



**DC COMPASS
AI VALUES ALIGNMENT REPORT**

In February 2024, through Mayor’s Order 2024-028, Mayor Muriel Bowser created a robust process for documenting and publicly discussing the work DC agencies do to make sure the AI tools they deploy are well aligned with DC’s values. As agencies work through this process, they create AI Values Alignment Reports, which respond to each of DC’s AI Values articulated in the Mayor’s Order, and described in-depth in the publicly available DC AI Values Alignment Handbook, <https://techplan.dc.gov/page/hand-book>.

This AI Values Alignment Report explains OCTO’s deployment of the DC Compass AI tool with the help of Esri, a global leader in Geographic Information System (“GIS”) software. The DC Compass tool brings together DC’s over 1,400 Open Datasets, Esri’s powerful GIS tools, and the diverse AI tools available through Microsoft’s Azure OpenAI platform, to give every resident rich, immediate information about the spaces where they live, learn, and thrive.

DC has made these Open Datasets available to the public for years, as part of the DC Data Policy created by Mayor Muriel Bowser through Mayor’s Order 2017-115. However, data literacy gaps, significant time and resource demands, and other frictional barriers have limited some residents’ ease of access to these important datasets that describe their own streets and neighborhoods. The DC Compass tool is designed to remove these barriers, and to help shrink the digital divide by providing every resident with far faster, cheaper, and more meaningful access to DC’s geographic data than ever before.

Through DC Compass, residents can ask questions verbally or in writing, in the language with which they are most comfortable, and get faster and better cited answers than a whole team of data scientists could have provided in the past. DC Compass automatically creates maps and dashboards responsive to residents’ questions, and provides statistics and natural language summaries in the same language chosen by the resident. It can suggest related datasets the resident might be interested in, and can even identify specific DC government initiatives related to the resident’s query.

When deploying such powerful AI tools, it is important that agencies carefully consider each of DC’s AI Values. This Report proceeds step-by-step through these AI Values, explaining the care with which OCTO reviewed each, the efforts that went into anticipating potential risks of unintended harms which might arise from deployment of the tool, and all of the work involved in structuring the planned deployment to mitigate those risks.

For more information on AI, and how its deployment fits into DC government’s Tech Plan, please visit <https://techplan.dc.gov/page/dcs-ai-values-and-strategic-plan>.

DC AI Value

Clear Benefit to Residents

Everything DC government does, it does for the benefit of residents. The deployment of AI tools is no different. This section of the DC Compass AI Values Alignment Report explains how OCTO intends for the DC Compass deployment to benefit DC residents. It is organized around the Key Concepts identified in DC's AI Values Alignment Handbook.

I. The Purpose of the DC Compass Deployment

Through its Open Data program, DC makes more than 1,400 datasets available to the public in machine readable format, free, for any user. These datasets are accompanied by valuable data dictionaries, and rich metadata. Users have traditionally made use of this data through keyword searches and other manual, time-consuming techniques. Sophisticated users, with access to sufficient resources and training, have been able to use the data to construct maps, dashboards, and similar data. While the Open Data program has succeeded in opening up a large volume of valuable data to the public, OCTO hopes the DC Compass tool will help more equitably to expand that access to every resident.

Toward that end, DC Compass is designed to reduce three barriers:

- Overwhelming Data Volume
- Data Literacy Gaps
- Non-Obvious Relationships between Datasets and Initiatives

A. Addressing Overwhelming Data Volume

DC adds to its 1,400 Open Datasets every year. The volume of information is difficult to wade through, even for users accessing the data as part of their day job. The annual review of Open Datasets by agency legal staff, for example, represents a serious investment of personnel hours. Handling so large a volume of data is even more daunting for the average resident.

For example, if a resident searches for the keyword “procurement,” that search may return 12 different available datasets. To find out which, if any, of these 12 datasets contains the information the resident seeks, they will need to individually review each dataset. If the resident is comfortable working with metadata, they may be able to identify the correct dataset more quickly than manually reviewing the data it contains. If not, they are left to search by hand through unfamiliar, technical data. The more general the search, the higher number of irrelevant datasets will be returned, and the more time and effort a resident will have to spend searching through data.

B. Mitigating Data Literacy Gaps

The term “data literacy” describes a user’s ability to work with data and draw meaningful information from it. Users with a lower level of data literacy may struggle to draw information from structured datasets, and may be unable to put what they learn from the data into a meaningful context. While DC already provides web tutorials and other free resources to help close data literacy gaps, even freely available materials may require a significant investment of time that residents may not have available to spend exploring Open Datasets. Data literacy gaps are difficult to address, and they contribute to the persistence and negative effects of the larger digital divide.

C. Identifying Relationships between Datasets and Initiatives

Important relationships exist among datasets, and between datasets and DC government initiatives. These relationships can be difficult or impossible to recognize from a single structured dataset. For example, if a resident is interested in what their government is doing to help make DC’s roads safer, and they search for “crashes” using existing Open Data tools, those tools are not sophisticated enough to associate the resident’s interest in “crashes” with DC’s important Vision Zero road safety initiative.

DC Compass is designed to address each of these three barriers. Through greatly improved natural language information retrieval abilities, powerful GIS visualizations, and generative AI natural language summaries, residents can be brought closer to the information they seek, faster and easier. Users will be able to ask a question, get a plain language response, an interactive map, a dashboard, a chart, data formatted for a spreadsheet, and citations to the Open Datasets from which all this information was drawn.

II. Making Sure the Benefits Flow to Residents

Throughout the Open Data program’s life, the benefits of DC’s Open Datasets have stretched far beyond its geographic borders. While this breadth of benefit is a good thing, the DC Compass deployment is calculated to deliver even greater benefits more directly and more equitably here at home. It does this by empowering users of all levels of data literacy to cut through excess and irrelevant data, and by helping to surface the data and programming that users are seeking, quickly, easily, and for free.

Policy professionals, academics, journalists, and businesses have long had the resources necessary to make the most of Open Data, and through their work, a resident would indirectly learn more about their community, or have access to products or services better tailored to their needs here in DC. But however valuable these indirect benefits may be, residents should have more equitable *direct* access to their community’s data. DC’s Open Datasets are drawn from the community, developed by residents, by DC government employees, by DC government’s partners, and are paid for with public money. In a very real way, that data belongs to residents. Using powerful AI tools, DC Compass minimizes technical challenges around volume, data literacy, and integration, and puts over 1,400 Open Datasets directly within reach of those residents to whom the datasets belong.

More than would have ever been possible with legacy technology, the benefit of DC's Open Datasets will flow directly to every DC resident. Residents will be able to write or speak their questions for the DC Compass tool in whatever language is most comfortable to them, and they will receive an AI-generated summary response in the same language, plus clearly sourced hard data, and additional content like interactive maps, charts, and other data visualizations.

DC Compass will help use residents' natural language inputs to search the Open Data library and identify the content they are after. The DC Compass tool will search through metadata that is difficult for everyday users to parse, and it will show its work by providing links. Finally, after doing the hard work of transforming natural language questions into complex and technical queries, and returning the material sought, the DC Compass tool will take the next step and link to relevant DC government insights, policies, and initiatives associated with the responsive datasets.

What are some concrete use cases? Residents can find DC government services or facilities in their neighborhood. They can find the closest DMV location, or a specific branch of the library, and have all the Open Data associated with those locations immediately at hand. When a student wants to anchor their science fair entry or a social studies report in their own community, DC Compass will make researching DC as quick and easy as any search engine, with bespoke charts and visualizations cited to real datasets. If a resident takes an interest in DC's urban forestry, bicycle safety, or affordable internet, and needs hard data to support an idea they want to bring before DC government leadership, the DC Compass tool will help them lay a sound, data-driven foundation for their own civic engagement efforts.

And, though DC Compass is designed to improve equitable, direct access for residents, the AI deployment will also provide more benefits to sophisticated, well-funded users than before. Outside groups have long benefited from DC's practice of affirmatively making public huge volumes of high-quality data. These groups include national and international academics, journalists, public policy professionals, think tanks, businesses, and interested private persons the world over. Because DC makes Open Data available subject to standard open licensing terms, when these outside groups use Open Data to create new, "derivative datasets" from DC's Open Data, these groups are required both to share those datasets back with the public, and to credit DC's Open Datasets as their source.

What other sorts of benefits will DC Compass bring to organizations? Non-profits and businesses can find grant opportunities in every geographic region of DC. They can explore wards and neighborhoods to identify the best places to target new community services or to start a new business. Civic groups and university researchers can more quickly identify and process available data sources and examine statistics tied to any geographic region of DC they are studying. The DC Compass tool will deliver more data faster and easier to journalists and academics than far slower, far more expensive tools like Freedom of Information Act requests.

Finally, DC Compass will actually help to improve the Open Data program. By helping more people bring more scrutiny to the datasets of greatest interest among DC's 1,400 Open Datasets, DC Compass will help drive increased community interaction with the data. This increased

interaction will, in turn, mean more quality assurance, better error recognition, and in the years following deployment will improve the accuracy of DC’s Open Datasets simply by empowering a much larger, much more diverse group to handle the data they care about. Though OCTO’s Data Team works to improve data hygiene across the Open Datasets, the crowd will help identify inconsistencies and missing information at scale, providing a level of quality assurance that would be prohibitively expensive in a world where only paid employees interact meaningfully with the Open Datasets.

III. Considering Alternatives Tools

When a DC agency considers deploying an AI tool, it is important to pause, take a deep breath, and make sure we are not forcing an overengineered solution where simpler technology would do the same job just as well. Before OCTO settled on the DC Compass tool, we thought critically about just how sophisticated a tool was necessary for addressing the data volume, data literacy, and data integration challenges we set out to address. We decided the DC Compass tool was the best approach.

Three core technology functionalities were necessary to make DC Compass a success: natural language processing, retrieval augmented generation (sometimes called “RAG”), and Geographic Information System (“GIS”) tools. AI technology is necessary to provide each of these three functionalities at the level needed to bridge the data literacy divide. And the DC Compass tool specifically, was the only AI-powered tool we found capable of providing all three.

Broadly, the phrase “natural language processing” describes the task of making meaningful use of everyday language (as opposed to the structured language of computer code commands or Boolean search parameters). The DC Compass tool uses Microsoft’s Azure OpenAI platform to process natural language at two different steps in its workflow. First, it uses language modeling tools to allow users to enter their questions—as text or voice, in the language of the user’s choice—into the system as natural language. Later, it uses the same language modeling tools to generate a synthetic, natural language summary of the results the program provided. None of these tasks—voice to text translation, language to language translation, natural language recognition, natural language generation—is possible without powerful AI technology running on enormous neural networks. There are no non-AI alternatives in existence capable of even passable accuracy on any of these tasks, and successful performance of each task is necessary to achieve DC Compass’s equity objectives.

“Retrieval Augmented Generation” describes a generative AI technique that helps to limit the risk of wild hallucinations in synthetically generated content. Rather than asking a language model to try to deliver an answer to a question based only on its (very, very) large-scale statistical modeling of the words in its enormous training dataset, an AI tool using RAG only asks the language model to transform the user’s question from a natural language question into something more like a formal query for a structured data set. Then, instead of the language model functionally asking itself and its training dataset the question, it asks the question of a pre-defined outside dataset.

In the DC Compass deployment, that pre-defined outside dataset is the DC Open Dataset. What this means in practice is that the Azure OpenAI platform used by the DC Compass tool does not go looking for the answer to a question in its language model's own training data or on the publicly available internet. Instead, it goes looking for the answer in the carefully curated DC Open Datasets. This means, in turn, massively reduced risk of wholly made-up hallucinated answers, and clear citations to the sources of data. RAG, like general natural language processing, is impossible without powerful AI technology running on powerful neural networks.

Finally, GIS describes the two- and three-dimensional data modeling and visualizing tools that are necessary both to generate maps from row/column datasets, and to generate row/column datasets from mapped data. Existing LLM tools cannot perform GIS analyses or functions, and cannot generate summaries of GIS data natively, or independently. And while hand-programmed, non-generative GIS tools have existed for decades, there is no way to integrate GIS with natural language interfaces at the level needed to meet DC Compass's equity goals without custom built interfaces between GIS software tools and AI tools built on powerful neural networks.

OCTO tested different AI models with existing GIS technology to identify all options available for integrating these three functionalities. We determined that a limited subset of DC Compass's goals may have been achievable through a language model alone, but that without GIS integrations the data visualizations we hoped to provide would not be possible, and as a result the data literacy gaps we hoped to help close would largely be left unaddressed. We needed direct integration between GIS tools and large language models, and could find no alternative solution.

DC AI Value Safety and Equity

AI tools provide powerful means for improving the quality and efficiency of government services. When deploying powerful tools, though, safety and equity considerations are of heightened importance. This section of the DC Compass Values Alignment Report helps to explain all that OCTO has done to take safety and equity into account in its planning the DC Compass deployment. It has been divided into four key areas of considerations that correspond to the Key Concepts identified in the Safety and Equity section of the DC AI Values Alignment Handbook.

I. Addressing Risks of Direct Physical Harm

When deploying a new technology, an agency must take special care to minimize any risk that the tool might cause physical harm. Because the DC Compass tool only operates by receiving queries from the public, and returning responsive materials from an Open Dataset, there are no obvious risks that the tool might cause direct physical harm. To help protect against non-obvious risks, OCTO worked with Esri to build automated safeguards against misuse of the DC Compass tool that might lead to direct physical harm.

For example, DC Compass is designed to deliver predetermined or “canned” answers to user queries that suggest the user might need immediate help from emergency services. Instead of processing queries like, “who do I call for an emergency,” by searching Open Datasets and generating maps or synthesizing summary text, the DC Compass tool will respond with the canned answer, “for an emergency, please call 911.” This is intended to help make sure that a user who misunderstands the functionality of DC Compass does not find their access to emergency services delayed because they mistook the DC Compass tool for a 911 substitute. These canned answer lists can be updated over time by OCTO, in consultation with Public Safety and Justice Cluster agency partners, if and when misuse trends develop.

II. Minimizing Risks of Indirect Physical Harm

Because the DC Compass tool exists only to receive queries and return responsive material from open data sets, there are no obvious risks that the tool might cause indirect physical harm. As OCTO searched for non-obvious risks of indirect physical harm, it occurred to us that increasing general access to geographic information might affect the level of risk for residents who are at increased risk of being the target of violence, whether they are a highly visible political or community leader, or they are the target of specific threats from another individual. This risk is effectively mitigated, however, by DC Compass’s design.

DC Compass does not draw data directly from agency data sets. It can only draw data from a specifically pseudonymized and anonymized copy of those datasets, which is reviewed regularly by each contributing agency’s General Counsel. This means that existing safeguards and processes protecting those with need of special privacy considerations with respect to data that

otherwise would be included in the Open Dataset still enjoy the same protections from any searches or productions made by the DC Compass tool.

III. Safeguarding Fundamental Rights

DC Compass is a powerful new tool in promoting open government and increasing resident awareness and engagement with the government that serves them. Given its multi-lingual, natural language interface, the tool provides public access to DC's Open Datasets more equitably than legacy tools ever could. The tool does not play any role in any determination of access to any government service, or in the adjudication of any rights. Despite careful consideration, OCTO has been unable to imagine a scenario where this tool might act to undermine fundamental rights. If such a circumstance were to arise, OCTO retains the ability shut off access to the tool and return to the status quo while whatever problems that have arisen are dealt with. During that time, or at any point users are uncomfortable with the tool, users can continue to use traditional tools to locate and work with Open Datasets as they did before deployment of the DC Compass tool.

IV. Promoting Equity

DC Compass is an important, multilingual, natural language equity tool. If it proves as successful in live deployment as OCTO hopes, it may serve as a more proactive model for open government initiatives in other jurisdictions. It is difficult to imagine the tool having an opposite effect, negatively impacting equity in DC.

DC Compass searches, summarizes, and provides information based on datasets that have been affirmatively released to the public since 2017. If the tool works as planned, and increases public engagement with the Open Datasets, it may help reveal existing, non-obvious inequities. Revelations of this kind can be the first step in identifying and addressing the root causes of inequity. DC Compass may prove to be a valuable tool in tracking access to important infrastructure, community anchor institutions, and more. If DC Compass proves popular with residents, it may also help to improve digital literacy and data literacy at a time when those skills are expected to grow in value, year over year.

DC AI Value Accountability

As valuable as DC Compass’s AI functionalities are in making access to DC’s Open Data faster, cheaper, and more equitable, OCTO has been careful to make sure that increased automation does not lead to a decrease in accountability. The DC Compass tool is highly sophisticated, it is very powerful, but it is not “intelligent” in the general sense of the word. Neither it, nor any other AI tool, “understands” its data like a human would.

As a result, OCTO structured its deployment so that responsibility for the impacts the tool makes remains at all times clearly traceable to DC government officials. This section of the DC Compass AI Values Alignment Report focuses on accountability, and is organized around the Key Concepts identified DC’s AI Values Alignment Handbook.

I. Tracing Responsibility to Real People

Before DC Compass, the DC Open Data program relied on individual program staff and custodians of records to generate and process its various enterprise datasets. The program relied on the legal and privacy staff of DC government agencies to review their enterprise datasets for privacy, safety, and regulatory compliance before release through Open Data. The program relied on DC’s Chief Data Officer (“CDO”) to oversee the aggregation and publication of the Open Datasets. DC Compass does not change anything about this process. These responsible government officials remain directly accountable for the maintenance and release of DC’s Open Datasets exactly as they were before launch of the DC Compass tool.

To learn more about how DC manages this accountability, please review DC’s Open Data Handbook, available at <https://opendata.dc.gov/pages/handbook>. Immediate responsibility for the DC Compass program is traceable to the DC Compass program manager, who can be reached at Open.Data@dc.gov.

As already addressed in this Report, the increased community use of DC’s Open Data should lead to greater accountability for the accuracy and completeness of DC’s Open Datasets. Crowdsourcing reviews and scrutiny of records can greatly improve oversight and error detection at scale, at a net benefit to every member of the community.

II. Making Sure the DC Compass Tool Performs

OCTO is actively monitoring the performance of the DC Compass tool. The tool’s formal performance metrics include: success and failure response rates, accuracy of responses, speed of responses, and cost. Success and failure response rates, accuracy of response, and speed of response should improve as the language modeling tools available through Azure OpenAI improve, independent of the DC Compass program. Similar improvements should be expected if subsequent versions of the DC Compass program upgrade to higher performing language models available through the Azure OpenAI platform. These three performance metrics should improve,

too, as the integrations between the Azure OpenAI tools and the DC Compass GIS tools improve. Finally, accuracy of responses should improve as DC's underlying Open Datasets improve from the additional review and error correction that comes with higher volume engagement and more equitable access. OCTO anticipates cost increases over time, both as the cost of the DC Compass tool increases to enable improvements to the product itself, and as the pass-through cost of querying Azure OpenAI tools increases.

OCTO will curate a set of standardized queries to anticipate commonly asked questions, and focus attention on testing and improving the quality of DC Compass's outputs in response to these common queries. This is a common technique in optimizing AI systems both for improved accuracy and reduced cost.

OCTO's partner Esri provides feedback on DC Compass software updates, improvements, and bug fixes, maintaining a direct line of communication between OCTO and Esri developers. Esri provides OCTO with the data access necessary for OCTO to monitor the tool's performance on a query-and-answer pair level. This allows OCTO to examine why a particular dataset may be referenced in response to a question, enabling us to review metadata of datasets and correct as necessary for better reliability and performance. This query-and-answer pair visibility also helps OCTO to understand the scope of queries and questions coming into the tool, to fine tune performance to user preferences, to adjust the scope of standardized queries, and to make sure the 911/emergency services canned answers discussed earlier in this Report are functioning as intended.

OCTO will track and mitigate any values alignment drift primarily through public engagement facilitated by OCTO's State Broadband and Digital Equity Office ("SBDEO"). By monitoring usage volume, and by directly engaging with the community, OCTO plans to stay well informed on the status of the community's uptake of the DC Compass tool, and the equity outcomes of this powerful new asset. If the tool begins to show signs of values alignment drift, OCTO is capable of directly deactivating the system until any values misalignment can be identified and remedied. Learn more about the SBDEO's digital equity and community engagement work at <https://www.techtogetherdc.com/>.

DC AI Value Transparency

Transparency is essential to accomplishing the equity goals of the DC Compass deployment. OCTO hopes the tool will encourage residents to use DC's Open Datasets more, to encourage them to use these datasets to become better informed about how their government is working for them, and hopes the tool will increase residents' engagement with all DC agencies on a continuous basis.

This section of the DC Compass AI Values Alignment Report documents OCTO's public engagement efforts around the tool, and shows how the DC Compass deployment was planned to ensure transparency in the public's interaction with this new technology. This section is organized around the Key Concepts identified in DC's AI Values Alignment Handbook.

I. Engaging with the Public on DC Compass

OCTO began live testing of the DC Compass tool with an alpha testing phase involving internal stakeholders. This included data analysts, data scientists, program managers, and agency executives. The DC Compass deployment is currently in its public beta testing phase, which includes an extensive public outreach and community engagement effort.

OCTO started an "Open Data DC Campaign" in 2023, designed to promote the Open Data program, and to engage residents, businesses, students, and other groups. We plan to hold demonstrations and learning sessions as part of our campaign to further socialize the tool, and to receive feedback on its deployment. In previous engagements, we have partnered with local community groups interested in our data outreach efforts, helping with messaging and organizing events. Esri, agency Chief Data Officers, Geographic Information Officers, and Chief Information Officers from other jurisdictions have all been helpful partners along the way. The breadth of our partnerships and collaborators has helped swell attendance at our public Open Data events, and helped us incorporate diverse perspectives into every stage of planning and risk management involved in the deployment of DC Compass.

To attract direct user input, we have incorporated feedback mechanisms into the DC Compass tool. For every response DC Compass provides, users have the opportunity to supply direct feedback through an online survey. The survey allows for both positive and negative feedback using a category-based approach, with optional text entry for more detailed information about their experience. This allows users to note specific issues they may have encountered, and to capture their overall level of satisfaction. We intend to use this user feedback data, together with the performance metrics discussed earlier in this Report, to continuously fine-tune the DC Compass tool. The survey is anonymous, but includes encouragement for users to contact OCTO directly to share their responses and to follow-up on improvements.

Despite these robust efforts, it is important to document those circumstances that will limit the impact of our public engagement efforts. While OCTO's State Broadband and Digital Equity

Office continues to work to address the digital divide, there are still households in DC without reliable, affordable broadband internet access. A high-speed internet connection is not absolutely necessary to make use of the DC Compass tool, but residents with limited access to the internet or limited familiarity with digital tools are unlikely to make Open Datasets and AI tools a big part of their day. The SBDEO continues to work on addressing the digital divide in DC, and as those efforts continue to bear fruit, we anticipate similar improvements in our own public engagement with the DC Compass tool.

Additional factors will limit public engagement while the DC Compass tool remains in beta testing, which limits access to users with a registered ArcGIS.com account. However, these limitations are valuable from a privacy and cybersecurity perspective, as discussed later in this Report. OCTO bears these limitations in mind as it evaluates the DC Compass tool, to help contextualize the feedback flowing from technology inclined individuals who would create an Esri ArcGIS.com account. After we complete the public beta testing phase, we expect this limitation to be easy to address.

II. Labeling Synthetic Data

When using generative AI technologies, there is always a risk that users will mistake content created by an AI tool—text, images, video, audio—as content created by a human. DC Compass’s use of Retrieval Augmented Generation is a powerful safeguard against the tool generating inauthentic data. Every dataset surfaced by the DC Compass tool is an authentic, human-reviewed dataset from DC’s Open Data program.

In order to communicate with users in natural language, and in the language of their choice, however, the DC Compass tool has to use generative AI to prepare the natural language summaries that respond to the user’s questions. To help make sure users appreciate that these summaries are synthetic, OCTO will provide notice of this fact to users upon their initial login to DC Compass. OCTO will continue to monitor the effectiveness of this label during the public beta testing phase, and adapt this strategy as necessary in later stages of the deployment.

III. Disclosure that Users are Interacting with AI

As language modeling technologies continue to improve, it is important that residents be told when they are interacting with a person, and when they are interacting with a synthetic agent of some kind. Throughout this Report, OCTO has described in detail the various roles AI functionalities play in DC Compass’s workflow. The fact that it is an AI tool has been communicated early and often throughout our public engagement efforts, and will be affirmatively disclosed alongside the synthetic data label discussed above.

By way of further explanation, DC Compass incorporates AI tools from Microsoft’s Azure OpenAI platform, and GIS tools from ESRI into the following five-step workflow:

Step One: User Interaction

The user inputs a question by direct voice or text, or by selecting one of the frequently asked questions specifically tested and curated by OCTO.

Step Two: Interpreting the User's Input

When the user submits their query, an AI assistant activates what's called a "civic agent," which passes the query through Azure OpenAI language modeling tools to interpret what the user has provided. This is where the user's natural language query is translated for the DC Compass system into questions it understands, like, "how many," "where is," or "how often."

Step Three: Incorporating GIS Tools

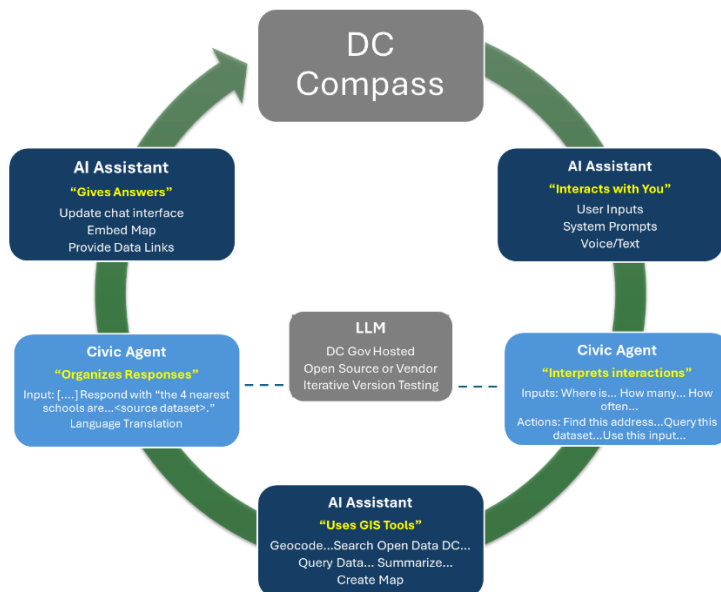
Based on the language model's interpretation of the user's query, the DC Compass system will select the appropriate GIS tools to use, and Open Datasets to search for responsive material. This can include geocoding, searching, querying, or summarizing. This is also the portion of the process where maps and other data visualizations may be generated.

Step Four: Formatting the Response

The civic agent is engaged to format the retrieved information into a coherent response. This is the phase where GIS visualizations are collated, responsive data is summarized with the help of the language model, and sources within Open Datasets are identified. This is also where translation back into the user's language of choice occurs.

Step Five: Delivering the Response

All responsive materials are provided back to the user. If the user asks any follow-up questions, they again cycle through the process.



DC AI Value Sustainability

New technologies can help OCTO to offer services that could not have been made available to residents before. But new technologies are not deployed in a vacuum. This section of the Report focuses on the sustainability of the DC Compass tool and is organized around the Key Concepts identified in DC’s AI Values Alignment Handbook.

I. Making Sure Costs Remain Sustainable

Maintaining control of costs can be especially difficult with AI tools. This is because many specific AI tools—including DC Compass—rely on integrations with enormously expensive, foundational, dual-use models. Only a half-dozen or so of the world’s largest companies have succeeded in building Large Language Models (“LLM”) powerful enough to match industry-leading performance. These tools depend on enormous amounts of data, computing power, and professional expertise that none but the largest companies and most powerful national governments could afford to duplicate. As a result, there are only about a half-dozen LLMs that DC Compass could switch to, if its current LLM provider were to propose significant price increases. The most effective way to manage such costs in a market that does not offer a ready supply of competing providers is to closely track a given program’s performance needs and the associated costs, so that strategic adjustments can be made as necessary.

To illustrate how these market conditions impact cost sustainable planning, consider a simplified hypothetical. If DC Compass were to use version 3.5 of a big, commercial LLM, OCTO may pay a cost \$1.00 per 100 user requests, and users may experience a response time of 20 seconds per query, with a given level of inaccuracy built into the limitations of version 3.5 of the LLM. Alternatively using version 4 of the same LLM may cost \$10.00 per 100 requests, with a response time of 5 seconds per query. OCTO could work with its partners and with end users to determine where DC Compass should land in this cost/performance trade off, anytime unilateral price increases were proposed by the LLM provider. Any change in LLM version or provider, though, would also entail the costs of re-design and re-testing before live deployment.

Another tool available to OCTO for maintaining cost sustainability for the DC Compass tool is use volume restriction. These restrictions can include limits for the overall number of queries in a given billing cycle, Tokens per Minute (“TPM”) rate limits to slow use and help to spread it more equally across a billing cycle, and more. Total query limits and TPM rate limits can help avoid accidental or wasteful overuse, but may also limit resident uptake, as someone who has a request denied due to a rate limit may not return to the tool in the future.

II. Environmental Sustainability Considerations

AI tools’ negative environmental impacts come primarily from their energy usage. AI tools use the vast majority of their energy during training of their neural networks, not during operation of the tools. Operation of the DC Compass tool does not require any additional neural network

training. Its use of LLMs can be fully supported by other-than-leading versions of most major LLMs. Currently, DC Compass does not use the most advanced LLM available through the Azure OpenAI platform. This helps to limit the negative environmental impact of the DC Compass tool.

DC Compass's ability to run on other-than-leading versions of major LLMs is owed both to the relatively low-risk domain in which the tool is employed, and to its retrieval augmented generation ("RAG") design. The low-risk domain (access and summary of public GIS data) helps DC Compass to run on other-than-leading versions of major LLMs because the risk to the public that might arise from a confusing or inaccurate DC Compass answer are very low. Considering a far riskier domain, by contrast, if an airport tried to replace air traffic controllers with a chat bot, even the slightest confusion or error would be accompanied by catastrophic loss.

DC Compass's RAG design helps it to run on other-than-leading versions of major LLMs because, as explained earlier in this Report, RAG uses a pre-identified data set to answer questions. A RAG only asks the LLM to perform the relatively easy task of turning natural language into a database query, and then summarizing in natural language the answer the database provides. Several versions of major LLMs are capable of performing this relatively minor task reliably. For similar reasons, OCTO's curation of frequently asked questions should help to limit the demands DC Compass puts on its commercial LLM, and in turn limit the DC Compass tool's negative environmental impacts.

III. Sustainable Job Quality for Impacted Employees

As a primarily public-facing tool, DC Compass is not expected to impact the working conditions of DC government employees. Though current employees' efforts at data collection, review, and publication are essential in making DC Compass work safely and effectively, these workers do not actually interact with DC Compass as part of the DC Open Data program. When an agency works with a given enterprise dataset, even if that dataset is later included in the DC Open Datasets, the employee almost always works with the agency's own record copy of the data. As explained elsewhere in this Report, DC Compass does not interact with the record copy of agency datasets. Instead, DC Compass draws exclusively from pseudonymized and anonymized copies of agency data made freely available through the DC Open Data program. To illustrate, if a DC government employee in education is working with data from schools, they would be working with a confidential and federally protected copy of the relevant dataset, not with a publicly available copy hosted through Open Data and accessible to DC Compass.

If, however, a DC government employee were to use DC Compass as part of their day-to-day work, DC Compass would have a positive impact on their job quality. DC Compass does not replace any of the fun or rewarding parts of being a government data scientist. It only removes the drudgery of long, technical query generation, and unnecessary sifting through irrelevant data. It could be likened to libraries' transition from paper card catalogue systems to software-based book inventory tracking.

IV. Carefully Considering Risks of Worker Displacement

DC Compass is not expected to displace any of the current DC workforce. DC Compass is not a faster or cheaper substitute for work already being done by humans. It does novel work, providing services that were never before offered to residents. If deployment of the DC Compass tool has any impact on DC's larger workforce, it will likely be a strongly positive one, if it proves successful in engaging the public and improving data literacy—a professional skill growing rapidly in importance in the age of AI—among residents.

DC AI Value Privacy

The DC Compass tool is designed to increase public access to 1,400 valuable datasets, so it makes sense to pay special attention to safeguarding important privacy interests. The primary privacy protection built into the design of DC Compass is its limitation exclusively to pre-screened Open Datasets. But these are not the only privacy safeguards incorporated into the design of the DC Compass deployment. This section of the DC Compass AI Values Alignment Report documents various relevant privacy concerns, and the steps taken by OCTO to mitigate them. It is organized according to the Key Concepts identified in DC's AI Values Alignment Handbook.

I. Identifying the Information Relevant to the DC Compass Deployment

There are two classes of data relevant to the DC Compass deployment. The first, DC's 1,400 Open Datasets, has been discussed extensively in this Report. The Open Datasets are not the record copies of agency enterprise datasets, they are carefully anonymized and pseudonymized, and they are subject to annual review by agency legal and privacy staff.

The second class of data includes the DC Compass user data—query and response pairs, ArcGIS.com account information, and user IP addresses. DC Compass users will be notified at log in that this user data may be retained and used to improve the DC Compass system. In planning the DC Compass deployment, OCTO has carefully accounted for both classes of data.

II. Maintaining Relevant Data in a Safe Location

The Open Datasets accessible by the DC Compass tool are stored in entirely separate locations from the agency record datasets from which the Open Datasets are drawn. The Open Datasets are hosted by OCTO or by Esri on a cloud infrastructure. Updates to the Open Datasets do not flow back to the record copies of agency datasets, which are stored on an entirely separate physical infrastructure.

The DC Compass user data is stored on Esri cloud infrastructure, and question/answer pairs that are passed from the DC Compass civic agent to the Azure OpenAI tools are maintained in the Azure environment in accordance with the Azure OpenAI terms available at <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?context=%2Fazure%2Fai-services%2Fopenai%2Fcontext%2Fcontext>.

III. Appropriate Data Use Restrictions

OCTO has long made Open Datasets available subject to a Creative Commons license that helps to ensure DC will receive attribution when the data is used, and that derivative datasets which users might develop from the Open Datasets must also be made available to the public. Because

OCTO is legally required to make Open Data freely available to the public, further restrictions on use of the Open Datasets would not be appropriate.

DC Compass user data is not made publicly available. The limits of Esri's use of DC Compass user data are still being defined while DC Compass is in its beta testing phase. When finalized, those limits will be recorded in an End User License Agreement. The limits of Microsoft's use of those question/answer pairs passed to the AI tools are subject to the Azure OpenAI privacy terms available at <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?context=%2Fazure%2Fai-services%2Fopenai%2Fcontext%2Fcontext>.

These standard privacy terms confirm, *inter alia*, that the question/answer pairs are not available to other customers, not available to OpenAI, not used to improve OpenAI models, not used to improve any Microsoft or third-party products or services.

IV. Technical Privacy Protections

The most effective privacy protections are those incorporated early in the design process. To help ensure DC Compass did not create new, unmanaged privacy risks, the design of the DC Compass tool incorporated five primary technical safeguards for the class of data accessed and surfaced by the tool:

First, as discussed throughout this Report, DC Compass does not reach beyond the pre-anonymized and pseudonymized Open Datasets. It cannot draw data from agencies' non-public, record datasets. It cannot write changes to those datasets. DC Compass effectively draws from a sandbox world where no non-public data exists.

Second, DC Compass is not directly exposed to anonymous web users. Instead, users can only access it after first registering for an ArcGIS.com account (registration and use is free to users). This technical design decision enables an important legal protection. By limiting DC Compass access to users who have agreed to ArcGIS terms and conditions, those terms and conditions become legally enforceable. This distinction—between data provided on the open internet and data placed behind a free registration wall that imposes terms and conditions—can be outcome determinative if a data provider (here, DC government) later tries to prevent an outside party (here, any group, organization, or individual) from flouting the terms of a license or use restriction.

Third, users have no ability to upload data into the DC Compass tool. This is an important limitation not just for protecting the accuracy and integrity of the Open Datasets, but also for preventing the DC Compass tool's misuse by users to disseminate data that has not gone through an appropriate privacy review.

Fourth, as previously discussed in this Report, OCTO has reserved the ability to activate volume or rate limitations on the DC Compass tool. While the primary benefit of these kinds of controls is to maintain cost sustainability, a secondary benefit is that these controls can be implemented to help guard against high-volume, automated queries that might be used to undermine performance of the DC Compass tool. It is important to note that these limits do not make Open Datasets less

available to those organizations who would like to process the Open Datasets with high-volume, automated tools. The limits would simply prohibit their misuse of the DC Compass tool, while leaving them free to download any or all of DC's 1,400 Open Datasets directly, through traditional means.

Fifth, OCTO provides agencies with tools for anonymizing and synthesizing data before it reaches the Open Datasets, to help maintain privacy without sacrificing consistency in linkage through similar datasets. DC procurement datasets offer a good example of how this technical safeguard promotes privacy. An identified business might appear across multiple datasets with different names, which might be due to data entry issues or that business's use of multiple doing-business-as ("d/b/a") names. Hypothetically, the same company might go by "Main Avenue Tech, LLC" for one line of its business, and "Main Street Tech Solutions" for another line of business.

On the non-public side, appropriate agencies have no trouble linking these multiple d/b/a names through the business's Federal Employer Identification Number ("FEIN"). Where the business is a sole proprietorship, or for some other reason lacks an FEIN, this same linkage can be made through the owner's Social Security Number ("SSN"). Instead of guessing whether Main Street Tech Solutions and Main Avenue Tech, LLC are close enough to be the same vendor, an agency with access to non-public datasets might simply compare the FEIN or SSN.

For obvious privacy reasons, DC cannot release FEINs or SSNs in Open Datasets, no matter how confusing d/b/as and other pseudonyms might be. To fix this problem, OCTO often indexes linked data to an otherwise meaningless pseudonymous identifier, so that Open Data users can still track any single business (or other meaningful cluster of seemingly unrelated data) with its meaningful relationships among data still functionally intact.

For the second class of data relevant to the DC Compass deployment—DC Compass user data—technical privacy safeguards are available at <https://trust.arcgis.com/en/privacy/privacy-tab-intro.htm> and <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?context=%2Fazure%2Fai-services%2Fopenai%2Fcontext%2Fcontext>.

V. Legal Protections for Relevant Data

The first class of data associated with the DC Compass deployment—the DC Open Datasets—are protected by a Creative Commons license, available at <https://creativecommons.org/licenses/by/4.0/deed.en>, and the ArcGIS terms and conditions, available at <https://www.esri.com/content/dam/esrisites/en-us/media/legal/ma-full/ma-full.pdf>, and <https://www.esri.com/content/dam/esrisites/en-us/media/legal/product-specific-terms-of-use/e300.pdf>, to which every DC Compass user agrees before they are permitted to use the DC Compass tool.

The second class of data, DC Compass user data, is subject to two separate legal protections. All of the user data is protected by the ArcGIS privacy terms available at <https://trust.arcgis.com/en/privacy/privacy-tab-intro.htm>, and will be protected by the end user license agreement still being negotiated between OCTO and Esri. Additionally, the subset of that

data which is handed off by the DC Compass civic agent to the Azure OpenAI language model is protected by the Azure OpenAI privacy terms, available at <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?context=%2Fazure%2Fai-services%2Fopenai%2Fcontext%2Fcontext>.

VI. Notifying Stakeholders of Potential Privacy Impacts

Users sometimes like to know before they interact with a tool whether using the tool will mean giving up some of their privacy. As discussed throughout this Report, the DC Compass deployment involves two classes of data: the DC Open Datasets from which it draws, and the DC Compass user data generated when users log in and submit queries to the tool. The first class, as has already been addressed, is exclusively public data, and therefore the DC Compass tool will not cause any change in the privacy posture of people who at some point upstream from the DC Open Datasets might have been associated with a given bit of information.

Concerning the second class of data, DC Compass users all receive notice of the ArcGIS privacy policies when they create an account. As part of its general public engagement work, OCTO will listen for any indications from the community that these notices are not sufficient for users' comfort. If further notice proves necessary, OCTO is able to add click-through notices concerning the treatment of DC Compass user data at log-in.

DC AI Value Cybersecurity

DC’s AI tools must be deployed in a way that promotes the confidentiality, integrity, and availability of DC’s information technology assets. Effective cybersecurity means robust encryption, access controls, and regular security audits to protect against unauthorized access, data breaches, unintended sharing, and malicious tampering. It means maintaining up-to-date software, conducting continuous monitoring for vulnerabilities, and conducting regular and thorough risk assessments. This section documents how OCTO planned the DC Compass deployment to promote cybersecurity. It is organized according to the Key Concepts identified in DC’s AI Values Alignment Handbook.

I. Secure Configuration

DC Compass is hosted on a combination of DC and Esri infrastructure. Virtual machines hosted within OCTO’s data centers provide the infrastructure for the ArcGIS Enterprise software deployments. This enables OCTO to publish RESTful web services for both geospatial and non-geospatial datasets. These datasets are stored within relational databases in OCTO data centers. The web services are then registered with the Esri-hosted ArcGIS Online platform, which contains an open data module and is the primary component for Open Data DC. Most datasets follow this architecture, with the exception being several datasets which have large numbers of records, such as property data. For these exceptions, we directly host and update the data within the ArcGIS Online platform to improve performance. These registered or hosted web services are the data sources DC Compass queries.

DC Compass itself is developed in LangChain and is embedded within the Esri ArcGIS Hub solution. Esri currently utilizes Azure OpenAI’s GPT 3.5 LLM for DC Compass’s AI functionality. This comprises the AI Assistant component of DC Compass, while the GIS functions are performed by Esri ArcGIS Enterprise and ArcGIS Online platforms. This enables geospatial functions such as query, point in polygon intersection, summary, search, geocoding, and map creation. The API key for Azure OpenAI will be managed through the Microsoft API Management Key Vault component, which enables use of the API key without exposing the actual API key. Esri is responsible for the uptime of the ArcGIS Online and ArcGIS Hub solutions, and OCTO is responsible for the uptime of DC’s ArcGIS Enterprise environment.

The following table illustrates the distribution of responsibility among OCTO and Esri staff, expressed in the common “RACI” format: R = Responsible, A = Accountable, C = Consulted, I = Informed.

Task/Role	Chief Data Officer (OCTO)	Senior Data Analyst (OCTO)	Product Manager (Esri)	Product Developer (Esri)
Initiation				
Establish Vision	R/A	C	R	I
Planning				
Use Cases	R	R	R/A	C
Cost Evaluation	C	I	R/A	R
Model Evaluation	I	C	R/A	R
Execution				
Design UI	I	C	A	R
Development	I	C	A	R
Testing	I	R	A	R
Security	C	C	R/A	R
Document	I	I	A	R
Closure				
Cost Reevaluation	C	I	R/A	R
New Use Cases	R	R	R/A	C
Model Reevaluation	I	C	R/A	R

Esri will be providing a method for OCTO staff to review queries from users to assist with improvements to underlying metadata and tool debugging for enhancements and fixes. These logs will not contain any specific user identification but will contain server-side logging of IP addresses and requests. The treatment of such DC Compass user data is discussed earlier in this Report. Its collection and retention is required for security investigations where logging is necessary to investigate queries or activity associated with risk, harm, or abuse.

II. Designing a User Interface that Promotes Security

DC Compass does not directly interface with any other systems outside of Esri’s ArcGIS Online platform and OCTO’s GIS infrastructure. It has no access to applications or data sources other than those registered for Open Data. OCTO manages the end-to-end workflows for data acquisition, data storage, data transformation, metadata publication, and web service publishing and registration. The team maintains Extract-Transform-Load processes to update datasets on a designated frequency in line with the dataset records in the Enterprise Dataset Inventory. Updates to web services occur on a regularly scheduled basis to avoid disruption to operations.

III. Data Security Considerations

The DC Compass tool will only access DC’s Open Datasets. All available datasets can be found at:

https://opendata.dc.gov/datasets/76a28737a6f84b3c92a421114accca2_5/explore?filters=eyJJEQVRBU0VUX0NMQVNTSUZJQ0FUSU9OIjpbIkxldmVsIDAgT3BlbiJdLCJJU09QRU5EQVRBIjpbIlkiXX0%3D.

IV. Special Consideration for Public-Facing AI Tools

DC Compass is a public-facing AI tool, accessible exclusively through ArcGIS.com registered accounts. Users of DC Compass do not have direct access to the incorporated LLM. The underlying LLM consumes tokens, and the ability to execute a single query or multiple queries is predicated on the number of tokens the model has available and dedicated to the tool. There will be no additional scrubbing, filtering, or pseudonymizing of outputs from DC Compass. This is because DC Compass is a Retrieval Augmented Generation tool which draws exclusively from the DC Open Datasets, which are regularly reviewed to prevent the inadvertent release of private or sensitive data.

V. Planning for Active Support

Esri and OCTO will work together on any updates or successor versions of DC Compass. Esri can provide sandbox versions of the DC Compass tool which utilize different models for testing purposes. All software used in operation of the DC Compass tool is actively supported by Esri or OCTO. OCTO's infrastructure undergoes regular security patching for operating systems, and application patches to the GIS software are applied as they are released and relevant to the version of the software currently in use. The cloud hosted platform, ArcGIS Online, undergoes quarterly upgrades. To minimize disruption to end users, the maintenance windows tend to occur in the evenings and overnight, depending on the platform and types of patching.

VI. Managing Risk

OCTO's DC Compass team is in constant contact with OCTO's Security Operations Center, and follows OCTO security policies, including OCTO's cybersecurity risk management policy. We will provide an appropriate email address for users to report specific concerns and any incidents.

The OCTO GIS infrastructure has multiple redundancies in place, for both databases and servers that host the datasets made available, including data center failover capabilities. The hosted environment from Esri is highly available and can failover to multiple cloud providers in the event of a major outage. Esri's ArcGIS Online platform is certified to be FedRAMP Moderate.