

**SAN FRANCISCO SUPERIOR COURT**  
**RFQ 38-24-002-IT**  
**Custom Application Development for Public Use**  
**FAQs**

*NOTE: Similar questions will be phrased to capture all inquiries*

#	QUESTION	ANSWER
1	Can you give an extension of 1-2 weeks to submit proposal? <i>(Similar questions asked.)</i>	The updated timeline is in the revised RFP that will be posted May 13, 2025 or May 14, 2025.
2	If the published answers give rise to additional questions and needs for clarification, would the [Court] be willing to allow a very brief round of follow-on inquiries? Maybe a very brief 24–48-hour window?	The Timeline in the RFP states possible interviews.
3	Please provide an indication of the institutional roles that will be involved with proposal evaluation (ex. technical IT specialists, communications, etc.). This will help to ensure that bidders are targeting their proposals appropriately.	The Court always has subject matter experts on their RFP evaluation teams.
4	Who is the incumbent vendor currently providing similar services? <i>(Similar question received.)</i>	There is no incumbent vendor as this is a new need for the Court.
5	Are incumbent vendors eligible to participate in this procurement?	There is no incumbent vendor as this is a new need for the Court.
6	What was the annual spend for the previous year on this Project?	This project has consisted of several phases. Prior efforts during earlier phases won't provide a reliable benchmark for current spending. The phase of the project in this RFP is new.
7	If this is a new Contract, what is the annual Budget for this? Or can we at least get a sense of the anticipated budget ballpark? <i>(Similar questions received.)</i>	This will be a new contract if awarded and while the total funding is confidential, we expect the initial build to take 8–10 months, with a preferred average monthly cost cap within the range of \$20,000-\$30,000.
8	Is it anticipated that the resulting contract will take the form of a firm fixed price agreement?	While the final contract structure is still to be determined, we are currently operating under the presumption that there will be no overbilling or timeline overruns. A firm fixed price agreement aligns with this expectation and may be

		the preferred structure, depending on the proposed scope and deliverables.
9	Will this contract rather entail the provision of services on an hourly (T&M-style) basis?	No.
10	Is there a particular pricing format that bidders should use?	A pricing template was not provided; however, a line-item price format is preferred that shows most if not all costs separated out.
11	The RFP indicates that bidders should specify \$0 for any "line item" for which there would be no cost. Is there a particular list of such line items that the Court would want bidders to address?	No.
12	To the extent that the RFP does not present a clear and definitive solution scope (in terms of desired product functionality), are we correct in thinking that the Court simply wants bidders to provide a list of positions and corresponding rates?	The overarching scope is outlined in the RFP. A complete proposal should detail the technical approach, team qualifications, project timeline, and proposed cost as well as showcase how proposer would build a public-facing application that aligns with the Court's needs.
13	Who are the primary decision-makers/ stakeholders for approvals and for collaborating with the vendor, and what are their roles? ( <i>Similar questions asked.</i> )	The primary stakeholders collaborating with the vendor will be members of the Court's internal project team that includes individuals with both technical and business-oriented expertise. This team will serve as the key point of contact, providing guidance on user needs, functional requirements, technical feasibility, and overall project alignment.
14	How often are we expected to engage with the stakeholders (e.g., daily/weekly standups, demos)?	During the execution phase, the Court expects daily Agile standups to maintain momentum and ensure transparency. The Court's internal project teams will also hold weekly meetings with key vendor contacts to review progress and align on next steps. Additional meetings may be scheduled on an as-needed basis to address specific issues or decisions.

15	Work will be onsite or remote? (Similar questions received.)	The Court expect some on-site presence during the project kick-off-- primarily from team leads and senior resources. Over time, the Court may be comfortable with a remote setup, though the Court emphasizes the importance of having senior resources onshore and onsite when needed.
16	Are you open to a hybrid delivery model with a mix of offshore and onshore resources? (Similar question received.)	Yes, the Court is open to a mix of on-shore, near-shore, or offshore project teams. However, the Court expects key contacts to be onshore, (and ideally on-site during project kick-off).
17	Must all work be accomplished by US citizens; are there any specific visa or residency requirements (e.g., U.S. citizen, green card holder)? (similar question asked)	Per the Terms and Conditions of the Agreement, provision 10.9 Governing Law; Jurisdiction; and Venue states: "this Agreement and performance under it will be exclusively governed by the laws of the State of California without regard to its conflict of law provisions."  For offshore team members, work authorization in the U.S. isn't necessary.
18	The contract for building the application will approximately start from June 2025 up to March 2026. The first iteration is expected to be ready for go-live within eight months, with the remaining time allocated for refinements, issue resolution, and release management. Is this Project Period fixed for 10 Months or is the Court flexible based on the requirements?	The project period is expected to span 8–10 months, with go-live in the first 8 months and the remainder for refinements. While there is some flexibility, the Court anticipates adherence to this timeline, with a preferred monthly cost cap as stated in this FAQ form.
19	What are the core pain points and challenges you are trying to address through the Wayfinder application?	The Wayfinder application is envisioned as an informational tool that extends the Court's Self-Help Center services. It will provide factual, procedural guidance to self-represented litigants through an interactive roadmap tailored to their situation.  A built-in chatbot will address basic, recurring queries and help users navigate general processes. However, for complex or nuanced questions, the application will redirect users to existing service delivery channels, such as in-person assistance, LiveChat, or IVR.  Importantly, the application will not support electronic filings, nor will it

		integrate with internal court systems. Its sole function is to help users understand the process, log their progress, and receive factual, rules-based guidance.
20	What exactly is the intended outcome of this contract?	The intended outcome of this contract is the design, development, and deployment of a custom, standalone application that provides self-represented litigants (SRLs) with factual, procedural guidance through an interactive platform.
21	What platform or technology stack (if any) is currently in use or is the Court trying to purchase a new system (fresh implementation) with a particular scope of functional and non-functional capabilities? (Similar questions received.)	This will be a new system that does not exist today. The Court is intending to have a custom-built application with a specific and well-defined scope.
22	Is the Court trying to put in place a new master contractual vehicle through which the Court can avail itself of the contractor's services, in an unlimited and as needed manner, for a yet-to-be-determined scope of technical assignments that might arise at some point in the future?	No, this RFP is not intended to establish a master services agreement or an open-ended contractual vehicle. It is specifically and solely for the design, development, and deployment of a singular custom application as outlined in the scope of work.
23	How many awards does the Court expect to make in relation to this RFP; will there only be a single award; or is the Court rather intending to create a pool of prequalified service providers, who will then be asked to bid on individual task orders?	The RFP did not state multiple awards therefore the award, if any, will go to one vendor.
24	In relation to the RFP's apparent requirements, has the Court received any presentations/ demonstrations from any prospective providers?	No.
25	The RFP devoted substantial attention to matters of engineering. In some ways, the RFP was more clear and detailed about the desired engineering than it was about the desired scope of functionality. There is even some specificity on DevOps tooling (ex. Docker and Kubernetes). What are the concerns that prompted the inclusion of such engineering specificity in the guidance?	The RFP lists technologies as examples—not requirements. Given the Court's limited in-house development capacity, the solution must use widely adopted, standard frameworks to ensure easy maintenance, updates, and vendor transitions, supporting long-term sustainability and cost efficiency. Additionally, the specificity in the RFP stems from the need for the future solution to align with the Court's existing Information Architecture, ensuring coherence.

26	If the vendor will be working independently to deploy and support a solution in the cloud, and if most users will simply be accessing the solution via standard web browser and/or mobile device, why would such engineering decisions matter to the Court?	The RFP's technology stack is illustrative, not prescriptive. The specificity in the RFP stems from the need for the future solution to align with the Court's existing Information Architecture, ensuring coherence.
27	Will it be necessary to accommodate some level of technical collaboration, or is there some other reason the technical engineering decisions will be of consequence?	The need to accommodate some level of technical collaboration is not anticipated, but if such a need arises it will at most be with the Court's internal technical team.
28	Will the Court be providing any existing content, datasets, user pathways, or legal rules to seed the initial rule engine logic, or should vendors anticipate collaborating with subject matter experts to define these from scratch? (Similar questions asked)	The Court will provide foundational materials to support the development, including existing logic, documentation, and sample user pathways. Additionally, vendors will have the opportunity to collaborate with subject matter experts (SMEs) identified by the Court to refine.
29	Is the Court currently using any case management systems, APIs, or data sources that the Wayfinder application will need to publish/ expose/ integrate with for real-time updates, document sync, user validation, or functionality? (Similar questions received.)	Although the Court currently uses CMS solutions, the expectation is that this application will operate independently (without integrating with existing data sources) in Phase 1 but may be a consideration in future iterations.
30	For the dynamic rule engine, should vendors plan to implement a logic testing sandbox within the admin interface, or would the Court prefer a more traditional development/ staging environment for rule validation?	The Court is predisposed towards the admin interface route but is open to vendor recommendations, as long as the approach supports scalability, usability, and integration, and aligns with our long-term maintenance and cost objectives.
31	Will the Court provide testing or access to representative users (e.g., self-represented litigants or court staff) during the prototyping and user acceptance testing phases? (Similar questions asked)	While the Court cannot guarantee access to members of the public, including SRLs, at this stage, the Court will make a good-faith effort to facilitate access to representative users as appropriate. Court staff who have been designated as testers will be available to participate in prototyping and user acceptance testing. For the AI Based Chatbot, the Court will provide foundational materials to support the development, including existing logic, documentation, and sample user pathways. Additionally, vendors will have the opportunity to collaborate with subject matter experts (SMEs) identified by the Court.

32	Would the solution also need to handle other SRL tasks/ processes (ex. traffic citations, small claims, adoptions, etc.)?	The solution will not need to handle other SRL tasks/ processes in this iteration.
33	What is the scope of SRL functionality that is already available?	The Self-Help Center (ACCESS Center) supports Self-Represented Litigants through the following main service delivery models: in-person triage, IVR, LiveChat (with a human agent), and online intake forms for after-hours access. There is more information on the Court's website: <a href="https://sf.courts.ca.gov/access-legal-self-help-center">https://sf.courts.ca.gov/access-legal-self-help-center</a>
34	Would you list off the names of any products that are being used to deliver the SRL functionality.	Information about resources available is found on the Court' website: <a href="https://sf.courts.ca.gov/access-legal-self-help-center">https://sf.courts.ca.gov/access-legal-self-help-center</a>
35	Has the Court opted to pursue a course of custom development for any of the components that are currently being used to deliver the SRL functionality or are there any existing court-focused applications or platforms that have shaped or influenced your requirements or expectations for this project? ( <i>Similar questions asked.</i> )	The Court does not have any applications currently providing SRL functionality. The Court is looking for full development.
36	What aspects of the existing SRL functionality were custom built?	None.
37	What are the names of the vendors who have been involved with the implementation and sustainment of the existing functionality?	There are no vendors because currently there are no applications.
38	What is the general scope of SRL functionality with which each vendor has been directly and predominantly involved?	The Self-Help Center (ACCESS Center) supports Self-Represented Litigants through the following main service delivery models: in-person triage, IVR, LiveChat (with a human agent), and online intake forms for after-hours access. The intended outcome of this contract is the design, development, and deployment of a custom, standalone application that provides self-represented litigants (SRLs) with another platform to access self-help services.
39	Will any preference be given to companies that can document prior experience and involvement with SRL-related systems and functionality?	The RFP evaluation criteria lists on what the proposals will be evaluated.

40	If the SRL profile functionality is provided by an off-the-shelf product/ platform or was custom developed, what is the name of the product/ platform/ vendor? <i>(Similar question asked)</i>	No existing SRL profile creation functionality is in place.
41	If the SRL profile creation functionality already exists, will the awardee of this project necessarily need to develop parallel new functionality as a part of this project? Or would integration with the existing functionality be an alternative option?	No existing SRL profile creation functionality is in place. The selected vendor, if any, will be responsible for building this capability from scratch as part of the application. The profiles and data will be fully self-contained within the new public-facing system, with no integration or data sharing with existing Court systems in this iteration. The sole purpose is to provide better, more targeted informational support to the public.
42	Does the existing SRL profile application expose an API that would be sufficiently robust to support SSO, and any other functionality that would be required, in the event that a course of integration is pursued?	There is no existing SRL profile application or API, so integration is not applicable. As this tool is intended for public, unauthenticated users, enterprise-level Single Sign-On (SSO) is not relevant. Instead, the application will require a simple and secure login mechanism, such as OAuth-based authentication (e.g., email, phone, or third-party login) to support self-contained user profiles used solely for delivering personalized informational guidance.
43	Would authentication need to be integrated with an existing SSO system/ provider? If so, please explain (and include the name of the relevant SSO system).	This is an external application where the users create an account/profile to login which requires secure login capabilities.
44	Does the Court have existing SSO (Single Sign-On) infrastructure or requirements for integrating with their identity provider and if so, what agency is responsible for the management of the existing SSO? <i>(similar question asked)</i>	Custom application is an external application where the users create an account/profile to login which require secure login capabilities.
45	Would the existing SSO infrastructure support SAML and/or OAUTH?	The Court's existing SSO infrastructure is designed primarily for internal use, and while it may support SAML and/or OAuth protocols, it is not intended for public-facing applications like Wayfinder. For this project, vendors should propose a lightweight, secure login mechanism (e.g., OAuth-based email or phone authentication) suitable for public users, without relying on internal SSO systems.

46	If the existing SRL profile application does NOT expose an appropriate API, would the awardee of this project be expected to provide the necessary API(s)? Or would the Court address this shortcoming with the vendor who is/was responsible for the implementation and sustainment of the SRL profile application?	There is no existing SRL profile application or API, so there is no dependency on external systems or third-party vendors. The awardee of this project, if any, will be expected to design and implement the profile functionality as part of the new, self-contained application. This ensures that the solution remains isolated from internal Court systems and is fully tailored to support public users seeking factual guidance.
47	Does the Court, have any existing identity management systems in place that would need to be targeted for integration? If so, please explain.	There are no existing identity management integrations anticipated as this is planned as an external application for public usage.
48	Does the Court have preferences or constraints regarding the underlying cloud platform (e.g., AWS, Azure, GCP) for hosting the application and associated services? <i>(similar question asked)</i>	While the Court's preferred platform is AWS, the Court is open to recommendations for other Cloud Providers.
49	How many administrative users does the Court expect to manage and interact with the admin portal, and will there be distinct user roles or permissions to implement and a brief description of the functionality that should be available to each? <i>(similar questions asked)</i>	The Court expects up to 750 active external unique users in the first year. No more than 25 active internal user profiles are anticipated at any given time, spanning both standard and administrative roles. We envision four primary roles: public user, general internal user, read-only internal user, and administrator. Internal roles include general internal user, read-only internal user, and administrator. These may evolve based on project needs. Internal access can remain web-only, as staff are not expected to use internal functions on mobile devices.
50	Are there existing systems within the Court that will require integration?	This application will initially operate in isolation, aimed at enabling the public to access accurate information about court processes and track their own progress based on self-entered updates, so any data persistence would only have to be within the application.
51	Will the Court provide its own hosting environment (e.g., on Azure, AWS) or should the vendor propose and manage a secure cloud infrastructure that is included in the proposed pricing? <i>(Similar questions asked)</i>	The Court remains open to both deployment models—Court-hosted or vendor-hosted—and welcomes recommendations. A final decision will be based on factors such as cost, security, scalability, and support



		requirements.
52	What is the expected daily or monthly user count and their roles interacting with the system, including the mobile application. <i>(similar question asked)</i>	The Court expects up to 750 active unique users for the first year. The system should be designed to scale efficiently, but high-volume usage is not expected at this stage.
53	How many internal users are expected to access the platform in a given month for scalability planning?	The Court anticipates no more than 25 active internal user profiles at any given time, including both standard and administrative privilege levels, for scalability planning purposes.
54	Which platforms are prioritized: Android, iOS, web, or all simultaneously? Please rank them in order of your preference.	Android, iOS, and web are all equally prioritized and the webpage should be responsive UI for all form factors.
55	Are there specific integrations expected third party APIs/tools? List out the applications that THE COURT plans to integrate with.	There are no expected integrations at this moment in this iteration.
56	Do you have a list of existing infrastructure/application we should align with?	The Court prefers the application to be AWS-based, but there are no additional requirements regarding alignment with existing infrastructure or system integrations at this time.
57	Wayfinder System: Is this based on real-time data or static paths?	The Wayfinder system will use dynamic paths, with recommendations generated based on underlying logic provided by the Court.
58	Are there specific devices/ browsers/ platforms to be prioritized in testing? Please list your preference.	The application is intended to function as both a mobile and web application. While specific device and browser priorities are still to be determined, cross-platform compatibility will be essential during testing.
59	What is the expected training group size, and should the training be conducted virtually, on-site, or in a hybrid format?	The anticipated training group size will be small and the Court is open to vendor recommendations about being conducted virtually, on-site, or hybrid.
60	Are the current ecosystem and technology stacks/ frameworks listed in the RFP the ones we are expected to use, do you have a technology stack preference, or can we propose alternative relevant technologies? <i>(similar questions asked)</i>	The technology stack is illustrative, not prescriptive. The Court remains open to vendor recommendations as long as they align with the Court's existing Information Architecture, ensuring coherence.

61	What are the languages expected to be supported for this project both at the application level and L1/L2 support?	On the internal user-side, only English is in the scope. For the public facing application, for this iteration, English will be the primary language--with Spanish included if the scope permits.
62	What are the types of roles which are applicable in this application (applicable to both web and mobile)?	The Court anticipates four primary user roles for the application: public user, general internal user, read-only internal user, and administrator. These roles may evolve over time based on project requirements and implementation feedback. Internal roles are not required to function on mobile; they can remain web-only, as court staff are not expected to use the staff-facing components of the application on mobile devices.
63	Are there any specific modules envisioned as part of the solution?	Yes, the RFP outlines key functional modules expected in the solution.
64	What are your expectations / scope regarding response times, SLAs, and support hours? (Or is it just divorce?)	While the initial scope focuses on divorce-related pathways, the system should be designed with scalability in mind for future expansion.
65	Will any preference be given to companies that can document prior experience and involvement with divorce-related systems and functionality?	The RFP evaluation criteria lists on what the proposals will be evaluated.
66	What other systems exist that would include data on the status and disposition of respective divorce cases? Please list the systems by name as applicable and for each system, please also provide an indication of the sort(s) of information that the system would be storing relative to a divorce case.	This will be an independent application and is not expected to integrate directly with existing court systems.
67	In terms of the wayfinding functionality that would be expected, what is the list of key steps/ stages according to which the Court would want divorce cases to be tracked?	The Court will be providing the underlying logic and process flow that defines the key steps and stages for divorce cases. This logic will guide the wayfinding functionality, including how user inputs determine next steps, relevant documents, and procedural requirements. Vendors will not be responsible for defining the legal process itself, but rather for implementing the logic in a dynamic and user-friendly interface.
68	Can the Court please provide a diagram of the normal process through which each divorce case would flow?	This information will be made available after the contract, if any, is awarded.

69	Would there be any wayfinding functionality that would be required beyond the presentation/ tracking of divorce case status? If so, please explain.	Yes, wayfinding will go beyond case status tracking. The system will not provide live case updates, as it's not integrated with the Court's internal CMS. Instead, it will use user inputs to guide SRLs through next steps based on Court-provided logic, without filling out forms or accessing active case data.
70	Is the Court presently aware of any off-the-shelf products that could potentially satisfy the functional requirements of this contract?	At this time, the Court is not aware of any off-the-shelf products that fully meet the functional requirements outlined in the RFP.
71	What existing systems does the Court already have in place that are used to manage case information or desired functionality? Please list the systems and briefly explain the scope of functionality that each system already delivers. <i>(similar question asked)</i>	This will be an independent application and is not expected to integrate directly with existing court systems.
72	Does the Court have any other major systems in place that it might consider using to deliver the desired functionality? (ex. Salesforce, or some other major ERP implementation)	This will be an independent application and is not expected to integrate directly with existing court systems.
73	Would the operations of this contract need to be carried out in coordination with other technical development initiatives in San Francisco? If so, please provide some sense of the overarching context. (i.e., how many development teams involved, how many developers, what tech stacks, any critical coordination/ participation expectations, etc.)	No.
74	If the configuration of a role's permissions must be possible, what is the extent of configurability that the Court has in mind? Can an example of few rules be shared for better comprehension of the complexity of the system? <i>(Similar questions asked.)</i>	<p>The workflow must be dynamically configurable through a rule engine that allows administrative users to make changes without requiring a full redeployment.</p> <p>Adaptive Rule Logic: If court rules change (e.g., eligibility criteria are updated or new required documentation is introduced), we must be able to update the rule engine logic to reflect those changes with minimal friction.</p> <p>For example, if a user indicates they are filing a claim and have limited income, the system should recommend exploring a fee waiver application and begin asking relevant eligibility questions. The rule engine must support adaptive logic so that if court rules change—such as updates to eligibility</p>

		criteria or required documentation—the underlying logic can be easily modified with minimal friction.
75	Are there particular [role] permissions for which the Court would want to be able to adjust settings?	No. The Court does not need any detailed setting adjustments—just basic settings: color, font size, etc.
76	If configuration of a role's permissions must be possible, does this need to happen via an administrator interface or would it be acceptable if this could only be accomplished at the development level?	The admin should have permission to add/ remove internal users and change the internal role.
77	Will any preference be given to companies that are situated (headquartered) locally in San Francisco?	No.
78	Could you provide any additional detailed functional requirements or business requirement documents (BRDs) beyond those outlined at a high level in the RFP?	Additional project-related documents cannot be shared at this time to ensure fair competition and confidentiality. This application is intended to serve as an additional service delivery channel that complements the offerings of the Self-Help Center: <a href="https://sf.courts.ca.gov/access-legal-self-help-center">https://sf.courts.ca.gov/access-legal-self-help-center</a>
79	To the extent that the desired “wayfinding” functionality might necessitate integration with other systems, do the other systems all expose sufficient API functionality to accomplish the objectives?	The Court does not anticipate external integrations as the wayfinding will be driven by a rule engine using Court-provided logic. While no API dependencies are required, the Court welcomes vendor recommendations that could support future integration opportunities.
80	Would the responsibility for creation of the necessary APIs fall upon the awardee of this contract?	While no API dependencies are required, the Court welcomes vendor recommendations that could support future integration opportunities.
81	Please clarify the sorts of rules that the system would be managing. (Would these be legal procedure rules? Or rules for experience personalization? Or rules for service recommendations? Or rules for something else? Please explain what the Court has in mind and provide some possible examples of the sorts of rules that might need to be supported.)	For clarification, see this example: Initial Input: A user indicates that they are filing a claim and have limited income. Recommendation: The system suggests exploring a fee waiver application and begins asking eligibility-related questions. Adaptive Rule Logic: If court rules change (e.g., eligibility criteria are updated or new required documentation is introduced), we must be able to

		update the rule engine logic to reflect those changes with minimal friction.
82	Please explain the scope of envisioned tasks/actions that might need to be triggered by rules. Just emails and notifications? Some manner of data sharing? Please explain.	The rule engine is needed to guide the user step-by-step based on the questions answered and the Court does not need any email/ notifications rather it's a series of questions of questions asked and the path varies based on the answers provided by the users. For example, if a user indicates they are filing a claim and have limited income, the system should recommend exploring a fee waiver application and begin asking relevant eligibility questions. The rule engine must support adaptive logic so that if court rules change—such as updates to eligibility criteria or required documentation—the underlying logic can be easily modified with minimal friction.
83	What exactly is the scope of configurability that would need to be supported in relation to rules?	Configurability is needed for the dynamic nature of the workflow that will guide the users to different path based on their case type and answers to the questions. For example, if a user indicates they are filing a claim and have limited income, the system should recommend exploring a fee waiver application and begin asking relevant eligibility questions. The rule engine must support adaptive logic so that if court rules change—such as updates to eligibility criteria or required documentation—the underlying logic can be easily modified with minimal friction.
84	Will this system need to provide interfaces that would allow for manual management/administration of cases? If so, please explain the required scope of functionality in sufficient detail to enable responsible estimation.	This will be a self-help guide. No interface needed for manual management of cases.
85	Regarding existing system that already provides case tracking and management functionalities, who are the vendor(s) who have worked on the system(s) and how will the new system (i.e., the solution implemented under this contract) stand in relation to the existing system(s) and to what degree will there be functional overlap? <i>(similar question asked)</i>	The Court does not have an existing system. This is a new external self help guide application. This will not integrate with any internal case management systems.

86	Regarding existing system that already provides case tracking and management functionalities, is there already a sufficiently robust API (or set of APIs) that would allow the new solution (the system delivered under this contract) to execute all desired operations? And if not, would the awardee of this contract be expected to create the necessary API(s) for the old/existing system(s)?	The Court does not have an existing system. While no API dependencies are required, the Court welcomes vendor recommendations that could support future integration opportunities.
87	Regarding existing system that already provides case tracking and management functionalities, will any of the existing system(s) need to be sunset after this new solution comes online?	The Court does not have an existing system, so no sunset is needed.
88	Are there particular case management workflows that would need to be supported by the new system that is delivered under this contract? If so, please explain.	This is not a case management system, this is a self-help guide that the court is providing for the external users.
89	Does the Court envision a chatbot solution that is more oriented towards contextually aware, free form, natural language interpretation; or will the chatbot solution merely need to handle a more modest and limited scope of requests, conforming to a relatively fixed decision tree? <i>(Similar questions asked.)</i>	The Court envisions a context-aware, free-form LLM-based chatbot using fine-tuned or RAG methods, grounded in publicly available content and supplemental documentation. Accuracy is critical, and the chatbot will operate with a low temperature setting to minimize hallucinations and ensure responses are factual, consistent, and based on verified sources.
90	Should bidders assume that the dialog options can be hard coded; Or should bidders assume that the dialogue options will need to be dynamic, and stored in a database?	The rule engine is a core component of the Wayfinder system. It will drive dynamic decision-making and user guidance based on structured logic provided by the Court. Rather than static workflows, the rule engine must support flexible, input-driven navigation paths, where user responses determine the next steps, resources, or guidance presented.
91	Will the Court need an interface to directly manage the dialog options? If so, does [the Court] envision that this would be done through the solution provided under this contract? Or would the necessary dialog management functionality be arranged through a separate contractual agreement?	The Court does not have dialog management functionality.
92	Will the chatbot need to support multi-lingual interaction? If so, in which particular languages? <i>(Similar questions asked)</i>	For this iteration, English will be the primary language, with Spanish included if the scope permits.

93	With what channels would the chatbot need to be integrated? Just a primary web portal? Or would it be necessary to enable chatbot access in other contexts as well?	The chatbot will be integrated directly into the Wayfinder Application and must be accessible through both the mobile app and the web app. Currently, no additional channels are required beyond these two primary platforms.
94	For the AI chatbot, are there any preferred models or constraints regarding hosting?	The Court is open to both API-based integrations and local model deployments, including on platforms like Bedrock or LangChain, provided the solution aligns with our cost constraints and infrastructure capabilities.
95	For the AI chatbot, would use of third-party APIs such as OpenAI be acceptable if appropriate safeguards and compliance measures are in place?	Yes, the Court understands that leveraging a third-party model is the most practical approach.
96	For initial chatbot training, does the Court anticipate providing a source dataset (e.g., legal FAQs, procedural guides)	Yes, the Court will provide a comprehensive mix of pre-existing public information already published on our websites, along with additional documentation as needed. Ensuring the accuracy of the model is of utmost importance to us, and the Court is committed to going above and beyond to ensure its knowledge base is properly reinforced.
97	Should vendors propose a CMS for managing multilingual content, rules, and chatbot training data, or does the Court anticipate managing these elements directly through the admin interface of the application?	The Court is open to vendor recommendations for managing multilingual content, rules, and training data. While the Court expects to use the admin interface for some tasks, the Court is also open to a CMS if it offers clear benefits in scalability, usability, and integration, and aligns with long-term maintenance and cost goals.
98	Does the Court envision a phased rollout (e.g., MVP launch with subsequent features added) or is the goal to have all core features (Wayfinder system, chatbot, and admin portal) fully operational at initial go-live?	The MVP is expected to include all core features, but the Court is open to a phased delivery approach if timelines remain feasible.
99	Will the application need to support user-uploaded documents, and if so, are there specific file types, size limits, or document retention requirements the Court expects us to plan for?	In this first iteration, the Court does not anticipate user-uploaded documents as part of the core features—except potentially within the LLM-based chatbot, where users might upload files for review. In such cases, the Court is open to vendor recommendations on how those documents should be processed, stored, and managed.

100	Are there any analytics or reporting requirements (e.g., usage tracking, common issue trends, chatbot engagement stats) that should be included in scope for the first iteration?	Ideally yes but limited to key metrics—The Court would like to monitor DAUs, MAUs, and other standard KPIs such as application health, session duration, feature usage, and chatbot engagement to keep a pulse on overall performance with the ability to expand in the future.
101	Does the Court currently have internal APIs or data feeds available for integration into the Wayfinder or chatbot system?	The application will operate independently in its initial iteration, and the Court does not anticipate any direct integration with existing systems at this stage.
102	Is the AI chatbot expected to answer case-specific questions or only general legal queries? <i>(Similar questions asked.)</i>	The AI chatbot is intended to provide general, factual legal information only—not legal advice. Its responses should be based solely on publicly available court resources and documentation that we compile and supply for fine-tuning purposes.
103	Will the Court provide legal corpus data for chatbot training?	The AI chatbot is expected to provide general, factual legal information only—no legal advice. Its responses should be based on publicly available court resources and documentation that the Court curates and use for fine-tuning.
104	If the Court is providing training data for the chatbot, in what format will it be made available (e.g., PDFs, CSV, database export, API)? <i>(Similar question asked)</i>	The Court will provide documentation but does not anticipate the need of Database exports or APIs. Whatever legal information that is relevant and will be needed to ensure responses are accurate will be within privacy parameters.
105	Will the system need to connect with the Court's internal APIs (e.g., for court date lookups, document uploads, or user status sync)?	This will be an independent application and is not expected to integrate directly with existing court systems.
106	Is multilingual support required for both UI and AI chatbot responses? If yes, which languages are priority?	The priority for this iteration is the AI chatbot. English will be the primary language, with Spanish included if scope and resources permit. Full platform-wide multilingual support may be considered in later phases.



107	What is the scope of the chatbot—FAQ-style, conversational, task-driven? Will the chatbot integrate with live agents for escalations?	The Court envisions a context-aware, free-form LLM-based chatbot using fine-tuned or RAG methods, grounded in publicly available content and supplemental documentation. Accuracy is critical, and the chatbot will operate with a low temperature setting to minimize hallucinations and ensure responses are factual, consistent, and based on verified sources. While it is not expected to handle complex interactions or tasks, it should be able to escalate to a human agent when appropriate by directing the user to seek assistance from the Self-Help Center directly.
108	RE AI Chatbot: Request to provide volume of documents involved in training the AI to answer legal queries (volume of documents can be in overall number of pages / input PDF documents count)	While it's difficult to provide a precise document count or page volume, the AI chatbot will be trained primarily on publicly available resources hosted by the Court's Self-Help Center. The types of FAQs, guides, and procedural documents the Court intends to include are on the Court's website: <a href="https://sf.courts.ca.gov/access-legal-self-help-center">https://sf.courts.ca.gov/access-legal-self-help-center</a> .
109	What is the expected Chatbot accuracy% of the AI responses in terms of relevance to the question answered?	Given the narrow and focused scope of this iteration, along with extensive training documentation and expert feedback during development, the Court expects very high chatbot response accuracy—ideally above 95% in terms of relevance and factual correctness. There is no tolerance for inaccurate or misleading responses within this context.
110	Please just confirm that the deployed solution will merely be expected to link to (or embed) the referenced help center functionality?	The application is envisioned as an additional channel for the public to access the Court's Self-Help Services—not a replacement for any existing service delivery models. It will link to or embed the referenced Help Center resources where appropriate, serving as a complementary, informational tool to enhance public access.
111	Please confirm that the deployed solution will not, directly and in itself, be expected to provide functionality that would allow for live (or "in person") interaction?	Confirmed. The deployed solution will not be expected to provide live or in-person interaction capabilities. It is designed solely as a self-service, informational tool to guide users

		through procedural steps using features like a Wayfinder interface and chatbot.
112	Please confirm that the awardee will not be expected to provide any manner of live case guidance/assistance to prospective end users?	Confirmed. The awardee, if any, will not be expected to provide any form of live case guidance or assistance to end users. The vendor's role is limited to designing, building, and delivering the application, which will provide automated, factual, and rules-based guidance.
113	Please confirm that there will not be a need to provide live representation in any end-user facing context (chat or video conferencing)?	Yes, confirmed. The vendor will not be required to build any form of live representation in this iteration.
114	If the vendor and/or the desired solution does actually need to support live interaction with end users, what is a clear explanation of the functionality and level of support expected? (Please include information about the number of prospective users, the locations of the users, the desired experience flow/routing pattern(s) for incoming inquiries, and any needs for special technical capabilities (ex. session handoff/transferring, as might need to occur between customer service representatives).)	The desired experience would be a guided flow where users interact with an automated system (e.g., chatbot or virtual assistant) to answer questions and be directed to relevant resources. If escalation is needed, users should be routed to court staff or self-help center representatives.
115	Regarding multilingual support, are there target languages beyond English that the Court would like to prioritize during initial rollout? ( <i>Similar question asked.</i> )	Beyond English, the Court would like Spanish to be a focus of our multilingual efforts. However, this is considered a "nice to have" rather than a strict requirement for the initial rollout. Support for additional languages is within scope for future phases.
116	Initially assume left-to-right languages only?	Yes.
117	Should vendors plan to incorporate human-reviewed translations or rely on automated systems? ( <i>Similar question asked.</i> )	The Court is comfortable with using automated translation systems, provided there is a human-in-the-loop process to review and validate translations during the initial phases. Over time, we expect the process to become more autonomous, guided by feedback and iterative improvement.
118	Should vendors anticipate setting up CI/CD pipelines and managing infrastructure post-deployment, or will the Court assume responsibility for hosting and maintenance after go-live? ( <i>Similar questions asked.</i> )	The Court expects vendors to establish a CI/CD pipeline within our environment and provide comprehensive post-deployment support, including knowledge transfer to enable eventual in-house infrastructure management. Ongoing communication for issue

		resolution is required. Extended post-delivery support is valued and may enhance a vendor's proposal.
119	What are the expectations and timeline for ongoing support and maintenance after the system is implemented?	While the Court aims to eventually manage infrastructure independently, the Court expects vendors to provide well-planned multi-level post-deployment maintenance and support.
120	Is there an existing application from where historic/ current data is expected to be migrated? If yes, is data migration in scope? <i>(similar question asked)</i>	There isn't an existing application—this project is to build a new custom application.
121	Would there be any specific interoperability standards applicable to the data? If so, please explain.	No specific interoperability standards are currently required, as this is a standalone application with no planned integrations. However, the solution should be designed with flexibility in mind to accommodate potential future standards or integrations if needed.
122	Does the Court anticipate that any part of the user data collected through the Wayfinder platform (e.g., profiles, case progress) will need to be persisted and retrieved across different court systems or departments?	This application will initially operate in isolation, aimed at enabling the public to access accurate information about court processes and track their own progress based on self-entered updates, so any data persistence would only have to be within the application.
123	Will it be necessary for the system deployed under this contract to share data to, or otherwise trigger events within, other downstream systems? If so, please explain the nature of the system interactions that will need to transpire.	No, the system deployed under this contract is intended to operate as a standalone application. It will not be required to share data with or trigger events within any downstream systems.
124	In terms of personal account/ case data, please clarify what is the scope of personal information that would need to be captured?	The system will collect only self-supplied, case-related details—such as general information about the user and spouse, key dates, property, and children—to guide users through the divorce process. These inputs will be gathered through a structured Q&A flow based on a logic tree already developed by the Court.
125	In terms of personal account/ case data, please clarify are we merely collecting self-supplied data points?	Yes, in this iteration, the system will only collect self-supplied data points provided voluntarily by users to guide them through the workflow. This may include basic demographic and case-related information necessary for delivering relevant next steps.

126	In terms of personal account/ case data, please clarify will the system need to handle/ capture official records (e.g. health, residency, special education, or guardianship)? Please explain.	No, the system deployed in the project is currently intended to operate as a standalone application for this iteration.
127	In terms of personal account/case data, please clarify what aspects of the collected data would [the Court] deem sensitive?	The system will collect only self-supplied, case-related details—such as general information about the user and spouse, key dates, property, and children—to guide users through the divorce process.
128	Are all modules required in both the web and mobile versions?	Yes, all public-facing modules will be required in both the web and mobile versions to ensure a consistent user experience across platforms. However, the admin and court user interfaces will be web-based only and are not required to be accessible via the mobile application.
129	Is there a required level of offline functionality or resilience for mobile users (e.g., form entry, progress tracking), or will the app be expected to operate only in an online state?	For the first iteration, the Court is comfortable with the solution operating in an online state.
130	Is there a particular minimum scope of mobile devices that would need to be supported and web (browsers/versions)? <i>(Similar question asked.)</i>	The Court is open to vendor recommendations, but the solution must ensure broad compatibility with older iOS and Android devices, given our user base may not have the latest technology. Similarly, web compatibility should extend to commonly used browsers.
131	Does the solution necessarily need to entail two separate native mobile applications (for iOS and Android)? Would the Court be prepared to instead allow mobile functionality to be delivered in the form of response web pages, or progressive/hybrid application functionality?	The Court remains open to vendor recommendations as long as the solution provided aligns with our cost constraints and infrastructure capabilities.
132	Are there any accessibility standards beyond WCAG 2.1 AA that the Court expects the application to comply with, particularly given its public-facing nature? Are there specific audit processes planned? <i>(Similar questions asked.)</i>	At this time, there are no additional formal standards required beyond WCAG 2.1 AA. However, given the public-facing nature of the application, the Court strongly encourages incorporating best practices for inclusive design. There are no Audit Processes planned at this moment.
133	Bidders will only be responsible for technical configuration and layout issues?	The Court will provide the required documents for the project. The bidder will be responsible for design, development and deployment.

134	Bidders will not be responsible for textual content issues, such as linguistic readability?	The Court will provide the required content for the project.
135	Is the final solution expected to be exclusively cloud-hosted, or will any hybrid/on-premises components be required?	The first iteration of the application is expected to operate solely in an online/cloud-hosted environment. The Court does not anticipate hybrid or on-premises components at this time.
136	Is there an anticipated project extension opportunity beyond the initial go-live? If yes, under what circumstances?	As of now, there is no planned extension beyond the initial go-live phase of the project.
137	What are the expected system hours of operation (e.g., 24/7, 9–5 PST)?	24x7.
138	What level of support coverage is required (e.g., 24x7, 18x5, 24x5)?	The Court envisions a 12x5 support coverage model (Monday to Friday, 6:00 AM to 6:00 PM PT) during the initial stabilization period following deployment. Support expectations beyond that phase are negotiable and will be defined based on system needs and performance.
139	Will the vendor be expected to provide ongoing maintenance post-deployment? If yes, what is the expected duration (e.g., 1–3 years)?	While the Court aims to eventually manage infrastructure independently, we expect vendors to provide well-planned multi-level post-deployment maintenance and support.
140	Is there a defined Service Level Agreement (SLA) or escalation protocol that the vendor must adhere to?	Vendors should propose SLA and escalation procedures consistent with industry standards. These will be reviewed and finalized during contract negotiations.
141	Will the contract allow for amendments or change orders post-award, especially if new modules or integrations are introduced?	As of now there are no amendments to the current scope and none are anticipated.
142	How many internal Court users / role types (staff, admin, support agents) will need access? ( <i>Similar questions asked</i> )	The Court anticipates around 25 internal users in the first year, across roles such as admin, read-only, and support. The Court anticipates four primary user roles for the application: public user, general internal user, read-only internal user, and administrator. These roles are expected to evolve over time based on project needs and implementation feedback.

143	Is there any requirement to interface with hardware-based identity tools such as RSA SecureID, biometrics, etc.?	There are no current plans to interface with hardware-based identity tools in the first iteration.
144	Is the vendor free to use cross-platform mobile frameworks (e.g., React Native, Flutter), or are native iOS and Android apps mandatory?	Vendors are free to use cross-platform frameworks provided the solution delivers a responsive and accessible experience across devices.
145	Are there any specific compliance standards beyond NIST SP 800-53, such as CJIS, HIPAA, or SOC 2, that the solution must meet?	At this time, NIST SP 800-53 provides the primary compliance baseline. Additional standards may be considered in future iterations based on feature expansion.
146	Could the Court share expectations around backup, disaster recovery, and data retention, including RTO (Recovery Time Objective) and RPO (Recovery Point Objective)?	The Court expects the RTO/RPO to be around one hour. These can be finalized before go-live.
147	Are there required load performance benchmarks (e.g., concurrent users, latency) the system must be tested against?	The Court expects vendors to design for moderate concurrency and responsiveness under public usage. Specific benchmarks can be refined during the technical design phase.
148	Will the Court provide access to existing security policies or frameworks, or is the vendor expected to propose its own?	This is an external facing application and vendors are expected to propose baseline security policies aligned with best practices. There would be limited interaction, if any, with internal systems.
149	What is the expected frequency of functional and security updates post-launch?	Initially the Court expects updates once a month however, the frequency could increase to quarterly.
150	Is the vendor expected to define and follow a formal change management process (e.g., ITIL, DevOps-based)?	Yes, the Court expects vendors to define and follow a change management process that aligns with DevOps or ITIL principles and supports long-term maintainability.
151	Are there any UI/UX guidelines, court branding assets, or content templates that must be followed?	There are no formal UI/UX guidelines, branding assets, or content templates currently in place. Vendors are expected to follow established best practices for inclusive, accessible, and user-centered design throughout the application.
152	For the admin user interface, do we have any approval/hierarchy flow that needs to be implemented? If yes, then please share more details on this.	Yes, an approval and hierarchy flow will be required for the admin interface. The system should support configurable approval workflows, both at launch and post-release. Specific details will be provided after project kickoff to guide implementation.

153	Regarding the development of the Case Management dashboard, you have mentioned developing customs or use the existing system. Here our presumption is that you are referring to your existing system. If so, then do the existing systems have API that we can use to get relevant data of Case?	In this context, the Case Management Dashboard is a simple, user-facing interface that provides a basic summary of key case details directly from the application. It is not envisioned to be integrated with the court's main Case Management System (CMS) and will only display data already available within the application—we don't believe external APIs are needed in this context in the first phase.
154	Regarding development of Rule Engine for Personalized Recommendation, could you please share a few use case scenarios to help us better understand this feature, along with the key parameters we should consider for generating the appropriate output?	The goal is for developers to implement logic as a rule engine that supports easy updates to reflect changes in court rules or procedures—without requiring a full redeployment. For example, if a user indicates they are filing a claim and have limited income, the system should recommend exploring a fee waiver application and begin asking relevant eligibility questions. The rule engine must support adaptive logic so that if court rules change—such as updates to eligibility criteria or required documentation—the underlying logic can be easily modified with minimal friction.
155	Do you require audit logs and user consent tracking for compliance purposes? Are there specific security standards (e.g., ISO 27001, SOC2) we must follow for storing Audit logs?	The Court requires user consent tracking. Audit logs are up for discussion.
156	What is your preferred data retention/ deletion policy for chatbot interactions/ logs (e.g., retention period, anonymization)? <i>(Similar question asked.)</i>	The Court has non-confidential agreements in place with users, and the data policy will align with those agreements. User profiles will have an initial retention period of 7 months, with the option for users to extend their profile validity if needed. After this period—if not extended—all associated data, including past interactions and progress, will be deleted and not retained.
157	Are there regional dialects or local legal terms we need to consider for multilingual users?	No, the Court does not support regional dialects.

158	Do you require GDPR/CCPA compliance for all translated data and interactions processed via third-party APIs (e.g., Google Translate/AWS Translate)? <i>(Similar questions asked.)</i>	Since the Court is using LLMs, the Court believes there is no need for a third-party API.
159	Is legal terminology translation accuracy critical, or will a generic translation service suffice?	The terminology should be within the public source of information, which will be provided during the training.
160	Should the chatbot detect and translate language dynamically, or will users pre-select their preferred language?	For this iteration, the language is pre-selected. English will be the primary language, with Spanish included if the scope permits.
161	Is there a preference for using pretrained models (e.g., BERT, RoBERTa) for legal domain Q&A instead of training from scratch?	The Court welcomes vendor recommendations. The Court remains open to working with all major models.
162	Who will validate the accuracy and legality of model-generated responses through AI?	<p>The AI model must rely strictly on vetted sources of truth, which will be provided and maintained by the Court. During the pilot phase, our subject matter experts will validate outputs and provide structured feedback to refine performance.</p> <p>Post-launch -</p> <p>Users should be able to report inaccuracies directly within the system to support ongoing quality control.</p> <p>Importantly, the Court cannot monitor each response for legal validity, as this is for self-help purposes.</p> <p>The model must operate at a low temperature to minimize variability and must not offer legal advice or recommendations—it should strictly present factual information derived from approved sources.</p>
163	Would you prefer an open source, on-premises solution like Rasa for better data control and local NLP processing with legal intent recognition? <i>(Similar questions asked.)</i>	The Court welcomes vendor recommendations. The RFP's technology stack is illustrative, not prescriptive and the Court remains open to working with all major models.
164	Are there any data privacy restrictions that limit the use of cloud-based LLM APIs (e.g., OpenAI)?	The Court has non-confidential agreements in place with users, and the data policy will align with those agreements. A cloud-based LLM is expected to be sufficient for this use case.



165	Is fine-tuning GPT-4 essential, or would prompt engineering using the base model suffice?	The model will be trained solely on publicly available Court validated information and curated documentation provided by the Court.
166	How is the court currently using SMS reminders in your workflows, and what improvements or changes would you like to see?	The Court does not currently use SMS reminders in relations to Self-Help. In the future state, reminders are expected to be delivered via SMS, email, or push notifications through the mobile app. The final decision will be based on a review of the intuitiveness and feasibility of each option.
167	Approximately how many users do you anticipate will need administrative or content editing permissions?	Around 5-10 admin users is anticipated will need administrative or content editing permissions.
168	Are there any specific privacy concerns regarding AI use, especially in relation to sensitive user data? Do you currently have policies or guidelines addressing this?	The model should not access, store, or process any confidential case information. All data used for training or reference will comply with the Court's data privacy standards and be non-disclosure exempt. Relevant documentation will be shared with the vendor to ensure response accuracy within these constraints.
169	If a user is logged in, would you want the AI to retain a memory of their previous interactions to improve continuity and personalization?	The Court is open to vendor recommendations on retaining conversational context for logged-in users, provided it aligns with our data privacy policies, does not involve personally identifiable or case-specific information, and remains non-disclosable under public records laws.
170	Will this application also serve as a repository for commonly used documents that are accessible to end users?	No, the application will not serve as a document repository or support e-filing in this iteration. However, it should be able to guide users to the appropriate resources by linking to or fetching downloadable forms from the Court's existing public-facing websites.
171	Please describe the current Case Management Dashboard including the technology it is built on and any current pain points.	Although the Court currently uses CMS solutions, the expectation is that this application will operate independently—without integrating with existing data sources—in Phase 1 but may be a consideration in future iterations.

172	How is this project being funded?	The court is not compelled to disclose the source.
173	Are there any hard deadlines (e.g., go-live dates, funding expiration, or legal mandates) that must be met during the project?	The Court expects the initial build to take 8–10 months.
174	Does the Court anticipate a phased rollout or a single deployment?	The current RFP is for Phase 1 with single deployment for external users; multiple deployment is allowed for internal testing.
175	Are there any soft deadlines or target timeframes the Court has in mind for project milestones such as kickoff, discovery, development, testing, and deployment?	Yes, the Court expects multiple project milestones and soft deadlines.
176	Looking ahead to one year post-launch, what specific outcomes or indicators would you consider essential for this project to be deemed a significant success?	Success will be measured by high user engagement, reduced staff burden, reliable system performance, and improved user confidence in navigating complex legal processes without a lawyer and with greater ease.
177	What would your ideal post launch support look like? Do you value flexibility in the arrangement? If so, what options would you most be interested in? <i>(Similar questions asked.)</i>	<p>The Court envisions 12x5 support (Monday–Friday, 6:00 AM–6:00 PM PST) during the initial stabilization period. Support beyond that is flexible and negotiable, based on evolving system needs.</p> <p>While the Court aims to eventually manage infrastructure independently, we expect vendors to propose a well-structured, multi-tiered post-deployment support plan. Strong post-delivery support is favorably weighted in proposal evaluations.</p>
178	Can you please explain in more detail the requirement for “Integration with External Services” specifically the API connections with self-help centers? Are these existing platforms/systems with already established APIs? Is the integration expected to support data exchange only, or also include authentication, user interface components, or real-time interaction?	“Integration with External Services” in the RFP reflects a possible future-state consideration, not a requirement for the current scope. In this current iteration, the application will operate completely independently, with no integration or interaction with existing Court systems, platforms, or APIs. The Court will issue a revision to the RFP to clarify.
179	In reference to “Rule-changing system (for modifying app rules and configurations dynamically),” can you give examples of how this would work in your ideal solution? What are you looking to customize with it after the software is delivered?	The workflow must be dynamically configurable through a rule engine that allows administrative users to make changes without requiring a full redeployment. For example, if a user indicates they are filing a claim and

		have limited income, the system should recommend exploring a fee waiver application and begin asking relevant eligibility questions. The rule engine must support adaptive logic so that if court rules change—such as updates to eligibility criteria or required documentation—the underlying logic can be easily modified with minimal friction.
180	Can you please clarify what the scope and process of the Third-Party Risk Assessment will entail, including what areas will be evaluated (e.g., cybersecurity, financial stability, data privacy), and whether the vendor will be required to cover any associated costs?	Third-Party Risk Assessment may be required at the Court's discretion for security assurance, with a focus on cybersecurity and data privacy. There is no cost to the vendor.
181	Can you provide functional requirements or specific use cases for the AI Chatbot and Rule Execution Engine components? <i>This will help us better understand the user interactions, data flow, and intended outcomes.</i>	Initial Input: A user indicates that they are filing a claim and have limited income. Recommendation: The system suggests exploring a fee waiver application and begins asking eligibility-related questions. Adaptive Rule Logic: If court rules change (e.g., eligibility criteria are updated or new required documentation is introduced), we must be able to update the rule engine logic to reflect those changes with minimal friction. <a href="https://selfhelp.courts.ca.gov/divorce/start-divorce">https://selfhelp.courts.ca.gov/divorce/start-divorce</a>
182	Are there any performance expectations for the AI Chatbot and Rule Execution Engine?	Yes, the AI Chatbot and Rule Execution Engine are expected to perform with minimal latency and provide reliable, consistent responses.
183	Can you expand on the Visualization requirement (C.1.ix of the Attachment A – Requirements)?	The visualization requirement refers to creating a clear, intuitive visual interface for the Wayfinder—not just a progress bar. It should visually guide users through steps using simple, accessible design elements. Low-fidelity mockups may be shared at later stages to illustrate the intended user experience.
184	Is a built-in Content Management Systems also expected as part of the deliverables?	A built-in Content Management System is not strictly required, but the Court is open to vendor recommendations if it offers clear benefits for scalability, usability, and long-term maintenance.

185	What is the ambit of "user-specific legal circumstances"? i.e. Are there any legal documents that should serve as a base for rules definition?	Refer to this link as an example:  <a href="#">Start a divorce case   California Courts   Self Help Guide</a>
186	What is the size and format in which the legal data shall be made available?	While an exact volume is difficult to quantify, the legal content for the AI chatbot will come primarily from publicly available resources on the Court's Self-Help Center website: <a href="https://sf.courts.ca.gov/access-legal-self-help-center">https://sf.courts.ca.gov/access-legal-self-help-center</a> . The data will be provided in computer-readable formats such as PDFs, Word documents, or Excel files.
187	What is the estimated number of these self-help centers?	The Court has one self-help center.
188	How many different types of such self-help centers integrations are in scope?	At this stage, only the divorce self-help center integration is in scope for this phase of the project, however, the system should be designed with scalability in mind for future expansion.
189	The RFP mentions integration with a self-help center and potentially an existing admin portal/case management dashboard. Are there other known Court systems (e.g., e-filing portals, identity management services) that this public-facing application might be expected to interface with, even if through hyperlinks or SSO in the initial phase?	The application is expected to operate independently in its first iteration, with no integration to existing court systems such as e-filing portals or identity management services. It is a public-facing, external tool where users will create their own accounts or profiles. While secure login capabilities are required, SSO will not be implemented at this stage. Any references to existing systems (e.g., self-help centers) will be handled through hyperlinks only, not system-level integration.
190	Please provide the anticipated number of users and expected load benchmarks for key features such as the AI chatbot, notifications, or rule engine (e.g., users per second or maximum concurrent sessions)? And should the system support only San Francisco County or be designed for potential statewide expansion? <i>(Similar questions asked.)</i>	The Court anticipates low load, around 30-40 concurrent sessions, but open to scaling on the public facing application. In the court user side, the Court anticipates around 20 active sessions but the rule-engine will only be admin-managed so that will be lower. This solution is only being designed for San Francisco but making the solution scalable will be beneficial.
191	Should accessibility testing include assistive technologies (e.g., screen readers like NVDA or JAWS) on both web and mobile?	There are no plans in Phase 1 for accessibility testing to include assistive technologies.

192	Should the rule engine be extensible for future legal domains beyond divorce (e.g., claims, evictions)?	Yes, the Court wants it to be scalable to support possible future expansions.
193	Regarding the "Testing framework for validating rule modifications," would the Court expect this to include features like rule simulation, a staging environment for rules, and an audit log of changes before rules are pushed to the live application?	The Court is open to vendor recommendations for managing rule changes. as long as the approach supports scalability, usability, and integration, and aligns with the Court's long-term maintenance and cost objectives.
194	Could the Court provide any target Service Level Agreements (SLAs) for application uptime and performance (e.g., response times for key features) that the deployed solution will be expected to meet?	Vendors should propose SLA and escalation procedures consistent with industry standards. These will be reviewed and finalized during technical discussions.
195	How do you currently track case timelines and filings?	The Court uses internal tools to track cases, but this public-facing application will be independent, relying solely on user-provided inputs for tracking progress.
196	What documents or evidence do you need instant access to during hearings?	The Court does not require instant access to documents during hearings. This tool is solely intended to provide public-facing information and guidance and will not be used to support or participate in court hearings. It will operate independently of internal court systems and processes.
197	How do you manage client updates and correspondence with the court?	The Court uses internal tools to track cases, but this public-facing application will be independent, relying solely on user-provided inputs for tracking progress.
198	What security and privacy features do you expect regarding confidential case files?	Confidential case files will not be accessed, stored, or processed by this application.
199	How do you prepare for a court session, and what information do you need quick access to? ( <i>Similar questions asked.</i> )	This application is not envisioned to help with court sessions.
200	How important is offline access to documents?	The application will exist in an online-only state in its first iteration
201	Beyond general GDPR/CCPA compliance, are there specific California Judicial Branch data retention schedules or policies that must be implemented for	User profiles will have an initial retention period of seven months, with the option for users to extend their profile validity if needed. After this period—if not

	user data, interaction logs, and any documents potentially stored or referenced by the application?	extended—all associated data, including past interactions and progress, will be deleted and not retained.
202	For the Divorce Process & Wayfinder System, we identified the following steps for the divorce process: Residency Check, Filing Petition (FL-100), Serving the Spouse, Respondent's Answer (FL-120), Temporary Orders, Financial Disclosures (FL-140/142), Negotiate Settlement, Trial (if contested), Final Judgment (FL-180). Does the Court confirm these steps, or are there additional/modified steps we should include?	These steps are partially correct but all the steps to a Divorce are not relevant for proposal submission.
203	Should the admin panel allow dynamic editing of rules (e.g., residency duration dropdown with options like 4/6/9 months) to reflect legal changes? Are there predefined rule templates the Court requires?	Yes, the workflow must be dynamically configurable through a rule engine that allows administrative users to make changes without requiring a full redeployment. The rule engine must support adaptive logic so that if court rules change—such as updates to eligibility criteria or required documentation—the underlying logic can be easily modified with minimal friction.
204	Is the Court's case management system accessible via API for fetching court dates/filing deadlines? If yes, will API documentation be provided?	The application is expected to operate independently, without integration to existing court systems in its first iteration.
205	Legal Status & Case Tracking: How should the application track legal status updates? Is there an existing API, or will this rely on user input?	The system will not provide live case updates as it's not integrated with internal CMS. Instead, it will use user inputs to guide SRLs through next steps based on Court-provided logic, without filling out forms or accessing active case data.
206	How will the 'Link to Self-Help Center' function? Will it connect to a live chat, require a user login, or redirect to a static webpage? Is there an existing ticketing system or API for handling self-help queries?	The 'Link to Self-Help Center' will redirect users to existing service delivery models via information prompts and hyperlinks. It will not connect to a live chat or integrate with a ticketing system or API.
207	Please clarify if AI-Powered Chatbot, Appointment Scheduling requirement is for the mobile app to access the court's calendar so that users can schedule appointments on the court's calendar?	No, there is not access to the court's calendar. The users will be directed to existing service delivery models where they can connect with court staff or subject matter experts through information prompts and hyperlinks.

208	If the awardee will decisively need to create new functionality for SRL profile management, does the Court have any existing identity management systems in place that would need to be targeted for integration? If so, please explain.	The Court does not have an existing system. This is a new external self-help guide application. This will not integrate with any internal case management systems.
209	Are there specific performance requirements (e.g., transactions per second, maximum latency)?	There are no formal performance benchmarks at this stage, but the system should deliver a smooth and responsive user experience under typical usage loads.
210	Are there restrictions on using external AI services?	The Court prefers a context-aware, free-form LLM-based chatbot using fine-tuned or RAG methods. Third-party models are acceptable and will be evaluated based on accuracy, cost, privacy, and performance. All major LLMs will be considered.
211	Does the Court prefer an existing rules management system or a custom engine? Any format preferences (decision tables, JSON/YAML)?	This is a new system, and the Court is open to vendor recommendations—whether leveraging an existing rules engine or developing a custom solution. The chosen approach, if any, should support scalability, usability, integration, and align with the Court's long-term maintenance and cost objectives. There is no strict format preference.
212	What is the technical skill level of staff using the no-code interface?	Non-Technical Staff users are the skill level of staff using the no-code interface.
213	What capabilities are expected (e.g. input hypothetical scenarios to simulate rule outcomes)?	Initial Input: A user indicates that they are filing a claim and have limited income. Recommendation: The system suggests exploring a fee waiver application and begins asking eligibility-related questions. Adaptive Rule Logic: If court rules change (e.g., eligibility criteria are updated or new required documentation is introduced), we must be able to update the rule engine logic to reflect those changes with minimal friction.
214	Should the app display simple milestones or the full official case record?	The application is expected to show personalized milestones.

215	Should the app enable scheduling in-person appointments or simply display referral info?	The application is expected to display referral information only in its first iteration.
216	Regarding notification triggers: How will key dates be sourced—user input, static rules, or external system integration?	Key information needed for future notifications will be sourced primarily through user input.
217	Beyond AES-256 encryption and GDPR/CCPA, are there additional standards (penetration testing, OWASP ASVS, FedRAMP)?	Yes, beyond AES-256 encryption and GDPR/CCPA compliance, the solution is expected to align with OWASP ASVS best practices and undergo regular third-party penetration testing.
218	Is there a specific reason for cancelling the previous RFP that was issued under a different name but covered the same scope? Note there were other questions asked about the cancelled RFP.	The Court is not compelled to specify the reason for cancelling an RFP and will not answer questions pertaining to a cancelled RFP.
END OF FAQs		