☑ **AMENDMENT OF SOLICITATION** ☐ **MODIFICATION OF CONTRACT** | Page 1

1a. Amendment No.	1b. Effective Date	3. Issued By
0001	7-16-25	U.S. House of Representatives CAO Office of Acquisitions Management 5110 O'Neill House Office Building Washington, DC 20515 Office Phone: 202-225-2921
2a. Solicitation Number	2b. Dated	
RFI OAM25072S	6-30-25	

4. Name and Address of Offeror	For Information, Contact:	
	4b. Name:	4c. Phone:
	4d. Email:	

5. The solicitation in Block 2a above is amended as set forth in Block 6 below.

The time and date specified for receipt of Offers ☑ is not extended
☐ is extended through

Offers must acknowledge receipt of this Amendment in Block 10 of the Solicitation, Offer and Award Form to be submitted prior to the time and date as amended or specified in the solicitation.

6. Description of Amendment

THE PURPOSE OF THIS AMENDMENT IS TO PROVIDE ANSWERS TO THE QUESTIONS RECEIVED FOR THIS RFI.

PLEASE SEE THE BELOW PAGES.

7a. Name and Title of Authorized Signer *(type or print)*		8a. Name and Title of Contracting Officer *(type or print)*	
7b. Offeror *(Authorized Signature)*	7c. Date Signed	8b. U.S. House of Representatives Kevin Morris Digitally signed by Kevin Morris Date: 2025.07.16 13:29:34 -04'00' *(Signature of Contracting Officer)*	8c. Date Signed

RFI-OAM25072S – GraphQL QUESTIONS with RESPONSES

1. **Hasura and GraphQL Endpoint Management**
   Are you considering alternatives to Hasura for GraphQL endpoint management, or is continued use of Hasura preferred?

A There is no preference. The RFI merely states our current approach and technology stack for context and comment.

2. **GraphQL Federation Strategy**
   To what extent is GraphQL federation under consideration, and do you anticipate federating services across agencies (e.g., GPO, Library of Congress)?

A Federation is likely to be a critical component of our future architecture. It is more likely to be internal to House systems, but multiple systems contribute data and behavior to the House data ecosystem.

3. **Hybrid Architecture Support**
   Would the House be open to hybrid architectures (e.g., combining REST microservices and GraphQL for specific workflows)?

A For this RFI we are focused on GraphQL as an API layer. There may be practical or implementation concerns that require the use of REST APIs for some components, as occurs today in some systems.

4. **Configuration Automation**
   What level of configuration automation maturity is expected? Does the current Helm chart deployment model meet your goals, or are you seeking enhancements such as GitOps?

A The current implementation meets our goals, but we expect it to evolve. Our current application charts and deployments do not use a GitOps tool, but they are organized to eventually work with one.

5. **Schema Version Control Expectations**
   Is there a need for schema version control tooling? Should it integrate with Git, Helm, or other systems?

A Yes, schema versions and updated APIs will need to be controlled, and their configurations will need to be managed. Our current tooling allows us to track versions through the entire stack, from the Helm appVersion to a Docker tag to a Git repository tag. We would anticipate having similar traceability in any future system, including API or API service versions.

6. **GraphQL Tool Preferences for Developer Experience**
   Are there specific tools or frameworks currently under evaluation for schema evolution, testing, and validation?

A No. We currently use common GraphiQL-based tools for simple ad-hoc querying, and we have used tools like Python Notebooks for more complex data orchestration and testing. There is no preference for those tools.

7. **Governance Models for Schema Evolution**
   Would you be interested in recommendations for governance models that support schema evolution and prevent breaking changes?

A Yes

8. **Front-End Integration Tools**
   Do you have a preference between Apollo Federation, Relay, or other client-side GraphQL tools?

A No. We currently use the Apollo client and graphql codegen for Typescript integration.

9. **Open-Source Tool Usage**
   Are there any restrictions on the use of open-source GraphQL tools in production environments deployed behind the firewall?

A No

10. **Security Standards (FedRAMP, OAuth2)**
    Will proposed tools or services need to be FedRAMP-certified, or simply comply with internal House security controls?

A We comply with security policies and controls issued by the House CISO (Chief Information Security Officer).

11. External Data Integration Approaches

How are external data sources (e.g., Library of Congress, Senate systems) currently integrated, and should future tooling support dynamic ingestion or schema mapping?

A Current integration uses a range of approaches, from legacy file formats to APIs. Dynamically ingested data or schema integration is not a significant concern. Data and schemas across the Legislative Branch are very similar and there is a high degree of collaboration between organizations around schema management.

12. Schema Translation for Partners

Should proposed solutions support schema transformation or translation between internal and partner-facing GraphQL APIs?

A Yes, but this is not a significant concern. The differences between the internal and partner facing schemas are small and there are already well-defined schemas to base these on, such as the Congress.gov API.

13. Follow-On Solicitation or Pilot Phase

Is there a planned follow-on RFP, RFQ, or potential pilot project to follow this RFI?

A Nothing at this time has been identified. This RFI is for Market Research purposes only.

14. Project Timeline for External API Delivery

What is the anticipated timeline for delivering GraphQL services to external legislative partners?

A A timeline has not been established.

15. Regarding developer productivity, what quantifiable improvements are you targeting for API delivery speed from conception to deployment if you were to implement a new tool like Grafbase?

A Productivity metrics have not been established.

16. Can you provide metrics on the current effort or estimated time spent by your teams on managing schema changes and ensuring consistency with your existing Hasura GraphQL implementation, and how much would you like to reduce this?

A We do not currently have metrics for this activity. Since this is largely driven by our database modeling activities the overall effort is low.

17. For Configuration Management, what percentage of your entire system configuration is currently managed as code, and what are the primary challenges in achieving full automation for deployments using Helm charts with your current tools?

A Nearly all our systems are managed as code, from infrastructure to builds, and all key parts of our deployments are fully automated. The main challenge is that it is a lot of code.

18. Regarding operational simplicity, what specific metrics are you tracking for observability (e.g., MTTR, resource utilization) for your containerized GraphQL services, and what are your benchmarks for improvement with a new solution, especially as you migrate metrics collection to Prometheus?

A We track the usual resource utilization metrics (e.g. CPU and memory utilization). As our API surface expands and the internal service dependencies become more complex, we are especially interested in request traceability. We do not have specific benchmarks for metrics collection.

19. Could you elaborate on the most critical security considerations for an externally facing GraphQL API, particularly concerning OAuth2 and JWT bearer token integration, and what capabilities would a solution like Grafbase need to demonstrate for your self-hosted private network environment?

A Externally facing APIs would likely have fewer access controls and more protection from malicious usage, such as denial-of-service attacks. Future access control for internal or partner usage is likely to follow standard approaches and current patterns where OAuth and JWT tokens are used to authenticate requests and assert claims for access control purposes.

20. When considering flexibility, beyond developer control, are there specific types of data transformations or integrations you anticipate Grafbase needing to handle within your GraphQL layer that are challenging with your current setup?

A Data transformation is not a significant concern. Most application use relational databases for their core data management.

21. For data reuse, what are the most common data sources you anticipate integrating with a solution like Grafbase for both internal applications and external partners, and what volume of data do these integrations typically involve?

A **Almost all data sources are internal applications. Typically, these applications are built on top of a relational database.**

22. Beyond the technical benefits, what are the top 2-3 business outcomes or strategic advantages the U.S. House of Representatives aims to achieve by optimizing its GraphQL delivery strategy with a tool like Grafbase?

A **We want to provide flexibility to data consumers and integrators as they build applications with data we deliver. We also want to make it easier to integrate data sources from across the organization. Ultimately, we want to make access to data and data integration easier, more flexible and more self-service versus having custom data integrations and stove-piped data from multiple systems and organizations.**

23. What are the estimated operational costs (e.g., infrastructure, maintenance, licensing) of your current GraphQL stack, and what level of cost efficiency are you hoping to achieve with a future self-hosted solution like Grafbase?

A **We do not have a current estimate for our operational costs specific to GraphQL, but we anticipate it will grow as more data is delivered to more consumers.**

24. Are there any specific compliance or regulatory requirements that heavily influence your technology choices for data handling and API exposure that a solution deployed in an air-gapped environment like Grafbase would need to address?

A **No**

25. Are you currently working with a preferred partner and who is your main point of contact?

A **No**

26. Do you require professional services to assist with the initial deployment and are there any additional training you will need particularly around configuration?

a. **We have no specific plans for additional services at this time.**

27. Are there any formatting or document organization requirements for the RFI response beyond the 10-page limit?

A **No.**

28. Does the Office of the Clerk have specific pain points that it is seeking to alleviate by modernizing the current implementation of GraphQL?

A **No. We are looking for technology and approaches that make it easier to consume our data**.

29. In the Background section of the RFI, three technologies are specifically mentioned: GraphQL federation, streaming, and Relay. Does the Office of the Clerk have defined use cases in mind for these technologies, or are they cited as general examples? Specifically,

A **For federation – are there other systems for which the Office of the Clerk would like to create GraphQL endpoints to federate with the current system, or is the goal to break the current system into a collection of federated services?**

A We are interested in federation as a general approach to integrating data from multiple systems and APIs.

a. **For streaming – what type of data is expected to be streamed (e.g., audio/video or other large binary data)?**
   A: Documents are a common streamed data source, often in PDF, but also other binary formats. We do have some streaming video sources.

b. **For Relay – are there particular objectives or challenges the Office of the Clerk is aiming to address through the use of Relay?**

A: **No, it is just provided as an example of a newer GraphQL related technology.**

29. Figure 2 of the RFI shows multiple APIs. Does the Office of the Clerk intend for a distinct GraphQL API to be provided for each consumer shown (i.e., public, partners, clerk)?

A **This illustration is conceptual. There is no specific commitment to how the API will be structured, but there is a separation of concerns between these different consumers.**

30. Does the Office of the Clerk currently use role-based access control within Hasura for secure, authenticated access, or is another mechanism in place?

A **All access control is managed with JWT claims.**

31. What tools or methods does the Office of the Clerk currently use for automated testing of the existing GraphQL services?

A **Python Notebooks have limited use as an API automation tool for specific use cases.**

32. Given that GraphQL implementations are typically additive, does the Office of the Clerk anticipate any schema changes that would involve the removal of queries or attributes?

A **It is possible, but this should happen rarely.**

33. Are there specific performance or latency benchmarks that the modernized GraphQL APIs must meet?

A **No**

34. Can the Office of the Clerk provide an estimate of the anticipate number of transactions (queries and mutations) for the modernized GraphQL API?

A **No. In general, our data and transaction volumes are small, especially for internal or Legislative Branch-facing systems.**

35. What is the long-term vision for enabling non-technical users (e.g., legislative analysts) to query and analyze data from LIMS or partner APIs using natural language?

A **This is certainly a potential use case for APIs in the future. Data integration with partners is a higher priority currently. However, we anticipate robust natural language search for specific data (e.g. bill or Member information) to be a critical component of any API because it is a key requirement for applications.**

36. Are there any current or planned initiatives around building a semantic layer or knowledge graph across Clerk systems?

A Not at this time. In most cases this work would be done by other organizations and coordinated through the Congressional Data Task Force.

37. How are you addressing the challenge of onboarding new analysts or developers who need to learn the internal schema and business logic quickly?

A We have internal documentation and training for new staff, in addition to training on Congressional process and data through the Congressional Staff Academy.

39 Are you exploring AI-assisted automation for repetitive tasks within LIMS or across data integration workflows with partners (e.g., GPO, LOC)?

A Not at this time, but we are open to input about how AI could be used to improve the areas of interest outlined in the RFI.

40. Is there a roadmap for integrating AI to assist with legislative research, bill tracking, or committee report generation?

A These items are outside of the scope of this RFI.

41. How does the Clerk's office currently manage schema evolution and versioning across Hasura GraphQL endpoints?

A Most changes are backwards compatible, but updates are coordinated across development teams and tested before release.

42. **Are developers manually wiring up GraphQL to UI components, or is there an abstraction layer for reuse?**

A Queries are created manually. The graphql codegen tool is used to facilitate integration with our React codebase.

43. How are you handling observability and debugging when APIs fail or deliver unexpected results in production?

A Through log analysis and database analysis. Data is often replicated in developer environments to reproduce problems.

44. Is there interest in reducing the complexity of CI/CD testing for GraphQL API contracts and data transformations?

A This is a future concern. There is interest in learning about strategies for API configuration management, governance and testing more generally.

45. How are authorization policies enforced across different GraphQL endpoints or datasets shared with partners?

A Authorization and access control would be enforced by each service based on the authentication token (currently a JWT) presented to the endpoint in the request.

46. Are you exploring options for self-hosted AI solutions that operate within a private network and comply with federal data handling guidelines?

A Not within the scope of this effort. There are other efforts to examine the role of AI in House systems.

47. Do any teams currently run retrieval-augmented generation (RAG) or LLM-based pilots internally or with partner orgs?

A Not within the scope of this effort.

48. Have you evaluated vertical AI agents for legislative or administrative workflows, and if so, what gaps remain?

A Not within the scope of this effort.

49. Would you be open to a no-cost PoV where PromptQL delivers an interactive, NL-powered demo using mock (or redacted) LIMS schema to showcase use cases like: Given your data is already in Hasura (Now PromptQL) this seems like a natural progression where you can either make decisions faster OR facilitate automation.

A Not as part of a response to this RFI.

- **Example: "Summarize all amendments filed to HR 1234 last month."**
- **Example: "Track average time from bill introduction to committee action."**

50. What is the maximum number of concurrent GraphQL queries your platform has reliably supported in your production environment, and what sustained QPS (queries per second) benchmarks? Can you share the ratio of simple to nested queries of your current system?

A We do not have specific metrics to address this. Current GraphQL deployments are for internal systems and the highest levels of activity are relatively small. Scaling is not a current concern.

51. Can you provide your typical median, 95th-percentile, and 99th-percentile response times for read and write operations, and describe how these metrics change under increasing payload sizes?

A We do not currently collect these metrics routinely.

52. How does your solution detect and respond to sudden traffic spikes—what metrics drive autoscaling, what are your scale-up/scale-down latencies, and how is state managed across nodes during scaling events?

A Traffic spikes are minimal and have not caused performance degradation. Scaling is not a current concern.

53. What mechanisms (rate limits, quotas, circuit breakers, query-cost analysis) do you provide to prevent individual clients or queries to prevent them from degrading overall service performance?

A None. For current internal systems traffic is manageable and scaling is not a concern.

54. What is the current cost of your Cloud and On-Prem infrastructure to handle this project, including licensing units (per-node, per-core, per-request), and any volume or commitment discounts you're receiving as a federal entity?

A We do not currently collect these metrics specific to this project. Per unit costs are not a current concern.

55. What uptime and support-response SLAs are you looking for (e.g., for P1 incidents)?

A We do not have specific requirements for an SLA and would address this for specific tools, services or vendors.

56. What issues have you run into post-implementation of GraphQL?

A Very few. Aside from normal API integration concerns we have not experienced any performance, security or integration issues that are a current concern.

57. Outline the ideal timeline you're looking at for any proposal on each of the topics.

A Nothing at this time has been identified. This RFI is for Market Research purposes only.

58. Describe your current protections against common GraphQL threats (query-depth limiting, introspection control), support for OAuth2/JWT, and encryptions used.

A Authorization, in the form of a JWT token, is passed to the lowest level of the services to validate credentials and enforce access control. For current internal systems we do not enforce other limits on the GraphQL API. Systems are hosted on private networks and use normal encryption approaches based on HTTPS/TLS for all accessible endpoints.