Ref	Requirement		
	HEADER		
	ICAEW Technical Accreditation Scheme **Legal Information & Research Software Software Evaluation**		
	Law Cyborg		
			
	May 2026		
	© 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)*		

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:	Law Cyborg (LC) is an AI-powered legal and tax research platform for professionals across NZ, AU, and UK with the goal to make meticulous legal research accessible for every practitioner. It provides conversational research tools, legislation and case law search, and automated opinion drafting, allowing users to generate client-specific, context-driven commentary. With over 1,200 firms using LC across the world, it is fast becoming a necessary tool in an adviser's toolkit. AI outputs are presented as research assistance for professional review — not autonomous decision-making. The platform is delivered as a web-based SaaS application accessed via modern browsers, with no client-side software installation required.	

Ref	Requirement	Response	Reviewer Comments
1.09	Supplier background:	Law Cyborg Limited is headquartered in Auckland, New Zealand. The company has approximately 25 FTE and is backed by Blackbird VC. Started by an ex-Big 4 tax adviser as a solution to a problem that has long made tax advice tricky. Designed to make perfect research more accessible the team has grown to include lawyers and accountants who work closely with our expert engineers and data scientists. Law Cyborg holds a SOC 2 Type 1 report (December 2025), with SOC 2 Type 2 audited and notified of successful certification, just waiting final report as of May 2026. ISO 27001 certification also held from April 2026. Continuous compliance monitoring is maintained through Vanta. Security and compliance documentation is published at https://trust.lawcyborg.com/.	
1.10	Product background and suitability for the user:	Law Cyborg is designed for accountancy firms, law firms, and tax professionals who need efficient access to legal and tax research materials. The platform combines AI-powered conversational research with comprehensive legislation, commentary and case law databases covering NZ, AU, and UK jurisdictions. It assists professionals in researching complex legal and tax questions, finding relevant authorities, and drafting opinions. The product is suitable for firms of all sizes, from sole practitioners to Enterprise professional services firms. By using a proprietary vector database, LC is able to structure tax data in a way that enables research to be comprehensively conducted, analysing analogous situations and applying previous decisions to future problems. As an AI tool, it easily flexes to the needs of the user, whether that may be general research, client specific facts, learning a new area or consolidating and confirming existing advice, our (Net Promoter Score) NPS score remains at over 76.	
1.11	Add-on modules:	Law Cyborg's core product includes: a Legal Research Chatbot for conversational research queries, legislation and case law search across NZ/AU/UK. Additional Chatbot knowledge jurisdictions are available for subscription add-on.	
1.12	Typical implementation [size]:	Law Cyborg is a SaaS platform with no on-premises installation. Typical implementations range from individual practitioners to teams of 100+ users within professional services firms involves training for users and admins. Onboarding involves account provisioning and optional SSO configuration (Microsoft SSO available at no additional cost). No server infrastructure or client software installation is required by the customer.	

Ref	Requirement	Response	Reviewer Comments
1.13	Vertical applications:	Law Cyborg is purpose-built for the legal and tax research vertical. It serves accountancy firms, law firms, and tax advisory practices across NZ, AU, and the UK. The platform is specifically designed to support professionals conducting legal and tax research, not general-purpose AI or accounting software.	
1.14	Server platform and database:	Law Cyborg runs on Amazon Web Services (AWS) in the Sydney, Australia region (ap-southeast-2) with London region planned May 2026. Compute: AWS ECS (Fargate, containerised) and AWS Lambda. Database: Amazon Aurora PostgreSQL, encrypted at rest (AES-256). CDN/Edge: AWS CloudFront. All infrastructure is managed by Law Cyborg — no on-premises components.	
1.15	Client specification required:	Law Cyborg is a web-based SaaS application. Users require only a modern web browser (Chrome, Edge, Safari, Firefox — current and previous major versions supported) and an internet connection. No client software installation is required. The platform is responsive and supports desktop, laptop, tablet and mobile access.	
1.16	Partner network:	Law Cyborg is sold and supported directly by Cyborg Limited. There is no Value Added Reseller (VAR) or partner network. All sales, onboarding, training, and support are provided directly by the Law Cyborg team. The organisation mantains many formal and informal relationships with industry bodies and other groups, including CA ANZ, CPA, NZ Law Society, ATAINZ etc.	

Ref	Requirement		
2.	**ISSUES AND CONCLUSION**		
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**	
*	Law Cyborg initially established its operations in New Zealand and Australia, where it provides coverage across New Zealand and Australia (tax law specifically). The vendor is now in the process of entering the UK market, with current UK coverage focused on tax law, drawing on case law, legislation and commentary. At the time of accreditation, customer data is hosted in Sydney; however, the vendor has publicly committed to launching UK hosting in the AWS London region on 15 May, after which UK customer data will be hosted and processed within the UK.	3.27, 6.02, 7.13, 7.43, 7.36, 7.47, 7.48, 7.7.49	
*	The platform services English-speaking jurisdictions only, and English is the sole language supported by the software.	5.01	
*	As with all solutions built on large language models, the application operates within the constraints of the underlying model. At the time of assessment, Gemini 2.5 Pro supports a context window of up to 1 million tokens, with a maximum response length of 8,192 tokens. In extended conversations, accumulated history reduces the capacity available for new retrieved content, and users may be required to start a new chat. These constraints are expected to be suitable for most use cases, but customers should assess suitability for their intended usage patterns.	3.27, 7.28, 7.29, 7.30, 7.50,	
*	The level of access and authentication security is customer-determined. The platform defaults to username and password authentication, with optional Microsoft SSO (SAML/OIDC) and MFA (TOTP via authenticator app) available at no additional cost. As SSO and MFA are not enforced by default, organisations should assess whether this configuration meets their baseline security control requirements.	3.01, 3.09, 7.71, 7.72	
*	User password history is not retained by the vendor and, as a result, password reuse prevention controls are not currently in place at the time of this accreditation. The introduction of password history and reuse prevention is under consideration but is not included on the active product roadmap.	3.10, 3.14, 3.15, 3.16, 7.73, 7.74	

Ref	Requirement		
*	As Law Cyborg is in the process of entering the UK market, the vendor has identified a number of planned developments and enhancements on its future roadmap, which prospective customers should consider as part of their assessment: • Future integrations – Outlook integration by end of May; iManage integration by mid-2026 to support document upload and file management. • Source citation treatment – Enhanced positive and negative citation treatment planned for the UK in Q3 2026; already live in Australia and New Zealand. • Branding and UI customisation – Custom branding (corporate colour palette and logo upload) planned for June 2026. • Password management controls – Password history and reuse prevention not currently implemented; under consideration but not on the active roadmap. • Onboarding, training and user enablement – Vendor-led onboarding and training, supported by a help centre; structured prompting templates in development with a targeted mid-May release.	7.42, 7.69, 7.77, 7.74, 7.36	
Evaluation conclusion			
2.03	For the specific use-cases in support of assisting accountancy firms to make effective use of Legal Information & Research Software software, for which the product is designed, the solution appears to meet this criteria. It continues to be actively developed and enhanced. 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.	**ACCESS AND SECURITY**		
Access control			
3.01	What security features are included to control access to the application?	Access is controlled via username/password authentication with optional Microsoft SSO (available at no additional cost). Multi-factor authentication (TOTP via authenticator app) is supported where SSO is not used. Role-based access control (RBAC) is enforced at the application level. Rate limiting on authentication endpoints provides brute-force protection.	Noted, additional question asked in 7.71
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. Law Cyborg uses role-based access control (RBAC). Users are assigned roles that determine which features and data they can access. Account administrators can manage user roles and permissions within their organisation. Menu options and features are displayed based on the user's assigned role.	Confirmed
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 with unique credentials. Access is managed via user roles (e.g. administrator, standard user) that determine the permissions granted. Account administrators assign roles to users within their organisation.	Confirmed
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. Account administrators can view all users within their organisation, including their roles and access levels, via the administration dashboard.	Confirmed
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. The interface displays only those features and options that the user's role permits. Users without administrator privileges do not see administration functions.	Confirmed
3.06	Does security allow for access to be limited to: - Read only? - Read/write? - Read/amend/delete?	Law Cyborg's RBAC model controls access at the feature level rather than individual read/write/delete granularity on each data field. Standard users can create and view their own research queries. Account administrators have additional permissions to manage users, configure SSO, and manage organisational settings.	Noted
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?	N/A — Law Cyborg does not provide direct database access via ODBC or external report writers. All data access is mediated through the application layer, which enforces the same RBAC controls.	Noted
3.08	Does the system security integrate with Microsoft's Active Directory or other tools that provide a single sign-on?	Yes. Law Cyborg integrates with Microsoft Single Sign-On (SSO) at no additional cost to all customers.	Confirmed
3.09	Does the system provide multi-factor authentication (MFA)?	Yes. Multi-factor authentication is supported via TOTP (time-based one-time password) using an authenticator app. MFA is available where Microsoft SSO is not used. Where SSO is configured, MFA is managed by the customer's identity provider.	Noted
Passwords and access logs			
3.10	Is access to the software controlled by password?	Yes. Access requires either username/password authentication or Microsoft SSO. No anonymous or unauthenticated access is possible.	Confirmed
3.11	Does each user have a separate log on (user ID)?	Yes. Every user has a unique user ID (email address). Credential sharing is prohibited by policy. No shared or generic accounts are used.	Noted
3.12	If there is no password facility please state how confidentiality and accessibility control is maintained within the software?	N/A — password authentication is provided as standard. Microsoft SSO is also available as an alternative.	Noted
3.13	Are passwords masked for any user logging in?	Yes. Passwords are masked (shown as dots) during entry on all login screens.	Noted

Ref	Requirement	Response	Reviewer Comments
3.14	Is password complexity available and enforced?	Yes. Passwords must be a minimum of 8 characters and include at least one uppercase and one lowercase letter. Complexity requirements are enforced at account creation and password change.	Noted
3.15	How many previous passwords are retained / the password history?	None.	Noted, additional question asked in 7.73
3.16	Are passwords encrypted?	Yes. Passwords are stored using bcrypt with salted hashing (10 salt rounds). Passwords are never stored in plain text. All data in transit is encrypted via TLS 1.2+.	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?	Session management uses secure, HttpOnly cookies. Sessions expire after a period of inactivity. When SSO is configured, session management is governed by the customer's identity provider policies. No application data is viewable without an active authenticated session.	Noted
Deletion of transactions			
3.18	Is it possible to delete a transaction?	Law Cyborg is a legal research platform, not a transactional accounting system. Users can delete their own chat history (if chat history is enabled by the account administrator). Account administrators can manage user accounts within their organisation. There are no financial transactions to delete.	Confirmed
3.19	If so, then how are deletions controlled by the system?	Where chat history deletion is permitted, it is controlled through the application's user interface and subject to RBAC. Infrastructure-level audit logs (AWS CloudTrail) record system-level changes. The platform does not process financial transactions that require deletion controls.	Noted
3.20	Are deleted transactions retained in the audit trail (see below) and denoted as such?	N/A — Law Cyborg is not a transactional system.	Noted
Audit trails			
3.21	Does the system have an audit trail (log) which records all changes to transactions in the system?	Law Cyborg maintains audit logging at the infrastructure level via AWS CloudTrail (multi-region, with log file integrity validation) and application-level monitoring via Datadog. User authentication events, system changes, and administrative actions are recorded. As a research platform (not a transactional system), the concept of a transaction audit trail relates to system access and configuration changes rather than financial postings.	Noted
3.22	Does this log also record any system error messages and/or any security violations?	Yes. AWS GuardDuty monitors for security threats (with S3 and EBS malware scanning). AWS CloudTrail records API calls and security events. Datadog captures application errors, performance metrics, and anomalous activity. PagerDuty alerts the security team on critical events.	Noted
3.23	Is it possible to turn off or delete the audit trail?	No.	Noted
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?	Yes. CloudTrail log entries include unique event IDs, timestamps, source IP addresses, and user identity information. Datadog application logs include timestamps and user context. As a research platform, Law Cyborg does not generate sequential financial transaction reference numbers.	Noted
3.25	Are all master file changes recorded in the audit trail?	Yes. Changes to system configuration, user accounts, and organisational settings are recorded in audit logs. AWS CloudTrail captures infrastructure-level changes. Application-level administrative actions are logged via Datadog and system events.	Noted

Ref	Requirement	Response	Reviewer Comments
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?	Yes. Law Cyborg operates in compliance with GDPR, NZ Privacy Act 2020, and CCPA. The system facilitates compliance through: logical data segregation per customer in a multi-tenant architecture; data deletion upon contract termination; a published Privacy Notice (https://lawcyborg.com/privacy/); consent obtained via Terms of Use at account creation; a designated Data Protection Officer (Jacob Sidford, CTO); breach notification within 72 hours of awareness.	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
3.27	Describe your use of sub-processors if any?	Sub-processors that process customer data: - AWS (core infrastructure — compute, database, storage, CDN, monitoring) — Sydney, AU (UK May 2026) - Google Cloud Platform / Gemini (LLM provider — processes user queries) — USA, Zero Data Retention agreement - Stripe (payment processing) — USA - Datadog (APM and logging) — USA Corporate SaaS (no customer data under normal operation): - Microsoft 365 (email, documents, identity/SSO) — Australia - Vanta (compliance monitoring) — USA DPAs are in place with all sub-processors. The sub-processor list is reviewed quarterly and customers are notified of material changes.	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
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?	Yes. Law Cyborg's backup and recovery procedures are documented in the Business Continuity and Disaster Recovery Plan, available under NDA via the Trust Centre (https://trust.lawcyborg.com/). Backups are automated and managed entirely by Law Cyborg — customers do not need to configure or manage backups.	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
3.29	How often are backups taken and to what point can restores be done?	Automated daily backups with a Recovery Point Objective (RPO) of 24 hours. Backups are encrypted (AES-256) and retained within the AWS Sydney region. Recovery Time Objective (RTO) is 8 hours.	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
3.30	How does the software facilitate recovery procedures in the event of software failure? (E.g. roll back to the last completed transaction).	The application runs on AWS ECS (Fargate) with automatic container health checks and restart. Amazon Aurora PostgreSQL provides automatic failover and crash recovery. If a failure occurs, the database rolls back to the last completed transaction. AWS infrastructure provides high availability with automatic recovery from hardware failures.	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?	N/A — Law Cyborg is not a batch processing system. Each user interaction (research query) is processed individually. If a request fails, the user simply resubmits the query. Database transactions use ACID compliance to ensure data integrity.	Noted
3.32	What features are available within the software to help track down processing problems?	Law Cyborg uses Datadog for application performance monitoring (APM), error tracking, and distributed. PagerDuty provides real-time alerting to the engineering and security team on processing issues.	Noted

Ref	Requirement	Response	Reviewer Comments
4.	**DATA PROCESSING AND REPORTING**		
Input and validation of transactions			
4.01	Is data input controlled by self-explanatory menu options?	Yes. The user interface is designed around a conversational research paradigm. Users input natural-language queries via a chat interface, and navigate the platform through clear menu options and intuitive UI elements.	Confirmed
4.02	Are these menus user/role-specific?	Yes. The interface adapts based on the user's role. Account administrators see additional management options (user management, SSO configuration, organisational settings) that are not visible to standard users.	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. Account details (organisation name, user profiles, email addresses) can be amended through the administration interface using standard form fields and dialogue boxes. No system configuration or code changes are required for routine data amendments.	Noted
4.04	Does the software provide input validation checks such as: - [account] code validation? - reasonableness limits? - validity checks?	Yes. The platform validates user inputs including: - Email address format validation on account creation - Password complexity enforcement (minimum 8 characters, mixed case) - Input length limits on all form fields - Character validation to prevent injection attacks (OWASP Top 10 protections) As a legal research platform, the primary user input is natural-language research queries, which are validated for length and sanitised before processing.	Noted
4.05	What control features are within the software to ensure completeness and accuracy of data input?	Required fields are enforced on all forms (account creation, profile updates). Research queries are processed in real-time with immediate feedback — users can see their query and the AI response, allowing them to verify accuracy and refine their input. Input sanitisation and output escaping are applied systemically across the platform.	Noted
4.06	How does the software ensure uniqueness of the input transactions? (i.e. to avoid duplicate transactions)	User accounts are uniquely identified by email address — duplicate accounts for the same email are prevented. Each research query generates a unique session/conversation. As a research platform, the concept of duplicate transactions (as in accounting) does not directly apply — users may intentionally submit similar queries for different research purposes.	Noted
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. The React frontend performs client-side validation (field format, required fields, input length) before submitting data to the server, where the server also does validation checks. This provides immediate user feedback and reduces unnecessary server calls.	Noted
4.08	Is data input by users validated by routines running on the server before data files are updated?	Yes. All data submitted by users is validated by server-side routines (Node.js backend) before any database writes. Server-side validation is authoritative — client-side checks are a convenience layer, not a security boundary.	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?	Yes. All input fields enforce maximum length constraints. Input sanitisation prevents unaccepted characters and protects against injection attacks. These controls are applied both client-side (React) and server-side (Node.js) as part of the platform's OWASP Top 10 protections.	Noted
4.10	Are input errors highlighted?	Yes. Input errors are highlighted inline with clear error messages explaining what needs to be corrected.	Noted

Ref	Requirement	Response	Reviewer Comments
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?	Errors are rejected and an error message is displayed on-screen immediately. Invalid inputs are not accepted or posted to any temporary area. The user must correct the error before the	Noted
4.12	Are responses to erroneous data input clear so that they do not lead to inappropriate actions?	Yes. Error messages are written in plain language, clearly describing the issue and the required corrective action. They do not expose technical details that could be misleading or create security risks.	Noted
4.13	Does the software have an automatic facility to correct/reverse/delete transactions?	N/A in the accounting sense. Users can delete their own chat history (if enabled by the account administrator). There are no financial transactions to correct or reverse. Research queries can be resubmitted with amended wording.	Noted
4.14	If yes, are these logged in the audit trail?	Administrative actions (user account changes, configuration changes) are logged via System Events and Datadog. Chat history deletion events are subject to the same application logging.	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. Database operations use ACID-compliant transactions. If any part of a data write fails, the entire operation is rolled back. No partial records are created.	Noted
4.16	Are messages provided to users clearly explaining whether the data entry or file upload has been processed successfully or not?	Yes. Users receive clear on-screen confirmation when actions complete successfully (e.g. account created, settings saved, query submitted). If an operation fails, an error message explains what happened.	Noted
Import and export of data			
4.17	Can files/attachments be uploaded and stored against any transaction?	Law Cyborg is a legal research platform. Users submit text-based research queries. File upload functionality for attaching documents to conversation is an available feature	Noted - additional questions asked in 7 - 'Document upload / attachments'
4.18	Is there an additional charge made for storage of uploaded files? - If yes, please indicate the cost.	No.	Noted
4.19	Can data be imported into the system from multiple types of files, e.g. XLS, text, CSV?	N/A — Law Cyborg is a research platform, not a data management system. Users interact via the web interface by submitting research queries. Bulk data import from external files is not a standard feature	Noted
4.20	Explain how the system validates imports into the system and what happens to any import which fails?	N/A — see above.	Noted
4.21	Are imported /interfaced transactions detailed in the audit trail? [See also 3.27]	N/A — see above.	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?	Users can copy research results from the interface in PDF or Docx formats.	Confirmed - additional questions asked in 7 - 'Document upload / attachments'
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)?	N/A — Law Cyborg does not require sequential processing steps in the accounting sense (e.g. month-end routines). The research workflow is user-directed: submit a query, review results, refine as needed.	Noted
4.24	Does the software provide automatic recalculation, where appropriate, of data input? (e.g. VAT)	N/A — Law Cyborg does not perform financial calculations such as VAT. The Depreciation Rate Finder tool returns reference information on depreciation rates from official sources.	Noted
4.25	Is a month/period-end routine required to be undertaken?	N/A — Law Cyborg does not require month-end or period-end processing. It is a research tool, not an accounting system.	Noted
4.26	Is it possible to delete accounts if the balance is Nil but transactions have been recorded against the code?	N/A — Law Cyborg does not use account balances or financial transactions. User accounts can be deactivated or removed by account administrators. Upon contract termination, all customer data is permanently deleted.	Noted

Ref	Requirement	Response	Reviewer Comments
4.27	What is the size and format of reference numbers and descriptions within: - Ledgers? - Stock? - Currencies?	N/A — Law Cyborg does not manage ledgers, stock, or currencies. It is a legal research platform, not accounting software.	Noted
4.28	How does the software guard against/warn about duplicate account numbers on set up?	User accounts are uniquely identified by email address. The system prevents creation of duplicate accounts with the same email. The concept of accounting chart-of-accounts duplicate detection does not apply.	Noted
4.29	How does the software enable the traceability [from, to and through the accounting records] of any source document or interfaced transaction?	N/A — Law Cyborg does not process accounting records or source documents. For research queries, each conversation maintains a complete history (if chat history is enabled), showing the original query, AI response, and any follow-up exchanges.	Noted
4.30	What drill down/around functionality is available within the software?	Research results include citations and references to source legislation and case law. Users can follow these references to view the underlying source material. This provides traceability from AI-generated research to authoritative legal sources.	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?	Law Cyborg's legal and tax knowledge base is maintained and updated by the Law Cyborg team. Legislation changes, case law updates, and tax rate changes are incorporated automatically on a weekly basis. Customers do not need to maintain or update reference data themselves.	Noted - additional questions asked in 7.18
Report writer			
4.32	Does the system have an in-built report generator or is a third-party solution used (if so please specify)?	N/A in the traditional accounting report generator sense. Law Cyborg provides AI-generated research outputs directly in the interface. Users can view and export their research results with an in-built pdf/docx generator.	Noted
4.33	Is the report writer based on a standard SQL-type approach and is it flexible and easy to use?	N/A — Law Cyborg does not provide an SQL-based report writer or direct database query access. All data access is through the application interface.	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?	N/A — Law Cyborg does not generate financial or operational reports in the accounting sense.	Noted
4.35	Is a comprehensive data dictionary provided to aid field selection?	N/A — no user-facing data dictionary is provided, as the platform does not offer direct data query or report-building tools.	Noted
4.36	Does the system provide a library of reports and templates which can be amended, saved and re-run?	The platform provides research templates/prompts to guide users in formulating effective research queries.	Confirmed
4.37	Can users create their own reports? If so, what are the controls on users doing this?	N/A — users do not create reports in the traditional sense. Users formulate their own research queries and receive AI-generated responses.	Noted
4.38	Can users create saved searches/filters/queries?	Chat history (when enabled by the account administrator) retains previous research conversations, allowing users to return to prior queries and results. Results can be grouped under an "engagement", which can be used for filtering search results.	Confirmed
4.39	Can regular reports be added to user menus in the appropriate area of the system?	N/A — the platform does not use a traditional report menu structure.	Noted
4.40	Does the system support the production of on demand (interactive) and scheduled batch reports?	N/A — Law Cyborg does not produce scheduled batch reports. Research is performed on-demand by users.	Noted

Ref	Requirement	Response	Reviewer Comments
5.	**USABILITY**		
Ease of use			
5.01	Does the solution provide a multi-language user interface?	The user interface is in English. Law Cyborg serves predominatly English speaking jurisdictions. Multi-language UI is not currently offered.	Noted
5.02	Does the system allow for customisable branding and UI (e.g. corporate colour palate, upload company logo, etc)?	No, it is planned for Q2 2026.	Noted - additional questions asked in 7.77
5.03	Does the system have a similar look and feel and overall consistency between screens and modules?	Yes. The platform uses a consistent design language (built with React) across all screens and features. Navigation, layout, typography, and interaction patterns are uniform throughout.	Noted
5.04	Is data entry easily repeated if similar to previous entry?	Yes. Users can refine previous queries in a conversational flow, building on prior context within a chat session. Chat history (when enabled) allows users to reference and reuse previous research approaches.	Noted - additional questions asked in 7.20
5.05	Does the software prevent access to a record while it is being updated?	N/A — Law Cyborg is a research platform where each user's queries and sessions are independent. There is no concept of multiple users editing the same record simultaneously.	Noted
5.06	Is there locking at file or record level?	N/A — see above. The multi-tenant architecture with logical data segregation ensures each user's data is independent. Database-level concurrency control (PostgreSQL MVCC) handles any concurrent operations at the infrastructure level.	Noted
5.07	Does the software allow for the running of reports whilst records are being updated?	N/A in the traditional sense. Users can submit new research queries at any time without affecting other users or system operations.	Noted
5.08	Can timestamps or user comments be added to transactions?	Chat history entries are timestamped. Users can provide feedback on AI responses via the 'Give Feedback' feature. The platform does not have a general-purpose commenting system on transactions, as it is a research tool.	Noted
5.09	Is there the ability to store preferences and default values on a per-user basis e.g. department/team/user?	User preferences based on usage patterns, usage settings are account level rather than user level.	Noted
5.10	Does the system have the ability to provide user-defined fields with associated validation of data input?	N/A — Law Cyborg does not provide user-defined custom fields. The data model is managed centrally by Law Cyborg.	Noted
5.11	Can the system provide users with reminders and notifications e.g. workflows?	No.	Noted
5.12	If the system provides workflows, does it have functionality to substitute/delegate authorisations?	N/A — Law Cyborg does not include workflow or approval routing features. It is a research tool, not a process management system.	Noted
5.13	Describe the tools and features available for a power user to make configuration changes such as amending a workflow.	Account administrators can manage users, configure SSO, enable/disable chat history, and control the 'Give Feedback' feature at the organisational level. No code-level or workflow configuration tools are provided to end users.	Noted
5.14	Is there the ability for users to define and configure layouts of letters and forms?	There is not yet ability to define & configure layouts of exported research results, planned Q2 2026.	Noted
5.15	Can users save the parameters of searches?	Chat history (when enabled) preserves previous research conversations.	Noted
5.16	Does the system have a "universal search" option, allowing a search to be undertaken over all modules of the system?	Law Cyborg's core function is search — the conversational research interface searches across legislation, case law, and tax databases covering UK, AU, and NZ from a single unified interface.	Noted
5.17	Can the system store menu option 'favourites' on a per user basis?	System stores last used tools in session, as a 'preference' system.	Noted
5.18	Can a user open multiple windows accessing the same or different modules of the system?	Yes. As a web-based application, users can open multiple browser tabs to access different features or research sessions concurrently.	Confirmed
5.19	Can more than one software function be performed concurrently?	Yes. Multiple research queries can be run in separate browser tabs. The platform supports concurrent usage without conflict.	Noted

Ref	Requirement	Response	Reviewer Comments
User documentation and training			
5.20	Confirm whether a user manual / instructions is provided and how this is distributed?	help.lawcyborg.com is provided for instructions, along with an onboarding interactive app tour & email sequence. Most new users also go through a demo with a specialist consultant who shares the key features and limitations of the system.	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?	Yes.	Noted
5.22	Is context-sensitive help available within the system?	Tooltips and inline guidance is available within the platform.	Confirmed
5.23	Is the manual and/or help editable by the user (subject to the permissions matrix)?	No. User documentation is maintained by Law Cyborg. Users cannot edit the help content.	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)?	As a SaaS platform, source code escrow is not available.	Noted
5.25	Please detail the training options available?	Law Cyborg provides onboarding and training directly through interactive onboarding flows or with our team.	Noted - additional questions asked in 7.36
5.26	Who provides training: - Software House? - VAR?	Training is provided directly by Law Cyborg (the software house). There are no Value Added Resellers (VARs). All training and support is delivered by the Law Cyborg team.	Noted
Support and maintenance			
5.27	How is the software sold: - Direct from the software house? - Via a Value Added Reseller (VAR) or Integrator?	Direct from the software house (Law Cyborg / Cyborg Limited). There are no VARs or resellers.	Noted
5.28	How is the product supported: - Direct from the software house? - Via a Value Added Reseller (VAR) or Integrator?	Direct from the software house (Law Cyborg / Cyborg Limited). There are no third-party support partners.	Noted
5.29	Do VARs have to go through an accreditation process?	N/A — Law Cyborg does not use VARs or resellers.	Noted
5.30	Is the software sold based upon number of named users or a number of concurrent users?	Named users. The subscription is priced per user per month.	Noted
5.31	The supplier should detail the support cover options available, covering: - The hours provided? - Associated costs? - The global regions covered?	Support is provided via help.lawcyborg.com chat and email (support@lawcyborg.com). Support hours are business hours and included in standard subscription cost.	Noted
5.32	Detail the process by which customers raise support requests and how these can be viewed/managed?	Customers raise support requests via in-app chat or email.	Noted
5.33	Please note the methods of support available: - Telephone? - Internet chat? - Remote access to customer workstation? - Other, please specify?	- In-app chat (Intercom) - Email (support@lawcyborg.com)	Noted
5.34	Do you offer service credits for failure to meet performance around SLA and uptime (if applicable)	Service credits for failure to meet SLA are negotiated on a per-customer basis as part of enterprise contracts.	Noted
5.35	What is your escalation path for tickets which have not been resolved within a reasonable time?	Front-line support 48 hours -> CTO	Noted
5.36	How often are general software enhancements provided?	Law Cyborg is actively developed with frequent updates. As a SaaS platform, enhancements are deployed continuously, Releases are daily.	Noted
5.37	Will they be given free of charge?	Yes. All software enhancements and updates are included in the subscription at no additional cost. Updates are deployed automatically as part of the SaaS service.	Noted
5.38	How are enhancements and bug fixes provided to customers?	As a SaaS platform, all enhancements and bug fixes are deployed centrally and become available to all users automatically. No customer action is required to receive updates.	Noted

Ref	Requirement	Response	Reviewer Comments
5.39	Is “hot line” support to assist with immediate problem solving available?	In-app chat (Intercom) provides near-immediate access to the support team for urgent issues. Mid-market and Enterprise customers receive an account manager with phone number for immediate access.	Noted
5.40	If so, is there an additional cost involved?	Support is included in the subscription fee.	Noted
5.41	At what times will this support be available?	Law Cyborg operates across NZ, AU, and UK. 24/7 support is available	Noted. While Law Cyborg states that 24/7 support is available across NZ, AU and the UK, UK operations are expected and therefore will only be fully operational from mid-May.
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. All Law Cyborg features (research chatbot, legislation) are fully integrated within a single web application. No additional setup is required to use different features.	Noted
5.43	Are they integrated on real time basis or batch basis?	Real-time. All features operate in real-time within the same web application.	Noted
5.44	Can the software be linked to other packages e.g. word processing, graphics, financial modelling, to provide alternative display and reporting facilities?	Research results can be copied from the web interface and pasted into any application (e.g. Microsoft Word, email). Microsoft Word integration planned for Q2 2026.	Noted
5.45	Can definable links to spreadsheets be created?	N/A — Law Cyborg does not produce structured data suitable for live spreadsheet linking. Research outputs are text-based and can be copied to spreadsheets manually.	Noted
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.	Chat history (when enabled) provides secure storage of research conversations, encrypted at rest (AES-256) and in transit (TLS 1.2+) in the AWS Sydney region. Documents can be uploaded for chat context, these are encrypted and uploaded to a private AWS s3 bucket. The platform does not provide general-purpose document storage.	Noted
5.47	Can documents be scanned into a secure repository?	N/A — Law Cyborg does not provide document scanning or a document repository.	Noted
5.48	Does the system provide data migration tools for transactional and master data sets (e.g. employees customers, suppliers, journals, invoices).	N/A — Law Cyborg does not require data migration from other systems. As a research platform, there is no transactional or master data to migrate in. User accounts are provisioned directly.	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?	N/A	Noted
5.50	Does the system support mobile working?	Yes. As a web-based application, Law Cyborg is accessible from any device with a modern web browser, including tablets and mobile devices. The interface is responsive and adapts to different screen sizes.	Noted - additional questions asked in 7.35

Ref	Requirement	Response	Reviewer Comments
6.	**SaaS HOSTING**		
	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?	Law Cyborg uses Amazon Web Services (AWS) data centres in the Sydney, Australia region (ap-southeast-2) with London planned for May 2026. The service is SaaS. Law Cyborg controls the application stack running on AWS infrastructure. AWS manages the physical data centre facilities.	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
6.02	Does the customer get a choice of the jurisdiction in which their data resides?	Customer data is stored exclusively in the AWS Sydney (ap-southeast-2) region, London planned for May 2026. Customers do not currently have a choice of alternative jurisdictions.	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
6.03	What certification(s) do you or your platform operators hold relating to your data centres and your business operations?	Law Cyborg: - SOC 2 Type 1: completed December 2025 - SOC 2 Type 2: certification obtained, awaiting final report in May 2026. - ISO 27001 certification obtained in April 2026. AWS data centre certifications include: ISO 27001, SOC 1/2/3, PCI DSS Level 1, and others. Law Cyborg maintains continuous compliance monitoring via Vanta.	Noted
6.04	Do you or your platform operator have an SSAE16 (System and Organisation Controls) report available?	Yes. Law Cyborg holds a SOC 2 Type 1 report (December 2025). SOC 2 Type 2 is expected May 2026. The report is available for review under NDA via the Trust Centre (https://trust.lawcyborg.com/). AWS also holds SOC 1/2/3 reports for its data centre operations.	Noted
6.05	What are the physical controls over the:- - Premises? - Fileservers? - Communications equipment?	Physical security is managed by AWS, whose data centres are ISO 27001, SOC 1/2/3, and PCI DSS Level 1 certified. AWS facilities include: 24/7 staffed security, biometric access controls, CCTV monitoring, environmental controls, and physical intrusion detection. Law Cyborg does not operate any on-premises servers or communication equipment — all infrastructure is cloud-hosted.	Noted
6.06	Is the space in this/these data centre(s) shared with any other companies?	Yes — AWS data centres are shared multi-tenant facilities. However, AWS provides strong isolation between customers via virtualised compute, network isolation (VPC), and encryption. Law Cyborg's infrastructure runs within a dedicated AWS VPC with strict network security controls.	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?	Law Cyborg uses a multi-tenant architecture with logical security to partition customer data within a shared database. Customer data is logically segregated — each customer's data is identified and isolated at the application and database level. Customers do not share database tables with other customers' data visible.	Noted

Ref	Requirement	Response	Reviewer Comments
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?	Customer data segregation is enforced at the application level through tenant-scoped queries — every database operation is scoped to the authenticated user's organisation. Role-based access control ensures users can only access data within their own organisation. The architecture prevents cross-tenant data access by design.	Noted
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?	All database queries are scoped to the authenticated user's organisation (tenant). The application layer enforces tenant isolation — there is no mechanism for a user to query or access another customer's data. This is validated as part of the SOC 2 audit controls and penetration testing programme.	Noted
6.10	How is [Internet] communication traffic monitored to identify potential problems before they happen: - From a performance perspective? - From a security standpoint?	Performance: Datadog APM provides real-time monitoring of application performance, latency, throughput, and error rates. Security: AWS GuardDuty provides threat detection (including S3 and EBS malware scanning). AWS CloudTrail monitors API activity. AWS CloudFront provides DDoS protection at the edge. PagerDuty provides real-time alerting to the security and engineering team.	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?	All database operations use ACID-compliant transactions (PostgreSQL). If a connection is interrupted mid-operation, the database transaction is rolled back, preventing partial writes or data corruption. The web application handles network interruptions gracefully with appropriate error messages to the user.	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 Law Cyborg is encrypted. TLS 1.2+ is enforced on all connections. HTTPS is mandatory (HTTP redirected to HTTPS). HSTS (HTTP Strict Transport Security) headers are enforced. Authentication cookies have the Secure flag set. This covers all data, not just login credentials.	Noted
6.13	Is data on your servers encrypted at rest?	Yes. All data at rest is encrypted using AES-256, including the Aurora PostgreSQL database, S3 storage, EBS volumes, and backups.	Noted
6.14	What level of encryption is used?	Data at rest: AES-256. Data in transit: TLS 1.2+ (all connections). Password storage: bcrypt with salted hashing (10 salt rounds). Internal service communication (CloudFront to ALB, ALB to ECS, ECS to Aurora, ECS to GCP LLM): all TLS.	Noted
6.15	Is a staging environment provided that is an exact replica of production; which can be used for testing purposes?	Yes. Production and non-production environments are fully separated with separate credentials and access controls. A staging environment exists for testing before production deployment.	Noted
6.16	Is a test environment provided to test configuration changes? If so, is there an additional charge for this?	A separate test/staging environment is not currently available for testing configuration and code changes before production deployment.	Noted

Ref	Requirement	Response	Reviewer Comments
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?	Law Cyborg complies with GDPR, NZ Privacy Act 2020, and UK data protection legislation. Mitigations include: - Data encrypted at rest and in transit - Customer data stored in AWS Sydney (ap-southeast-2) with DPAs in place with all sub-processors - Privacy Notice published at https://lawcyborg.com/privacy/ - Data deleted upon contract termination - Designated DPO (Jacob Sidford, CTO) - Breach notification within 72 hours - User data not used for AI model training (Google Zero Data Retention agreement) - Appropriate controls being connected via Vanta for automated compliance monitoring of systems	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
6.18	Are you subject to any legal or regulatory requirements obliging you to retain a copy of customer data?	Law Cyborg does not retain customer data beyond the contract term for regulatory reasons. Upon contract termination, all customer data is permanently deleted. There are no legal or regulatory requirements obliging Law Cyborg to retain a copy of customer data after the contract ends.	Noted
6.19	Who will be able to access or see customer data?	Access to customer data is restricted to authorised Law Cyborg personnel who require it for service delivery or customer support. Production access uses just-in-time (JIT) elevated permissions for incidents and customer support only. All personnel with data access have signed confidentiality/NDA agreements and completed security training. Sub-processors (AWS, Google GCP, Stripe, Datadog) process customer data as detailed in the sub-processor list, all under DPAs.	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.	- Background checks completed before access to environments is granted - Confidentiality agreements signed by all employees and third parties - Production access requires JIT elevated permissions (not standing access) - Least-privilege IAM roles at infrastructure level - RBAC enforced at application level - All production access logged via AWS CloudTrail - Formal offboarding process revokes access promptly when staff leave	Noted
6.21	Explain the release management procedures in place and the associated segregation of duties ?	All changes are subject to formal analysis, testing, and approval before acceptance. Version control (Git/GitHub) with a consistent build process. A defined change management process covers recording, approval, implementation, testing, and versioning prior to production deployment. Impact of significant changes is analysed and approved by the IT Head. Production and non-production environments are fully isolated.	Noted
6.22	Is there sufficient segregation of duties preventing system developers from accessing and changing live applications and data files?	Yes. Production and non-production environments are fully separated with separate credentials and access controls. Developers do not have standing access to production data. Production access requires JIT elevated permissions, which are time-limited and audited.	Noted

Ref	Requirement	Response	Reviewer Comments
6.23	Explain the review and approval procedures covering system operations staff when emergency changes need to be made to live applications and data?	Emergency changes follow the same version-controlled deployment process but with expedited review and approval. All emergency changes are logged via Git, CloudTrail, and the deployment pipeline. Post-incident review captures the change, its justification, and any follow-up actions.	Noted
6.24	Is an audit trail always maintained of these emergency changes?	Yes. All changes, including emergency changes, are recorded in Git (version control), AWS CloudTrail (infrastructure audit trail), and the deployment pipeline. Log file integrity validation is enabled on CloudTrail.	Noted
6.25	What procedures are in place when members of staff leave to ensure that their system access is stopped?	Law Cyborg has a formal termination and offboarding process documented in a checklist. Physical and logical access is revoked promptly. Company equipment is returned. All active sessions are terminated. Where Microsoft SSO is configured, deprovisioning the identity provider account immediately revokes all application access.	Noted
Platform and service levels			
6.26	Which databases and servers are used to host the software?	Database: Amazon Aurora PostgreSQL. Compute: AWS ECS (Fargate, containerised) and AWS Lambda. CDN/Edge: AWS CloudFront. All hosted in AWS Sydney (ap-southeast-2).	Noted - additional questions asked in 7 - 'Hosting, UK market readiness & data residency transparency'
6.27	What forms of user authentication are supported e.g. user names, passwords certificates, tokens etc.?	- Username (email) and password - Microsoft SSO (SAML/OIDC) — available at no additional cost - Multi-factor authentication (TOTP via authenticator app) where SSO is not used - Secure, HttpOnly session cookies	Noted
6.28	What is the proposed product/service availability percentage?	99.9% uptime	Noted
6.29	What percentage availability has been achieved over the past 12 months?	99.8% uptime	Noted
6.30	Is a service level agreement ("SLA") offered regarding: - Service availability? - Data recovery?	Service availability is offered as part of Enterprise Contracts.	Noted
6.31	Is the service available 24x7 or are there downtime periods for maintenance?	The service targets 24/7 availability. Maintenance is typically performed during off-peak hours with minimal or no downtime, if downtime is required 2 weeks noticew ill be provided. As a SaaS platform on AWS, most infrastructure maintenance is handled transparently by AWS.	Noted
6.32	Is the customer made aware of maintenance periods in advance?	2 weeks notice if downtime is required.	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 web browser. No client software installation is required. Supported browsers include Chrome, Edge, Safari, and Firefox (current and previous major versions).	Noted
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 — Law Cyborg does not require downloading or running any executable programs. The application runs entirely within the web browser.	Noted
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. Law Cyborg uses modern, actively maintained technologies: React (frontend), Node.js (backend), PostgreSQL (database), and current AWS services. No obsolescent or end-of-life technologies are in use.	Noted

Ref	Requirement	Response	Reviewer Comments
Platform security			
6.36	What security steps are taken to prevent and detect intrusion attempts?	- AWS GuardDuty for threat detection (including S3 and EBS malware scanning) - AWS CloudTrail for API audit logging (multi-region, with log file integrity validation) - AWS CloudFront for DDoS protection - Datadog APM for application monitoring and anomaly detection - PagerDuty for real-time security alerting - Regular patching of dependencies in line with severity-based SLAs - Annual penetration testing by independent third party	Noted
6.37	Is firewall hardware and software used to protect the live systems from unauthorised access?	Yes. AWS provides network-level firewalling via Security Groups and Network ACLs. The application runs within a VPC with restricted ingress/egress rules. AWS CloudFront provides edge-level protection including Web Application Firewall (WAF) capabilities.	Noted
6.38	Which monitoring software is used to create alerts when intrusion attempts are suspected?	AWS GuardDuty monitors for unauthorised access, unusual API calls, and known malicious IPs. Datadog detects application-level anomalies. PagerDuty routes alerts to the security team in real-time.	Noted
6.39	Are designated staff responsible for receiving and urgently responding to these alerts?	Yes. The security team receives automated alerts via PagerDuty. Breach determinations are made by the CEO, CTO, and Security Delegate. The incident response process is documented in the formal Incident Response Plan.	Noted
6.40	Have clear procedures been established for identifying and responding to security incidents?	Yes. Law Cyborg maintains a formal Incident Response Plan covering identification, containment, eradication, recovery, and post-incident review. Notification to affected parties and regulators within 72 hours of initial awareness. Legal Counsel reviews all external breach notices before sending.	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.	Yes. Third-party dependencies are maintained and patched in line with severity-based SLAs. Known exploited vulnerabilities are prioritised immediately. Where patching is not possible, equivalent mitigations are implemented. AWS manages OS-level patching for managed services (Fargate, Aurora, Lambda).	Noted
6.42	List the procedures and software tools in place to prevent or detect and eliminate interference from malicious code, such as viruses?	- AWS GuardDuty with S3 and EBS malware scanning - Input sanitisation and output escaping across the application - Dependency vulnerability scanning as part of the build pipeline - No client software installation — reducing the attack surface - Developer training on OWASP Top 10 and secure coding practices	Noted
6.43	Is a system log maintained by the service provider that details - User access? - User activity? - Error messages? - Security violations?	Yes. The system maintains comprehensive logs covering: - User access: authentication events logged - User activity: application-level activity logged via Datadog - Error messages: captured by Datadog APM - Security violations: detected by GuardDuty and CloudTrail All logs are timestamped and include user identity where applicable.	Noted
6.44	Is this log available to the customer?	No, logs are not currently available unless requested by the account owner.	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. Law Cyborg has not experienced any security breaches in the past three years.	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?	Annual penetration testing is performed by an independent external specialist. The most recent test was completed in July 2025 with no outstanding high or critical findings. The penetration test executive summary is available under NDA via the Trust Centre.	Noted
6.47	Are procedures in place to ensure that any weaknesses found by penetration testing are addressed quickly?	Yes. Findings from penetration tests are prioritised by severity and remediated in line with severity-based SLAs. Critical and high findings are addressed immediately. The July 2025 test resulted in no outstanding high or critical findings.	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 — annual penetration testing is performed. In addition, continuous vulnerability management includes: dependency scanning, AWS GuardDuty threat detection, and Vanta continuous compliance monitoring.	Noted
6.49	Are security procedures regularly reviewed? Please indicate frequency of reviews.	Yes. Security policies and procedures are reviewed at least annually. Vanta provides continuous compliance monitoring against SOC 2 and ISO 27001 control frameworks. The ISMS policy (02-ISMS) is reviewed and approved annually by leadership.	Noted
6.50	What security reporting is provided demonstrating compliance against certification(s) and policy(ies)?	Law Cyborg's Trust Centre (https://trust.lawcyborg.com/) publishes security and compliance documentation. The SOC 2 Type 1 report, penetration test summary, and policy documents are available under NDA. Vanta provides continuous compliance monitoring and evidence collection.	Noted
6.51	How are security breaches communicated to customers?	The Incident Response Plan requires notification to affected parties within 72 hours of initial awareness. Legal Counsel reviews and approves all external breach notices before sending. Breach notification commitments are included in the Subscriber Terms and Privacy Notice.	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?	- How: Automated daily snapshots of the Aurora PostgreSQL database - Frequency: Daily - What: Full database and associated data stores - Media: AWS-managed storage (S3/EBS) - Location: AWS Sydney (ap-southeast-2) — same region as production - Copies: Multiple copies maintained by AWS for durability - Retention: Per the retention policy (available under NDA) - Access: Restricted to authorised Law Cyborg personnel via IAM with JIT access - Encryption: Yes, AES-256 at rest	Noted
6.53	How frequently is a test-restore of backups undertaken?	Annual DR testing, including backup restoration verification.	Noted
6.54	Can the provider restore from a backups that it has taken at a customer request?	N/A, Law Cyborg can restore data from backups at a customer's request.	Noted
6.55	Does a customer have the ability to undertake their own backups?	Customers do not have direct access to take their own database backups. All backups are managed by Law Cyborg although they can export all their data as per GDPR requirements.	Noted
6.56	If so, can a customer restore data a backup that they have taken?	No.	Noted

Ref	Requirement	Response	Reviewer Comments
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?	Law Cyborg maintains a Business Continuity and Disaster Recovery Plan covering: - Database/application corruption: restore from automated daily backups (RPO 24 hours) - Hardware failure: AWS provides automatic hardware replacement and failover - Fire, flood, disaster: AWS data centres have redundant power, cooling, and fire suppression; data replicated within the region - Communication failures: AWS CloudFront provides edge caching and DDoS protection; Aurora provides automatic failover RTO is 8 hours. The BC/DR Plan is available under NDA via https://trust.lawcyborg.com/.	Noted
6.58	How often are these plans tested?	Annually, including backup restoration verification.	Noted
6.59	How often are these plans reviewed and updated?	The BC/DR Plan is reviewed and updated at least annually, or more frequently when significant changes to infrastructure or services occur.	Noted
6.60	What is the longest period of time envisaged that service may not be available?	The Recovery Time Objective (RTO) is 8 hours. This represents the maximum target time from a disaster event to full service restoration.	Noted
6.61	What are your: - Recovery Point Object (RPO) standards? - Recovery Time Objective (RTO) minimum standards?	- Recovery Point Objective (RPO): 24 hours - Recovery Time Objective (RTO): 8 hours	Noted
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?	Timestamps are local to the user's browser.	Noted
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?	Data export is available for customers to extract their resultst & data.	Noted
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?	Data access for customers to extract records and make alternative arrangements for a period of time.	Noted
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?	Law Cyborg's infrastructure runs on AWS, which is operated independently. If Law Cyborg were to cease trading, customer data would remain on AWS infrastructure for the duration covered by the AWS service agreement.	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?	Processes are resilient with auto-scaling to meet demand automatically. For system outages there are runbooks, processes and oncall rotation in place to ensure service continuity.	Noted
Platform change management			
6.67	Describe your approach to upgrades including what option customers have not to take upgrades (if any)?	As a SaaS platform, all users receive the same version. Upgrades are deployed centrally by Law Cyborg. Customers do not have the option to decline or defer upgrades, as this is standard for SaaS delivery.	Noted
6.68	Are users able to test the application before new versions go into live use?	New features may be released with feature flags, allowing controlled rollout. Customers do not typically test new versions before they go live, as changes are deployed centrally after internal QA.	Noted
6.69	Are users given notice before application changes are applied to the live system?	Significant changes have emails sent in the week-prior.	Noted
6.70	Are changes delivered into the live environment "switched off" to enable users to test them before enabling them for their environment?	Feature flags are used for controlled rollout of changes. Not all changes are delivered in a 'switched off' state for customer testing — bug fixes are deployed directly.	Noted

Ref	Requirement	Response	Reviewer Comments
6.71	Describe what testing and QA processes are undertaken before upgrades and other changes are made live/available to customers?	All changes go through a defined process: development in isolated environments, code review, automated testing, manual QA in staging, approval, and then deployment to production. Production and non-production environments are fully isolated. Impact of significant changes is analysed and approved by the IT Head.	Noted
6.72	Explain the release management procedures in place and the associated segregation of duties?	Version control (Git/GitHub) with a consistent build process. All changes are subject to recording, approval, implementation, testing, and versioning prior to production deployment. Production and non-production environments have separate credentials and access controls. Developers do not have standing production access.	Noted
6.73	Are users informed when they next login of the application changes that have gone into live use?	Users do not see release notes upon login.	Noted
6.74	Do customer staff have to take any action (e.g. regression testing) when new editions, patches or upgrades are released? If so, please describe what they should ordinarily do.	No. Customers are not required to perform regression testing when updates are released. All QA and testing is performed by Law Cyborg before deployment.	Noted
Subscription options			
6.75	What is the minimum level of commitment must the customer sign up to, e.g. 36 months?	No-lock in monthly subscription. 1 month.	Noted
6.76	Where online payment is used, what type of security is used to protect sensitive information?	Payment processing is handled by Stripe, a PCI DSS Level 1 certified payment processor. Law Cyborg does not store credit card details. All payment data is transmitted directly to Stripe over TLS.	Noted
6.77	Where online subscription / payment is used, is an invoice provided to the customer and, if so, in what format?	Yes. Invoices are provided to customers via email.	Noted
6.78	When subscriptions need to be renewed, what advance notice is provided and what is the time limit for renewal?	Advance notice is not provided on a rolling 1 month contract.	Noted
6.79	Is there a procedure for late renewal and is there a time limit after which subscriptions cannot be renewed?	Account Owner can update billing details and their subscription will be continued at any time.	Noted
6.80	How soon after creating or renewing a subscription (if applicable) can the system / service be used?	Access is available immediately after account provisioning. For SSO-configured accounts, access is available once SSO integration is completed.	Noted
6.81	What notifications / confirmations are provided to the customer regarding subscriptions and payments?	Payment confirmation & Payment failure emails are provided upon each payment attempt.	Noted
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?	A user can access their data for export until they close their accounts. Data will be retained for 7 years.	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?	7 years.	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?	Upon contract termination, all customer data is permanently deleted from Law Cyborg's systems, including backups (once the backup retention cycle completes).	Noted
6.85	What is your processes regarding disposal of end-of-life and failed hardware devices that were used to operate your service?	N/A — Law Cyborg does not operate physical hardware. All infrastructure runs on AWS managed services (Fargate, Aurora, Lambda). AWS manages hardware lifecycle and disposal in accordance with its own security certifications (ISO 27001, SOC 2, etc.).	Noted
SaaS/Hosted reporting			
6.86	Are reports produced from the same software as the financial applications or is separate reporting software used?	All outputs are produced from the same web application. No separate reporting software is used. Research results, chat history, and any administrative views are all served from the same platform.	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. All outputs are viewable within a standard web browser. No additional software is required.	Confirmed

Ref	Requirement	Response	Reviewer Comments
6.88	What browser versions are support: - On desktop/laptop (PC, Mac, Linux)? - On Tablets? - On mobiles?	Desktop/laptop: Chrome, Edge, Safari, Firefox (current and previous major versions) on PC and Mac. Tablets: Supported via the same browsers. Mobile: Accessible via mobile browsers (responsive design).	Noted
6.89	Is access to the reporting facilities and data controlled by the same procedures as access to the main application?	Yes. All access (including viewing outputs, research results, and administrative views) is controlled by the same authentication and RBAC system.	Noted
6.90	If it's different, explain the user access control facilities available to ensure information is only viewed by users with appropriate authority?	N/A — the same access controls apply throughout the application.	Noted
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?	Research results are displayed as HTML within the web browser. Users can copy content to other applications or export to PDF or DocX format.	Noted
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?	Research results are generated dynamically and displayed in the browser. If chat history is enabled, conversations are stored on Law Cyborg's servers (encrypted, in AWS Sydney, UK in May 2026). Data is not stored on the user's local computer (beyond normal browser caching). Access to stored conversations is controlled by the same authentication and RBAC system.	Noted
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?	Standard web browser caching may temporarily store page content on the user's device. HTTPS is enforced with appropriate cache-control headers.	Noted
6.94	Are communications between the browser and the server encrypted for any report related communications?	Yes. All communications between the browser and server are encrypted using TLS 1.2+, including all report and research output data. HTTPS is mandatory (HTTP redirected to HTTPS). HSTS is enforced.	Noted
6.95	If reports are produced dynamically each time the user views them can historical reports be reproduced at any time?	If chat history is enabled, users can return to previous research conversations at any time. Without chat history enabled, previous conversations are not retained and cannot be reproduced.	Noted
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?	Research results can be navigated conversationally — users can ask follow-up questions to drill deeper into a topic. Citations and references in research results link to source legislation and case law. The interface does not support traditional tabular drill-down (columns, rows, time periods) as it is a research tool, not a reporting/BI platform.	Noted
6.97	Can report data be reliably copied and pasted direct from browser viewable reports to an MS Excel spreadsheet retaining any table layout?	Research results are text-based and can be copied and pasted to spreadsheets. As research outputs are prose rather than tabular data, the concept of retaining table layout is of limited applicability.	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?	If a network interruption occurs during response generation, the user sees a loading indicator or error message indicating the request did not complete. Partial responses are not displayed without indication — the user can clearly tell if a response is incomplete. The user will then be able to reload their conversation and the research will be returned from the server.	Noted

Ref	Requirement	Response	Reviewer Comments
7.	**LEGAL INFORMATION & RESEARCH**		
Gen AI Platform architecture			
7.01	Does the platform rely on a single large language model (LLM) interaction, or does it use a multi-step processing pipeline? (e.g. retrieval, ranking, generation, post-processing, etc)	Law Cyborg uses a multi-step processing graph, not a single LLM interaction. The pipeline consists of several distinct stages, including question decomposition, semantic analysis & Vector Retrieval, Ranking & Selection	Noted
7.02	Where a multi-step pipeline is used, which components are deterministic versus probabilistic? For the user facing response, which elements are templated or fixed and which are dynamically generated?	**Deterministic components:** - Document retrieval and search (vector similarity and keyword matching) - Source ranking and weighting (score normalisation, namespace weighting, temporal decay) - Citation validation and numbering (post-processing maps citation numbers to retrieved chunks) - Certainty score calculation (mathematical formula applied to retrieval scores) - Source categorisation (cited, see-also, reviewed-but-not-cited) **Probabilistic components:** - Question decomposition (LLM-generated) - Excerpt description generation (LLM-generated) - Final response text (LLM-generated) **Templated/fixed user-facing elements:** - Citation format (superscript numbers after full stops) - Source metadata (title, date, origin URL, document type) - Certainty score display (percentage) - Disclaimers and structural formatting rules	Noted
7.03	Which foundation model(s) is the platform based on?	Google Gemini 2.5 Pro	Noted
7.04	Can customers select a foundation model if they do not wish to use a particular model supplier?	No not currently.	Noted
7.05	Is a RAG (Retrieval-Augmented Generation) architecture used exclusively? If not, what triggers the use of RAG?	Law Cyborg uses RAG (Retrieval-Augmented Generation) for all legal research queries, including user uploaded document queries.	Noted
Data handling, retention and deletion			
7.06	How is customer data handled once submitted to the platform (including chat history, queries, results, saved searches/filters, and 'engagements'), particularly where content may be sensitive or confidential?	Customer data submitted to the platform is handled as follows: - Queries (prompts): Sent to the retrieval pipeline and then to the LLM provider (Google Gemini) for response generation. Queries are stored in the database as part of the conversation record only if chat history is enabled for the account. - Responses: Stored alongside the query in the conversation record (if chat history is enabled). Responses include cited sources, see-also sources, and the certainty score. - Chat history: Stored in the primary database (Aurora PostgreSQL), encrypted at rest. Only retained when chat history is enabled by the account administrator. - Uploaded documents: Stored in a private AWS S3 bucket, encrypted at rest. Document text is extracted, chunked, and embedded in the vector database for semantic search within the conversation. Documents are automatically deleted 30 days after last access.	Noted - reference to use of data and sensitive storage in 7.54 & 7.55

Ref	Requirement	Response	Reviewer Comments
7.06 cont.		- Saved searches/matters: Conversations can be grouped under a "matter" for organisation. Matter data is stored in the primary database with the same encryption and access controls.	Noted
7.07	What are the default retention periods for chat history, engagements, saved searches/filters, and results? Are these retention settings configurable by the customer/administrator?	Chat history: Retained indefinitely by default when chat history is enabled. Enterprise accounts can enable automatic deletion after 90 days via an account-level setting (`auto_delete_chats`). - Uploaded documents Automatically deleted 30 days after last access. The deletion timer resets each time the document is accessed in a conversation. This retention period is not currently configurable per customer. - Matters/saved searches Retained for the lifetime of the account. - Responses/results Retained as part of chat history (same retention as above).	Noted
7.08	What capabilities are available to export user activity (e.g. queries, results, saved searches/engagements), and in what formats can this be provided?	Law Cyborg provides a comprehensive data export capability (GDPR Article 20 — Data Portability). The export is in a JSON format and includes: - User profile (name, email, account memberships, permissions) - All conversations with exchanges (user queries, AI responses, accuracy scores, cited sources) - Uploaded documents (title, summary, metadata) - Login history (IP, browser, device)	Noted
7.09	If deletion is enabled by customer admin roles, does this deletion extend to backups?	When customer data is deleted (either via automatic chat deletion or account closure), the deletion applies to the primary database and S3 storage immediately. Uploaded documents are also deleted from the vector database. Database backups (automated daily Aurora snapshots) retain deleted data until the backup retention cycle of 7 days completes. Backups are encrypted at rest (AES-256) and retained within the same AWS region.	Noted
7.10	How is customer data segregated in your multi-tenant environment (logical segregation controls and any additional safeguards for sensitive data)?	Customer data is logically segregated using a strict multi-tenant architcture.	Noted
7.11	Does any customer content flow into observability tools (e.g., DataDog, APM/logging), and how is sensitive data prevented from being captured in logs/traces?	Logging is structured to capture psuedonymised operational metadata, no user-generated content.	Noted
7.12	What are your audit log retention periods for both infrastructure logs (AWS CloudTrail) and application logs (Datadog)?	AWS Cloudtrail - indefinitely. Datadog - 7 days.	Noted

Ref	Requirement	Response	Reviewer Comments
Sources, Jurisdiction coverage and updates			
7.13	What jurisdictions (UK etc.) and content types are covered (legislation, case law, regulations, guidance, manuals), and what are the key exclusions/limitations?	We cover NZ (All areas of law), Australia (Tax Law) and UK (Tax Law). LC uses three key source types; Case law, legislation and commentary. In the UK this covers all courts and the Tax tribunals, speciifc tax legislation (including devolved jurisdictions) and all guidance, manuals and TIINS published by HMRC.	Noted
7.14	Is the model limited to only using approved sources? How is this enforced?	Yes. The model is strictly limited to approved sources. Enforcement occurs at two levels: 1. Retrieval layer: Search queries are filtered by namespace (document source) and category (practice area). Only documents matching the user's licensed practice areas and jurisdictions are returned. This filtering happens in the search controller before any results reach the LLM. 2. Knowledge base: The vector database contains only documents that have been ingested through Law Cyborg's controlled data pipeline. Documents must pass validation, deduplication, and quality checks before entering the knowledge base. There is no mechanism for the LLM to access external sources during legal research mode. In non-legal mode, web search results may supplement context, but this mode includes explicit disclaimers and does not provide legal citations.	Noted
7.15	How are sources ranked by the platform?	Cosine similarity between the user's query embedding and document chunk embeddings, practice area alignment, document type (legislation, case law, guidance etc), temporal decay.	Noted
7.16	How does the system detect when there are insufficient or low-quality retrieval results? How does the model react and signpost this to the user?	Each answer displays a "certainity score" which indicates to a user how much of the answer is derived directly from primary sources, vs generated by the AI. Further, as each source is shown in the response, including reviewed but not cited, a user can empirically determine the use of the right sources.	Noted
7.17	Are there potentially relevant sources of information that are not included in the model's review? How is this signposted to the user?	Responses flag and share assumptions made and always include a suggested follow-up location to look at.	Noted
7.18	Is there a time-lag between legal change and its availability in the solutions? And if so, what is the typical time-to-update after changes to UK legislation/regulations/guidance, and how is this communicated to the user?	Information is updated weekly over the weekend. So maximum time-to-update is 7 days for any source. This is shared during onboarding and demos.	Noted
7.19	Can the platform handle effective dates, superseded versions, and historical snapshots (e.g., "law as at" a specific date)?	Not currently.	Noted
7.20	If the same query is run multiple times, how consistent is the response? (Including sources used, citations listed and estimated certainty rating.) How is this measured?	LC's system prompt has a very low temperature setting, so identical questions receive very similar if not identical responses. As we don't train on client data, we cannot measure this across users.	Noted

Ref	Requirement	Response	Reviewer Comments
7.21	How does the model treat contradictory sources?	By using our heirarchy and ranking system, LC looks at the nexus between seniority and recency to determine relevant sources. If not, it will flag the inconsistency to a user for their judgment.	Noted
7.22	How does the system respond when the legal position is known to be unsettled?	The certainty score mentioned above would indicate lower confidence on this (as there is less source data to rely on).	Noted
7.23	Are citations generated using direct linking to retrieved passages or via post-generation matching of text to sources?	The retrieval pipeline identifies and numbers specific document chunks before they are sent to the LLM. The LLM is instructed to cite these numbered extracts using superscript notation. During post-processing, citation numbers in the LLM response are mapped back to the original retrieved chunks. Any citation number that does not correspond to a retrieved chunk is flagged as invalid and excluded. This ensures every citation in the final response links directly to a specific, verified passage in the knowledge base	Noted
7.24	Can the model return citations that are syntactically valid but semantically unrelated?	No, our weighting system has minimum thresholds for consideration of semantic alignment.	Noted
7.25	How are citations validated before being shown to users?	Only verified, numbered document chunks from the knowledge base are included in the LLM context. The LLM can only cite chunks that were retrieved and provided.	Noted
7.26	Can a response be generated without citations?	Only if a user selects "Non-legal" mode, which includes disclaimers. Users may use this mode for drafting general emails, or other non technical tasks e.g. generating a marketing plan.	Noted
Model operation and auditing			
7.27	How does the system rank the importance of each section/chunk of a document? (Both user inputted and legal source documents)	For legal source documents, chunking is section-based — each statutory section, case law paragraph, or guidance section becomes a separate chunk with metadata (section code, title, page number, ordering). Chunk importance is ranked by the semantic relevance score (embedding similarity to the user's query), weighted by the document source type and practice area relevance.	Noted
7.28	What is the context-window for a single prompt? (Including documents submitted alongside the prompt)	The context window for a single prompt includes: the system prompt (instructions, formatting rules, jurisdiction context), up to 30 retrieved document chunks, any relevant user-uploaded document chunks, and the user's question. The total context is bounded by the LLM's context window (Gemini 2.5 Pro supports up to 1 million tokens). There is no separate, smaller limit imposed on a single prompt beyond the model's native capacity.	Noted
7.29	What is the context-window for a single response?	The maximum response length is 8,192 tokens (approximately 6,000 words). This is configured at the application level and applies to all LLM providers.	Noted

Ref	Requirement	Response	Reviewer Comments
7.30	What is the context-window for a continuous conversation?	For a continuous conversation, the full exchange history (all previous user messages and AI responses) is included in the context sent to the LLM. This means the effective context available for new retrieval results decreases as the conversation grows. When the total context exceeds the model's capacity, the user receives an error message advising them to start a new chat. There is a configurable maximum number of documents per conversation (currently 20 uploaded documents). There is no hard limit on the number of exchanges, but the context window constraint serves as a natural boundary.	Noted
7.31	What happens when the context window is reached or exceeded? (Including both the quality of answers and the warnings issued to end users)	Users receive an error message and encouraged to start a new chat.	Noted
7.32	Can the system explain how a decision/output was reached, including the model's reasoning steps? How can the user view this and is there an audit trail that can be downloaded or viewed retrospectively?	Yes. The platform provides cited sources used in forming a response but not the chain-of-thought reasoning.	Noted
User workflow, prompting and accessibility			
7.33	Does the product provide structured prompting guidance/templates for common tasks?	This is currently being designed for release in mid May.	Noted
7.34	Can users annotate, compare, refine responses, and ask follow-up questions while maintaining context?	Users can ask follow-up questions at any time, or edit the original question. Soon, when LC outputs an email specifically (on request), in app text editing will be possible.	Noted
7.35	What accessibility features are currently supported (e.g., screen reader compatibility, keyboard-only navigation, colour contrast, text resizing/zoom, and captions/transcripts)? If any are not currently supported, what is the planned scope to address them and the target delivery dates?	Users can use speech-to-text for inputs and also supports response audio reader.	Noted
7.36	Is the training limited to navigation, or does it also include guidance on query and prompt creation?	Training covers both navigation and prompt creation. Onboarding includes an interactive in-app tour and email sequence. The help centre (help.lawcyborg.com) provides guidance on effective query formulation. Structured prompting templates for common legal research tasks are currently in development, with release targeted for mid-May 2026	Noted

Ref	Requirement	Response	Reviewer Comments
Document upload / attachments			
7.37	What file types (including documents and audio) can users upload or attach to queries, and what are the maximum file size limits?	Users can upload or attach these document types: PDF, DOC, DOCX, and RTF. Limits: Up to 20 documents per conversation. Each document up to 10 MB. For audio, LC supports speech-to-text for prompt inputs	Noted
7.38	Can customers configure or disable document retention settings?	Document retention (30 days after last access) is a platform-wide default and is not currently configurable per customer. Enterprise accounts can enable automatic chat history deletion (90-day retention) on request, but this applies to chat history only — not uploaded documents. The document upload feature itself can be disabled per account via feature flags.	Noted
7.39	Are uploaded documents processed in full, partially extracted (e.g., text only), or transformed before being sent to any AI/LLM services?	Uploaded documents are processed through a text extraction and only the relevant information, decided semantically, is shared with the LLM	Noted
7.40	Where are uploaded documents stored, how long are they retained?	AWS S3 in the relevant region, stored for up to 30 days of inactivity.	Noted
7.41	Who can access uploaded documents (end users, admins, support), and how is access logged/audited?	Only end users can access documents they uploaded within conversations they own. We do not enable access to the documents through our elevated JIT permissions.	Noted
7.42	What is the current roadmap for future integrations (e.g. document management systems), with the expected scope and target timelines?	Outlook by end of May 2026 iManage by mid 2026 - to enable uploading of docs and use of file management	Noted
Hosting, UK market readiness & data residency transparency			
7.43	Are there UK paying customers today? If Yes: Confirm where production data is hosted, processed, and backed up, and from where the platform is administered and supported (including privileged access locations).	Yes, data is hosted in Sydney (ap-southeast-2), administered from Sydney & Auckland. Migration planned for May.	Noted
7.44	If "Yes" to the above: Where and how is the current hosting location (Australia) explicitly disclosed to customers/prospects (e.g., contract/DPA, privacy notice, onboarding materials, website)?	https://lawcyborg.com/privacy & in the Subscriber Terms (available on sign-up and from a User's account)	Noted
7.45	Please confirm if there are any sales, marketing or product materials currently describe the service as 'UK-hosted' or 'UK-operated' currently for UK customers.	No	Noted
7.46	What is the committed timeline and scope for the AWS London region launch, and what dependencies/risks could delay go-live beyond May 2026?	15th May is publically committed timeline for AWS London region launch, this will include full data sovereignty. Risk is key-person dependency but buffer is built in. Existing users will be migrated in weeks following with no timeline currently committed.	Noted
7.47	Once the London region is live, will UK customer data (if there are any) be fully hosted and processed in the UK, or will any data continue to reside in / be processed from Australia (including backups and support access)?	UK data will be hosted & processed in the UK with full data sovereignty, support access will still be done from Australia/New Zealand via UK vpc.	Noted
7.48	If there are any existing UK customers at the time the London region goes live, how will migration be handled (opt-in vs mandatory), and how will customers be notified in advance?	Mandatory, UK emails will be emailed 2 weeks in advance, migration will be done in an off-peak maintenance window, frictionless experience for the user.	Noted

Ref	Requirement	Response	Reviewer Comments
7.49	Given daily backups are currently retained in Sydney, do you plan to support UK-region backups and disaster recovery for UK customers? If so, please outline the target architecture (backup location(s), DR region) and expected timeline.	Yes, London hosted, full data sovereignty, same RTO's.	Noted
Sub-processors, international transfers and contractual controls			
7.50	Which sub-processors process customer content (e.g. Hosting, LLM including prompts, uploaded documents, generated outputs) and which are 'corporate SaaS only'.	AWS - Core infrastructure, compute, database, storage, CDN. Google Cloud Platform (Gemini) - LLM provider, Stripe - Payment processor.	Noted
7.51	Where Law Cyborg (or any of its sub-processors) processes or accesses UK customer personal data outside the UK (e.g., Australia or the USA), what international transfer mechanism(s) are relied upon to ensure UK GDPR compliance (e.g., UK IDTA or EU SCCs with the UK Addendum), and can you provide the relevant executed terms/templates on request?	Law Cyborg relies on Standard Contractual Clauses and/or the UK International Data Transfer Agreement (IDTA) for transfers of UK personal data to countries outside the UK (including Australia). Data Processing Agreements (DPAs) are in place with all sub-processors that handle customer data.	Noted
7.52	How do customers access the current sub-processor list and receive notice of material changes (and is there an objection/termination right)?	The sub-processor list is reviewed quarterly and updated when processors are added or removed. Customers are notified of material changes to the sub-processor list. The current list is available on request via trust centre and referenced in the Privacy Notice and Subscriber Terms.	Noted
7.53	Do any sub-processors (or offshore staff) have access to customer environments for support/troubleshooting, and how is that access controlled and audited?	No. Sub-processors do not have direct access to customer environments	Noted
AI/LLM usage, data use limitations and retention			
7.54	Is any customer data used for model training, tuning, evaluation, or product improvement? If yes, describe what is used, under what conditions, and can customers can opt out?	No	Noted
7.55	What customer data is sent to the LLM provider (e.g., prompts, snippets, uploads, metadata), and what controls ensure it is not retained or reused beyond immediate processing?	The user's prompt & semantic chunk from document uploads are sent to the LLM provider. No customer account data, personal information, or metadata beyond the conversation content is sent to the LLM provider. Google's Zero Data Retention agreement ensures no customer data is retained or used for model training.	Noted
7.56	How are user interactions (feedback, clicks, saved engagements, corrections) incorporated into the retrieval system (e.g., vector database)? And is this tenant-isolated?	User interactions are not incorporated into the retrieval weighting system. The interactions are scoped to the users account.	Noted
AI safety guardrails and responsible use			
7.57	Are prompts and responses logged to enable oversight and investigation of misuse?	Only if chat history is enabled.	Noted
7.58	What controls are in place to prevent, detect, or appropriately handle harmful, high-risk, or sensitive queries? Including those that could reasonably lead to unsafe or adverse outcomes if acted upon?	The LLM provider has built-in safety features, along with Law Cyborgs constrained knowledgebase.	Noted
7.59	Can firms restrict usage to supported jurisdictions and domains (e.g., UK tax only), and does the tool warn users when queries fall outside supported scope?	Yes. Each customer account is licensed for specific practice areas and jurisdictions (e.g., UK tax only). The retrieval pipeline filters sources by the customer's licensed namespaces and categories, ensuring only relevant jurisdictions are searched. If a user asks a question outside their licensed scope, the system will not return results from unlicensed jurisdictions. The user's practice area selection (or automatic categorisation) determines which sources are queried.	Noted
7.60	What measures minimise inaccurate citations or hallucinated authorities, and how are sources presented to support verification?	The LLM can only cite numbered extracts that were retrieved from the curated knowledge base and provided in context. It cannot fabricate source references.	Noted

Ref	Requirement	Response	Reviewer Comments
Testing, errors and corrections			
7.61	What testing and quality assurance of the product has been performed specifically for the UK jurisdiction?	Prior to launch, our internal experts were able to test UK specific questions and conduct checks for accuracy and relevance of citations. Further, we have engaged with a number of tax professionals in the UK to trial and use the product, and provide feedback.	Noted
7.62	Has the system been tested to the point of legal misinterpretation of the model due to reaching the limits of the context-window?	The system handles context window exceedance by returning an explicit error message to the user, advising them to start a new chat. The LLM is not permitted to generate a response when the context window is exceeded — preventing degraded-quality answers. The maximum response length (8,192 tokens) is set to ensure responses complete within the model's capacity. Conversation history grows with each exchange; long conversations naturally approach the limit and trigger the error before quality degrades.	Noted
7.63	How are errors picked up by users tracked and remediated? Is there an escalation path for legally significant errors?	Users can report errors via the in-app "Give Feedback" feature, which captures a rating, comment, optional error type classification, and the full conversation history. Feedback is reviewed by the Law Cyborg Product team and helps curate our content	Noted
7.64	What per cent of queries have been marked as inaccurate by end users over the past: 3 months; 6 months; 12 months? (Or similar tracked periods)	3 months /6 months /12 months, 0.024%, 0.032%, 0.04% respectively. These are marked "Negative" therefore assumed inaccurate	Noted
7.65	Does the system update users if something that has previously been generated is later found to be inaccurate?	No	Noted

Ref	Requirement	Response	Reviewer Comments
Reliability, transparency and user assurance			
7.66	Does the platform communicate limitations of AI-generated responses and encourage user verification where appropriate?	Yes. The platform communicates limitations through: - Disclaimers: The platform is positioned as a research assistance tool for professional review, not autonomous legal advice. This is communicated in the Subscriber Terms, Privacy Notice, and onboarding materials. - Certainty score: Each response displays a certainty percentage, alerting users when source coverage is low. - Source transparency: All retrieved sources are displayed (cited, see-also, and reviewed-but-not-cited), enabling professional verification.	Noted
7.67	Does the platform outputs include confidence/quality indicators (e.g., source strength/recency), and how would users interpret them? If not, is this planned (with target release date)?	Every answer includes a "certainty score" (as discussed in 7.16), this enables users to determine whether their question is easily answered by existing sources, or requires AI interpretation, the higher the score the more confident a user is that the output is backed entirely by primary sources. The back-end algorithm already scores for recency and hierarchy of source to account for quality of sources. Again, providing clear citations and links to sources users can interpret the answer themselves. For some jurisdictions law, we have also developed a treatment system for case law to track citation strength.	Noted
7.68	How are the estimated certainty ratings generated? Are they based on a separate statistical model, or generated by the AI model itself? Can these be audited / explained to the user?	The certainty score is generated by a deterministic post-processing algorithm, not by the AI model itself	Noted
7.69	Do sources/case treatments show reliability signals (e.g., positive/negative citing), and how would users interpret them? If not, is this planned (with target release date)?	Sources are displayed with metadata including: document title, date, document type (legislation, case law, commentary), and origin URL. Sources are categorised as cited, see-also, or reviewed-but-not-cited. Positive/negative citing treatment indicators are displayed for Australia & New Zealand, United Kingdom is planned for Q3.	Noted
7.70	Have the estimated certainty ratings been tested and shown to be stable over time? How is the stability measured?	The certainty score is deterministic	Noted
Authentication, access control and account security			
7.71	Can customers use the platform with username/password authentication only, with SSO and MFA as optional add-ons? If yes, what mitigations exist (e.g., enforced MFA for admins, conditional access, risk-based controls)?	Yes. Customers can use the platform with username (email) and password authentication only. SSO (Microsoft, via SAML/OIDC) and MFA (TOTP via authenticator app) are optional and available at no additional cost.	Confirmed
7.72	Can administrators enforce MFA for all users (and separately for privileged/admin roles)?	When SSO is configured, MFA enforcement is controlled by the customer's identity provider (e.g., Microsoft Entra ID conditional access policies). For non-SSO accounts, MFA (TOTP) is available but is currently user-initiated — account administrators cannot enforce MFA for all users via the Law Cyborg platform.	Noted
7.73	Can customer administrators increase password requirements above the current enforced baseline (8 characters + upper/lowercase)?	Not currently.	Noted

Ref	Requirement	Response	Reviewer Comments
7.74	You noted no password history is retained. Is there a roadmap to introduce password history / password reuse prevention? If so, will it be platform-wide or configurable per customer?	Implementing password history and reuse prevention is under consideration but it not on the roadmap currently.	Noted

Ref	Requirement	Response	Reviewer Comments
Commercials, query limits and pricing mechanics			
7.75	Are research queries subject to any frequency, volume, concurrency, or fair-use limits by subscription tier, and how do these vary across plans?	Users have unlimited queries at all plan levels, with a corresponding Acceptable Use Policy, which covers botting, fair use, account sharing and other abuse. No other limits are applied.	Noted
7.76	What happens if a customer exceeds limits (throttling, additional charges, plan upgrade), and how is usage measured and communicated to customers?	We have not yet experienced a scenario where a customer has exceeded limits. Our policy is to discuss with a user and understand what has happened. We measure by number of prompts per month.	Noted
Future plans and upgrades			
7.77	Is there a confirmed timeline and scope for the planned branding and UI customisation?	Custom branding (corporate colour palette, logo upload) is planned for June 2026	Noted