to an MS Excel spreadsheet retaining any table layout?	Report data can be exported directly to Excel (XLSX) and CSV retaining the tabular layout, which is provided in preference to manual copy-and-paste.	Noted.
6.98	If reports are incomplete, for instance due to a poor Internet connection, is sufficient information provided to enable the user to notice that some of the report is missing?	Reports and exports are generated from the complete underlying dataset, and screens show loading indicators while data is retrieved rather than rendering partial data, which avoids silently incomplete reports.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
7.	FX Hedging & Payments		
Global setup and organisation configuration			
7.01	What organisational structures are supported (e.g. entities, subsidiaries, business units or groups)?	The platform supports multi-organisational structures. A corporate group can be onboarded in full — the parent (group) company together with all of its subsidiaries — or, alternatively, a single subsidiary can be onboarded on its own, provided the individuals responsible have the authority to execute FX transactions on behalf of that subsidiary. Each organisation is a separately isolated tenant, enforced server-side by organisation ID, so the data, users and configuration of each entity remain segregated.	Confirmed.
7.02	Does the platform support centralised, decentralised and hybrid treasury operating models?	Yes. The platform accommodates centralised, decentralised and hybrid treasury operating models. In a centralised model, the group and its subsidiaries are onboarded together and administered centrally — a central treasury team holding administrator and approver roles across the group's entities. In a decentralised model, each subsidiary is configured as an independent organisation with its own workflows, permissions and policies, and a single subsidiary can also be onboarded entirely on its own. Hybrid arrangements combine the two: central oversight and consistent policy (applied at onboarding and copied from a common template where required) alongside independent per-entity execution. The balance between central and local control is determined by how administrator and approver roles are assigned across the entities.	Noted.
7.03	Can organisations configure key platform settings such as reporting currency, approval controls and available functionality?	Yes. Organisations configure key platform settings — including the reporting/base currency, the functional currencies in which they transact, and approval controls (including value-based approval limits). These settings are managed by the organisation administrator.	Confirmed.
7.04	Can individual entities or business units within the same group have independent configurations, workflows and user permissions?	Yes. Each entity within a group is configured independently and maintains its own configuration, approval workflows and user permissions. No shared configuration is imposed across entities — each is administered as a separate organisation.	Confirmed.
7.05	Can configuration settings be inherited from a parent organisation while allowing entity-specific overrides where required?	Configuration is not automatically inherited from a parent organisation; each subsidiary is set up independently. However, an existing entity's configuration can be used as a template and copied to a new subsidiary as a starting point, which is then tailored to that entity's requirements. This provides consistency at onboarding without enforcing ongoing inheritance.	Noted.
7.06	Can new entities, business units or regions be added without significant system reconfiguration?	Yes. New entities, business units or regions can be added without significant system reconfiguration. The principal prerequisite is confirming that our execution partner is able to support the new entity or region (for example, the relevant currency corridors and the entity's jurisdiction); once confirmed, the entity is provisioned through standard onboarding.	Noted.

Ref	Requirement	Response	Reviewer Comments
User management & access control			
7.07	Which authentication methods are supported (e.g. Single Sign-On (SSO), Multi-Factor Authentication (MFA), Microsoft Entra ID, OAuth, SAML)? (See also 3.08)	Enterprise Single Sign-On is supported via the customer's identity provider, including Microsoft Entra ID and Okta; email/password authentication is also supported. Where SSO is configured, multi-factor authentication is enforced by the customer's identity provider. Native in-application MFA enrolment/enforcement for email/password users is a planned enhancement (see questionnaire ref 3.09).	Confirmed.
7.08	Does the platform support role-based access control (RBAC)? Please describe the standard and custom roles available, and whether permissions can be restricted at a granular level, including by transaction type, approval limit/materiality threshold, entity, business unit, or other transaction attributes.	Yes. The platform uses role-based access control with five standard roles — Admin (org_admin), Approver, Trader, Viewer and Auditor. The functional access associated with each role is fixed: two users holding the same role within an organisation cannot be given different levels of functional access. Granularity is instead provided in two ways: (i) approval authority can be varied per user by value/materiality threshold — the organisation administrator sets the approval limit that applies to each user, so users in the same role may carry different approval ceilings; and (ii) because each entity or business unit is a separate organisation, permissions are inherently segregated by entity. Custom roles beyond the five standard roles are not defined.	Confirmed.
7.09	Can permissions be configured at user and/or role level?	Permissions are configured at role level and applied per user through their role assignment; in addition, each user's approval value limit can be set individually within their role by the organisation administrator. Functional access is therefore determined by role, while approval authority is refined per user by value threshold.	Confirmed.
7.10	Does the platform support auditor or read-only access?	Yes. Both the Viewer (read-only) and Auditor roles provide non-editing access to data.	Confirmed.
7.11	Does the system enforce segregation of duties between transaction entry, approval and administration activities?	Yes. Segregation of duties is enforced server-side. Transaction entry, approval and administration are separated by role, and a four-eyes rule requires that the approver of a deal is a different user from the arranger — a user cannot approve their own transaction. Role assignments and changes are recorded in the audit log.	Noted.
7.12	Can approval limits be configured based on user, role, and/or transaction value? Additionally, does the system maintain a comprehensive audit log of changes to permissions, user roles, approval limits, and other security-related configurations?	Yes. Approval requirements are configured by role and by transaction value, and approval limits can additionally be set per individual user by the organisation administrator, so users in the same role may carry different value limits. All changes to permissions, role assignments, approval limits and related security configuration are recorded in the immutable, append-only audit log, as is each approval event.	Noted.
AI narrative generation			
7.13	Does the platform use AI or Large Language Models (LLMs) to generate narrative commentary, explanations, recommendations or insights?	Yes. The platform uses large language model (LLM) technology to generate a digest summarising events and activity over the preceding 7 days, and uses machine learning to produce a 90-day cashflow forecast. These are the two AI-driven outputs currently provided.	Confirmed.
7.14	Which reports, analyses or outputs can be AI-generated?	Two outputs are AI-generated: (i) the weekly digest — an LLM-generated narrative summary of the organisation's activity and events over the previous 7 days; and (ii) the 90-day cashflow forecast, produced using machine learning.	Confirmed.
7.15	Which AI models or LLMs are supported, and can organisations select or configure the model used?	The platform uses large language models from Anthropic and OpenAI, each applied to the function for which it is best suited. The model used for each function is selected and managed by CrescitPay; it is not customer-configurable.	Noted.

Ref	Requirement	Response	Reviewer Comments
7.16	Does the platform use a single LLM interaction or a multi-step AI processing workflow (e.g. retrieval, validation, post-processing)?	The platform uses a multi-step AI processing workflow rather than a single model call. Relevant customer data is first retrieved and assembled; the model then generates the output; and validation safeguards are applied to the result (for example, checks to prevent double-counting — see 7.21) before it is presented to the user.	Confirmed.
7.17	What internal and external data sources can be used to generate AI-driven narratives and insights?	The AI-driven narratives and insights are generated solely from data provided by the customer. Outputs describe the current state of the customer's own organisation; the platform does not draw on external data sources to compare the customer against specific competitors or against generic market expectations.	Confirmed.
7.18	Can AI capabilities be enabled, disabled or restricted by organisation, business unit, role or individual user?	Yes. As with all platform features, the AI capabilities can be enabled or disabled by the organisation administrator, i.e. at organisation level. Because each entity is a separate organisation, this also provides entity-level control. The feature is not separately toggled per role or per individual user.	Confirmed.
7.19	Can users customise AI-generated outputs (e.g. audience, tone, level of detail or reporting period)?	At present the platform does not support customisation of AI-generated outputs; users cannot adjust parameters such as audience, tone, level of detail or reporting period. Outputs are generated to a standard format.	Noted.
7.20	Are AI-generated outputs supported by references, source data or explanations to help users validate the results?	Yes. AI-generated outputs are grounded in, and traceable to, the source data provided by the customer, which allows users to validate the results against their own records. The platform does not introduce external benchmarks or third-party reference data — it does not hold the customer's objectives against which such benchmarking would be meaningful — so outputs reflect only the customer's own data.	Noted.
7.21	Does the platform perform validation, reasonableness checks or confidence scoring of AI-generated outputs?	Yes. The platform applies validation safeguards to AI-generated outputs, including checks designed to prevent errors such as the double-counting of transactions or exposures, before results are presented to the user.	Noted.
7.22	Is an audit trail maintained for AI-generated content and subsequent user edits or approvals?	The immutable audit log records the platform features a user has interacted with, including the generation of AI outputs. AI-generated output cannot be edited from within the platform, so there are no in-platform user edits to the content; the 'subsequent user edits' contemplated by the question therefore do not arise.	Noted.
7.23	Is customer data used to generate AI outputs retained, used for model training or shared with third-party AI providers? Please describe the applicable privacy, security and governance controls.	Customer data used to generate AI outputs is not retained, and is not used for model training, by the third-party AI providers (Anthropic, OpenAI). Data is submitted to generate the relevant output and is not stored or reused by the provider beyond that purpose.	Confirmed.
Exposure management & forecasting			
7.24	Can the platform capture and manage forecast foreign currency exposures, including both payables and receivables?	Yes. The platform captures and forecasts both payables and receivables. Forecast exposures can be populated in three ways: automatically through integration with the customer's ERP system, by manual upload, or through pattern detection that identifies recurring flows from historical activity.	Confirmed.
7.25	Can exposure data be imported from external systems and manually maintained within the platform, and can users distinguish from exposures that were manually entered or imported externally?	Yes. Exposure data can be imported from external systems and also maintained manually within the platform, and users are able to distinguish between forecast and actual (confirmed) exposures.	Confirmed.
7.26	Does the platform support recurring and one-off forecast transactions?	Yes. The platform supports both recurring and one-off forecast transactions — modelling recurring and non-recurring flows is central to its exposure-management function.	Confirmed.
7.27	Can forecast exposures be categorised, weighted or prioritised based on certainty or probability?	Yes. Forecast exposures can be categorised, weighted and prioritised according to their certainty or probability, so that more certain exposures can be distinguished from more speculative ones when assessing hedging requirements.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
7.28	Does the platform provide visibility of hedged, partially hedged and unhedged exposures and, can reports be generated detailing hedging status?	Yes. The platform provides clear visibility of hedged, partially hedged and unhedged exposures. Reporting includes hedged items, unhedged items and value at risk, together with consolidated views of total exposure, hedged exposure, and upcoming transactions shown with or without a hedge in place.	Confirmed.
7.29	Does the platform support natural hedge analysis and net exposure calculations?	Yes. The platform supports natural-hedge analysis and net-exposure calculation, offsetting opposing currency flows to determine net positions. This analysis can be performed across group companies as well as at individual entity level.	Confirmed.
FX Rates, hedging & execution			
7.30	What FX rate sources are supported and how frequently are rates refreshed?	FX rates are provided by our execution partners — currently Visa and Ebury Partners — and refresh every second. When a transaction is being assembled, the customer is presented with an exchange rate and has 30 seconds to accept or decline it before it refreshes to the prevailing rate.	Confirmed.
7.31	Does the platform support spot, forward and other FX hedging instruments?	The platform currently supports spot and forward FX contracts. Other instrument types (for example options or swaps) are not offered at this time.	Noted.
7.32	Can users define and manage hedging policies within the platform?	Yes. Hedging policies are defined and managed within the platform by the organisation administrator, who configures the policy parameters that govern the organisation's hedging activity.	Confirmed.
7.33	Are changes to hedging policies maintained in a system audit log?	Yes. All changes to hedging policies are recorded in the system's immutable, append-only audit log, providing a complete and tamper-evident history of policy amendments.	Confirmed.
7.34	Can the platform generate alerts or automated actions based on market conditions, exposures or policy rules?	Yes. The platform can generate alerts and initiate automated actions in response to market conditions, exposures or policy rules, operating within the rules configured by the organisation administrator and subject to the applicable approval limits.	Confirmed.
7.35	Does the platform support automated or semi-automated execution of FX transactions?	Yes. The platform supports automated and semi-automated execution of FX transactions, within the policy rules and approval limits configured for the organisation.	Confirmed.
7.36	Can automated or semi-automated trade executions be restricted by user, currency or materiality?	Yes. Automated and semi-automated execution can be restricted by user, by currency and by materiality, in each case subject to the approval limits set by the organisation administrator.	Confirmed.
7.37	What controls exist to prevent unintended execution of trades or payments?	A number of controls prevent unintended execution. Trades and payments are subject to four-eyes approval (the approver must be a different user from the arranger) and to configurable value/approval limits. When a transaction is assembled, the customer must actively accept the presented exchange rate within a 30-second window before it refreshes, preventing execution on a stale or unintended rate. The platform also enforces permitted status transitions server-side and provides a review step before commitment, so a transaction cannot progress to execution without deliberate, authorised action.	Confirmed.
7.38	What controls prevent automated execution of an incorrect, unauthorised hedge or a hedge that is deemed out of policy rules?	The system actively blocks a hedge that would breach the organisation's policy rules — such a transaction is prevented rather than merely flagged. This is enforced in addition to four-eyes approval and the configured approval limits, so an out-of-policy or unauthorised hedge cannot be executed.	Noted.
Payments & treasury operations			
7.39	Can the platform initiate and manage international payments and settlements?	Yes. The platform initiates and manages international payments and settlements through its regulated execution and payment partners, currently Visa and Ebury Partners.	Confirmed.
7.40	Does the platform support beneficiary management and validation controls?	Yes. Beneficiary management is provided; as it concerns the regulated payment activity, it sits within the domain of, and is performed by, the execution partner.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
7.41	Does the platform perform beneficiary screening or require validation before authorised in the system?	Beneficiary screening — including sanctions and anti-financial-crime checks — sits within the domain of the FCA-authorised execution partner and is performed by them as part of the regulated payment process, before a payment is authorised.	Confirmed.
7.42	Does the platform support maker-checker approval workflows for payments and trades?	Yes. Maker-checker (four-eyes) approval applies to both trades and payments — the approver must be a different user from the initiator.	Confirmed.
7.43	What happens to trade execution if the payments provider is unavailable?	Both of our current execution partners hold e-money licences. As a result, if a payment cannot be initiated immediately — for example, if the payments provider is temporarily unavailable — the trade is still completed, and the resulting funds are held in a holding account under a virtual IBAN in the name of the customer until the payment can be initiated. The customer's completed trade and funds are therefore protected, and settlement proceeds once the payment route is available.	Noted.
7.44	Can users track transaction status throughout the processing/payment lifecycle?	Yes. Users can track the status of a transaction throughout its processing and payment lifecycle — from assembly and approval through execution and settlement — with the current status visible at each stage.	Confirmed.
7.45	In the event of trade execution or payment processing failure, what controls are in place to notify users, prevent duplicate processing and protect against loss of funds?	In the event of a payment failure, the customer is notified automatically by email. The notification is triggered as soon as the execution partner's systems identify an issue with the payment, so the customer is alerted promptly and can take corrective action. Fingerprint-based de-duplication additionally guards against duplicate transactions.	Noted.
Risk analysis and scenario modelling			
7.46	Does the platform provide dashboards showing current exposure, hedge coverage and risk positions?	Yes. A headline dashboard is presented on the home screen, showing current exposure, hedge coverage and risk positions at a glance, with the ability to view the underlying detail.	Confirmed.
7.47	Does the platform support scenario modelling and what-if analysis?	Yes. The platform supports scenario modelling and what-if analysis, allowing users to model the effect of different assumptions and potential market movements on their exposures and hedge positions.	Confirmed.
7.48	Does the platform provide stress testing using historical, hypothetical or statistical scenarios?	Yes. The platform provides stress testing against both historical and hypothetical scenarios. Six predefined historical scenarios are included: the Global Financial Crisis (September–October 2008), the 'Taper Tantrum' (May–June 2013), the Brexit referendum (24 June 2016), the COVID-19 market shock (March 2020), the Russia–Ukraine conflict (February–March 2022), and the 2022 US dollar surge (January–October 2022). Each applies the observed currency movements of that episode to the customer's own exposures. Users can also stress test against defined hypothetical movements, such as a 5%, 10% or 20% shift in rates.	Confirmed.
7.49	Can the platform quantify the potential financial impact of FX market movements?	Yes. The platform quantifies the potential financial impact of FX market movements, translating rate movements and stress scenarios into their estimated monetary effect on the customer's exposures and positions.	Confirmed.
7.50	Does the platform provide counterparty, concentration or exposure reporting?	Yes. The platform provides counterparty, concentration and exposure reporting, enabling users to see how exposure is distributed and to identify concentrations of risk.	Confirmed.
7.51	Can the platform limit how much FX exposure, cash or derivatives it has with any one bank, broker or trading counterparty?	Yes. The platform can limit how much FX exposure, cash or derivative position is held with any single bank, broker or trading counterparty. This limit functionality is configured by CrescitPay on request from the customer.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
Hedge accounting, compliance and documentation			
7.52	Which hedge accounting standards and regulatory frameworks are supported?	The platform supports hedge documentation and compliance reporting aligned with IFRS 9 and FRS 102, and EMIR transaction reporting. These are the principal standards and frameworks supported.	Confirmed.
7.53	If the system supports multiple standards, can these be set as different for separate business units or entities?	Yes. Because each entity is configured independently as a separate organisation, different entities or business units within a group can apply different hedge-accounting standards where required.	Confirmed.
7.54	Does the platform support hedge designation, effectiveness testing and documentation management?	Yes to all three. Hedge designation is fully supported: a formal designation record links the hedging instrument to one or more hedged items and captures the documentation required at inception — risk management objective and strategy, hedge type (cash flow, fair value or net investment), hedged risk, effectiveness testing method, hedge ratio, the forward-points-exclusion election, and a snapshot of the instrument's terms at designation. Designations follow a controlled lifecycle (designated, effective, ineffective, discontinued, matured). Effectiveness testing is performed within the platform, supporting both the critical-terms and dollar-offset methods, with results recorded against the designation. Documentation management is supported: designation documentation and supporting records are maintained in the platform and can be exported for audit.	Noted.
7.55	Can hedge accounting and compliance documentation be generated and exported for audit purposes?	Yes. Hedge-accounting and compliance documentation can be generated and exported (PDF, Excel and CSV) for audit purposes.	Confirmed.
7.56	Can historical transactions from external counterparties be incorporated into reporting and compliance processes?	Yes. Historical transactions executed with external counterparties can be imported into the platform and incorporated into reporting and compliance processes.	Noted.
7.57	If a forecasted transaction that has been hedged later changes, how does the system deal with the hedge and the associated accounting records?	Where a forecast transaction that has been hedged subsequently changes, the hedge designation is not adjusted automatically — it must be removed manually by a user with the appropriate authority. On removal, the designation moves to 'discontinued' in the designation lifecycle and the associated hedge-accounting treatment ceases from that point. Because de-designation is a manual step, it depends on the change being identified and actioned by the user.	Confirmed.
7.58	Does the platform calculate and monitor hedge effectiveness and hedge ineffectiveness, and can it generate the supporting reports, disclosure information and accounting entries required for IFRS 9 and FRS 102 compliance?	Yes. The platform calculates and monitors both hedge effectiveness and hedge ineffectiveness, and tracks designation status on an ongoing basis. It produces the supporting journal entries — the effective portion to other comprehensive income (OCI) and the ineffective portion to profit and loss — together with the related disclosure output.	Confirmed.
Integrations & data exchange			
7.59	Which ERP, accounting, banking and treasury systems can the platform integrate with?	The platform integrates with a range of systems. ERP and accounting integrations comprise Xero, Sage Business Cloud, Oracle NetSuite, Microsoft Dynamics 365 Business Central, Coupa, Tipalti, Odoo, Unit4, Basware and Zip. Banking connectivity is provided through the open-banking aggregators TrueLayer and Plaid. FX execution and payments are delivered through Visa and Ebury Partners. Dedicated treasury-management-system (TMS) integrations are not currently offered; TMS data can be exchanged via CSV import and export.	Noted.
7.60	What integration methods are supported (e.g. API, file upload, connectors, open banking)?	The platform supports API integration (via cloud functions and webhooks), file upload (CSV import and export), pre-built connectors, and open banking (PSD2).	Confirmed.
7.61	Which banking connectivity standards and protocols are supported (e.g. Open Banking, SWIFT, EBICS, host-to-host connectivity)?	Banking connectivity is provided through Open Banking / PSD2 (via TrueLayer and Plaid), with FX execution and payments delivered through the execution partners. SWIFT, EBICS and host-to-host connectivity are not used; connectivity is open-banking based.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
7.62	Which business objects or data types can be exchanged through integrations (e.g. master data, transactions, FX rates, payments, balances, journal entries or forecasts)?	The business objects and data types exchanged through integrations include master data (for example entity and beneficiary data), transactions (trades and payments), forecast exposure data (payables and receivables), and FX rates.	Confirmed.
7.63	What data import, export and reconciliation capabilities are available?	Data can be imported via validated CSV (with a pre-commit preview and fingerprint de-duplication) and exported in XLSX, CSV and PDF formats. Imported data is validated and de-duplicated on entry, and can be reviewed against source records before it is committed.	Confirmed.
7.64	What data import and export capabilities are available, including supported file formats and scheduling options?	Supported formats are XLSX, CSV and PDF for export, and CSV for import. Scheduling: ERP and banking data are synchronised automatically on a daily schedule once a connection is established, and can additionally be synchronised on demand by the user. Recurring report delivery by email can be scheduled by the user. Scheduling of arbitrary ad-hoc file imports or exports to a customer-specified timetable is not currently offered.	Noted.
7.65	Do integrations support real-time, near real-time and scheduled data synchronisation?	Both scheduled and on-demand synchronisation are supported. ERP and banking integrations synchronise automatically on a daily schedule, and a user can additionally trigger a synchronisation on demand at any time from the integrations screen. FX rates are provided on a real-time basis, refreshing every second (see 7.30). Continuous real-time or event-driven streaming from ERP and banking systems is not currently offered; where an ERP document changes between synchronisations, the change is picked up on the next scheduled or manual sync.	Noted.
7.66	What reconciliation capabilities are available to validate imported, exported or synchronised data?	Two reconciliation capabilities are provided. On import, data is validated row-by-row against the expected format, with a pre-commit preview and fingerprint-based duplicate detection, allowing users to confirm accuracy before data is committed. Separately, a bank reconciliation function matches imported bank transactions against trades and payments recorded in the platform, using an automated matching engine that proposes matches with a confidence indicator; users review, accept, amend or exclude each proposed match, and every transaction carries a matched, unmatched or excluded status. Synchronised ERP entries are additionally de-duplicated against the source document identifier on each sync, so re-synchronisation does not double-count an exposure.	Confirmed.
7.67	How does the platform detect, report and recover from integration or data exchange failures?	Detection and reporting operate at three levels. Each ERP and banking connection carries a status (active, expired or revoked) and a last-synchronised timestamp, so a failed or stale connection is visible on the integrations screen; an expired or revoked authorisation is surfaced for re-authentication. Where a payment fails, the customer is notified automatically by email as soon as the execution partner's systems identify the issue. Recovery is by re-authenticating the connection where credentials have lapsed, or by re-running the synchronisation on demand; because imports are de-duplicated on the source document identifier, re-running a failed synchronisation does not create duplicate records.	Noted.
7.68	How are data transfers secured (e.g. encryption in transit, authentication and certificate management)?	Data in transit is encrypted using TLS 1.2 or higher (TLS 1.3 preferred); integrations authenticate via the provider's API credentials/OAuth, with secrets held in a dedicated secrets-management system and cryptographic key/certificate management provided through the underlying cloud platform (see Information Security Policy §4.1 and §6.1).	Noted.
Reporting, dashboards and scheduled outputs			
7.69	Does the platform provide an executive or board report summarising FX exposure, hedges, risk position and policy compliance?	Yes. The platform produces an executive/board report summarising FX exposure, hedges, risk position and policy compliance, and this report is customisable.	Confirmed.
7.70	Can reports be generated for historical reporting periods and trends?	Yes. Reports can be generated for historical reporting periods and present trends over time.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
7.71	Do reports reflect live data, scheduled snapshots or both?	Dashboards reflect live data; reports can be produced on demand against live data or captured as point-in-time snapshots, so both live and snapshot views are available.	Confirmed.
7.72	Are dashboard and reporting outputs available by currency, counterparty, entity, due date and policy status?	Yes. Dashboard and reporting outputs are available by currency, counterparty, entity, due date and policy status (for example hedged, partially hedged or unhedged).	Confirmed.
7.73	Can dashboard and report data be filtered and drilled down to underlying transactions?	Yes. Dashboard and report data can be filtered and drilled down to the underlying transactions.	Confirmed.
7.74	Does the platform provide searchable transaction, payment, hedge and conversion history?	Yes. The platform provides searchable transaction, payment, hedge and conversion history.	Confirmed.
7.75	Can users customise the standard board report, including adding/removing sections or changing fields?	Yes. The board report is customisable, including adding or removing sections and changing the fields displayed.	Confirmed.
7.76	Can users customise dashboards by selecting, rearranging or configuring widgets, charts and KPIs?	No. Dashboards are presented in a standard layout and users cannot currently select, rearrange or configure widgets, charts or KPIs. Dashboard content can be filtered and drilled down to underlying transactions, and the board report is customisable, but dashboard layout itself is fixed. Configurable dashboards are a candidate enhancement.	Confirmed.
7.77	Which output formats are supported for reports (e.g. PDF, Excel, CSV)?	PDF, Excel (XLSX) and CSV.	Confirmed.
7.78	Can users schedule recurring reports to be emailed automatically?	Yes. Users can schedule recurring reports to be delivered automatically by email. Reports are sent to the user's own registered email address; the platform does not support automated delivery to third-party recipients.	Confirmed.
7.79	Can scheduled reports be configured by frequency, timing, recipients and included content?	Partially. Scheduled report delivery can be configured by frequency (daily, weekly or monthly), by included content (the subscriber selects which of the available reports are included), and by time zone. Delivery is to the subscribing user's own registered email address; configuration of arbitrary recipient lists, and configuration of a specific delivery time within the day, are not currently offered.	Confirmed.
Audit trail and data history			
7.80	Does the platform maintain a comprehensive audit log of user and system activity?	Yes. The platform maintains an immutable, append-only audit log of user and system activity.	Confirmed.
7.81	Which activities are captured in the audit log, e.g. sign-in, deal booking, approval, configuration change, export and policy change?	The audit log captures user and system activity, including deal booking, approval, cancellation, settlement and booking failure; payment approval, rejection and re-quotation; pipeline entry creation and deletion; organisation creation and subsidiary linking; user-management activity (invitations issued, accepted, cancelled and resent; user suspension, activation and removal; role changes; approval-limit changes; organisation suspension and activation); external deal import; and the platform features a user has interacted with (see 7.22). Authentication events are individually recorded in the application audit log — successful sign-in, failed sign-in and password-reset requests — in addition to being recorded at the identity-platform level. Report and export actions are not currently recorded as individual audit events. Each entry additionally captures the source IP address and device/browser string of the acting user.	Confirmed.
7.82	Does the audit log capture details such as user ID, timestamp, action type and affected transaction/item, and allow users to search and filter audit records by criteria such as user, date range, entity, transaction or activity type?	Each entry records the acting user, timestamp, action type and the affected transaction/item, and entries can be filtered, including by user.	Confirmed.
7.83	Can the audit log be exported? If yes, in which formats?	Yes. The audit log can be exported from within the platform in Excel (XLSX) format, with the applied filters reflected in the export.	Confirmed.
7.84	Is the audit log immutable, i.e. can entries be amended or deleted by users or administrators?	Yes. The audit log is immutable and tamper-evident — entries cannot be amended or deleted by users or administrators.	Confirmed.

Ref	Requirement	Response	Reviewer Comments
7.85	How long is audit log data retained, and can retention periods be configured?	Operational logs are retained for a minimum of 12 months and trade records for a minimum of 7 years (see Information Security Policy §4.3). Retention periods cannot be customised by the customer, as they mirror regulatory reporting requirements. Customers who wish to preserve a copy of their log for their own records can download it annually.	Noted.
7.86	How does the system support subject access requests or other data retrieval requests where audit log data contains personal data?	Subject access and data-retrieval requests are handled under CrescitPay's Data Protection Policy in accordance with UK GDPR; relevant records, including audit-log data relating to an individual, are retrieved and provided through that process.	Noted.
Training, help and ongoing support			
7.87	What implementation support is provided when a new customer is onboarded?	Structured onboarding brings the client live within approximately four weeks and includes: a welcome call; a configuration session (hedge policy, approval workflow, user permissioning and integrations); a training session; a first-trade walkthrough with the Customer Success Manager on call; and a 30-day review (see Customer Support Standards §6.1).	Noted.
7.88	Can customers use a sandbox environment during setup or training to avoid accidental live trades/payments?	A customer-facing sandbox environment is not provided. During onboarding, the platform is demonstrated and customers are walked through it on a shared screen using a demonstration environment, so that users are trained on the workflows without any risk of accidental live trades or payments.	Noted.
7.89	What training and documentation resources are available for end users and system administrators (e.g. downloadable user manuals, feature guides, documentation, videos, online courses or instructor-led training)? (See also 5.20 & 5.21)	Role-based training sessions are provided during onboarding, with role-specific deep-dives. An in-application user guide is available to all users, covering the platform's principal functions.	Confirmed.
7.90	Does the platform provide in-system help or guidance for users?	Yes. An in-application user guide is available to all users from within the platform, covering the principal functions and how to use them. Guidance is also embedded at the point of use, for example set-up notes on each integration connection card explaining the credentials required and how to obtain them.	Confirmed.
7.91	Does the system provide contextual hints or guided walkthroughs within each tab?	Contextual hints are provided in places — for example explanatory notes on integration set-up and on report and export options — but the platform does not currently provide guided walkthroughs or interactive product tours within each tab. Guidance is instead delivered through the in-application user guide and through onboarding training.	Confirmed.
7.92	Which customer support channels are available (e.g. in-system chat, email, phone, customer portal or dedicated account manager)?	Support is available by email (support@crescitpay.com), by telephone, and through a named Customer Success Manager assigned to each client. Standard support hours are 09:00–18:00 UK time, Monday to Friday (excluding bank holidays), with extended cover for Severity 1 issues (see Customer Support Standards §3).	Noted.
7.93	How are product feedback and enhancement requests captured and prioritised?	Enhancement requests and feedback are captured through the Customer Success Manager relationship and the quarterly client reviews, and are fed into the product roadmap and continuous-improvement process (see Customer Support Standards §10).	Noted.