May 13, 2025 Meeting - Seattle Community Technology Advisory Board

Topics covered included: One Seattle Data Strategy; Responsible AI and AI Governance

This meeting was held: May 13, 2025; 6:00-7:30 p.m., via Webex and in City Hall Room 370

Attending:

Board Members: Isabel Rodriguez, DeiMarlon Scisney, Coleman Entringer, Hailey Dickson, Omari Stringer

Public: Dorene Cornwell, Sanchit Gera, Call-in User_1, Robert Kruse, Arif Gursel, Desiree Walker, Temi

Staff: Trayce Cantrell, Steve Barham, Brenda Tate, Jon Morrison Winters, Vinh Tang, Mark Schmidt, Cass Magnuski

19 In Attendance

Omari Stringer: Lovely to see you all. I'm filling in for Phillip Meng, who is unable to be here this evening. We have a little bit of a light agenda tonight, including some presentations from some City staff folks, and then I will be doing a presentation on responsible AI governance. Then we can do some quick updates and close it out. So, I will go ahead and start with introductions.

INTRODUCTIONS

Omari Stringer: I think that's everyone on the call. Next up, we will roll through the agenda here with the approval of minutes first, and then the approval of tonight's agenda. I move to approve the minutes. Can I get a second?

Coleman Entringer: Second.

Omari Stringer: Great. All in favor of approving last month's minutes? Abstentions? Nays? Looks like motion carries. And tonight's agenda? Can I get a motion to approve tonight's agenda?

DeiMarlon Scisney: I move to approve tonight's agenda.

Omari Stringer: Can I get a second?

Isabel Rodriguez: Second.

Omari Stringer: All right. All in favor of approving tonight's agenda, say 'Aye.' Any opposed or abstentions? Great. Both motions carry. And with that, we will go right into our meeting for tonight. We have Steve Barham and Mark Schmidt from the City, presenting on the One Seattle Data Strategy. I will turn it over to you for your presentation.

ONE SEATTLE DATA STRATEGY

Steve Barham: Mark Schmidt is going to run the presentation from the room. I will kick it off. We are here to talk about what we call the One Seattle Data Strategy. I think we have about 40 minutes and we will run through 25 or 30 minutes of slides. We want to tell you what is in our strategy for the City, how we put it together, and what's coming next.

Before there was a data strategy, there was something called What Works: City Certification. This is basically a program that certifies cities according to criteria related to data-driven practices, transparency and impact. So, for the last couple of times that Seattle has tried this, we have gotten gold certification, which means that Seattle is actually. as a City compared to other cities, Seattle has been pretty good at things like using data, tracking progress with data, and making decisions. The certification above this is platinum. Maybe next time, we are going to go for platinum. There are only three cities that are platinum right now, New York, Phoenix, and Tempe, AZ. We were certified in 2023, and also in 2023, we began participation in the Bloomberg Philanthropies City Data Alliance program. This program is something where Bloomberg Philanthropies provides executive coaching and education and training, and support for cities to do certain things. And for us, it was creating a citywide data strategy. This program required cities to be pretty good at data, which we were, and then it also required strong executive participation. And so, our Mayor took the lead and actually went to data camp for several days by himself to strengthen his data practices and lead from above. And so we set out joining the second round in May, and our goal was to create a three-year data strategy.

It was basically the second half of 2023. We didn't have any template to work from, but we organized and convened almost 100 people. There were 85 contributors, data leaders and practitioners across the City that came to help us put together this data strategy that represented almost every operating department, so 24 departments and we were organized into eight different work groups, which were formed around impact areas that we will talk about in a second. But we ended up putting everything on the wall, like what could we to impact data strategy in the City. We came up with 21 things that were practical and feasible, things that we can do in the next three years, things that involve policies, changing City standards, building tools that people can use, and also looking at different capabilities and roles for using data across the City.

That was a lot of work for all of us for six months, Mark Schmidt can attest to that. At the end, we had a party, and a party for data people is a hack-a-thon. So we partnered with Tableau and Salesforce to host a hack-a-thon at their headquarters in Fremont. Part of

the hack-a-thon was looking at what can citizens and students do, using some of our open datasets. So, it was a really good way to highlight parts of the data strategy, especially advancing open data, using data to help solve some of the bigger issues in Seattle. Some of the datasets we used -- I think the Mayor really liked that we used data for the Downtown Activation Plan, and data on permits. We also used library data, which was a little bit less interesting. But this was a great way to launch our strategy, and at the same time, we also enacted our strategy through an executive order.

Like I said earlier, we didn't have a strategy on how to create an updated strategy. We just had some coaching and some ideas. But there is a template now and it is because Seattle is now the template for how to create a citywide data strategy. So we think it is nice to be (unintelligible), and this is how we are highlighted from the Bloomberg Center for Public Innovation who is now running the City data alliance.

That's how we got started, but why did we need a data strategy in the first place? I think it's really about unlocking the power of data. I work with a lot of projects and I've been with the City for a long time. There is so much data out there that we already have that we're not using. One of the reasons why we are not using it is also because we are working in silos. Mark Schmidt can attest that we also have to balance utility and risk. We are trying to get from doing things well to doing things really well, and some of our departments do things really well and we want that to be across the City. And the last thing is some of the challenges that Seattle faces -- we call them wicked problems -- but the complex multi-dimensional problems where we need to come together, it's really hard to do that without leveraging the use of data.

Some of the things that I've seen over the years in data projects is a project ends up in a drawer; you can't reproduce what your team has done; data hasn't been updated since the Sonics were playing in the Key Arena; departments are talking apples and oranges and can't work together. And sometimes, and this is what I see a lot is what you really need is maybe a fishing boat project but what you design is a cruise ship or an aircraft carrier. Or maybe it's the other way around. You need an aircraft carrier, but all you need is a little tiny fishing boat. And lastly, what's I've seen a lot are individuals or groups or teams that might have some data skills -- just enough to be dangerous, but lacking the context of working with the City and best practices. I call that the danger zone. But data strategy is meant to address a lot of those challenges. So, I'm just going to go through some highlights in the strategy itself, and maybe not go through all of it. But we started in 2023 working on it. It's a three-year strategy, so really we are in the second year of the three-year strategy.

This is how we visualize it into three different components. The first one is really around how we use data, how we analyze data, including evaluation and public engagement. The second one is how we govern and manage our data. And that's a really important piece. We can't do the first one without it. And then, the third one, data literacy and culture, is really upscaling and building the capacities and capabilities of our staff. We've branded the strategy with this color scheme and icon, but the black band in the middle represents the fact that one of the most important things is collaboration and coordination between us, so we need that to be strong. And then, there is a strong equity component when we put this together. Data equity, equity analysis, is a big part of the strategy. I will highlight a couple of those 21 elements that are actually in there. The next slide shows you what is in the first group on data use. For me, again, this is really unlocking the power of data, and some of the things that are part of the strategy what we have internally in the evaluation framework, so how you evaluate programs using data, a lot of resources on analytics, different style guides, and an emphasis on things that are public-facing, like dashboards and portals, and doing better at that. But really, the goal, when we think about it, is in terms of how can we achieve better outcomes. And so, when we are designing that, that is always the question that we are asking.

The next one around data governance and data management, part of this is looking at how can we have a citywide framework for that, advancing our open data practices. We have some tools and resources for City employees to use when they're thinking about data quality and other data standards. And then, also a big part of this is really looking at the roles and functions and responsibilities for people who are either working with data and using data. I like to think that every person who works at the City can be a data analyst. The last one is a segue into that in really empowering everyone. So, we want people who maybe are very domain-specific to be data analysts and then we want our data analysts to also be able to be trained in the different programs and issues that our communities face. So, part of this is people coming together and thinking about what is a good workforce development plan that we can have and developing the capabilities. For me, I'm a data scientist, so roles like data analysts and data scientists don't necessarily have very specific positions in the City for that. So, one way is really to think about what do other cities use, what is the industry standard. And then, another thing that we did to bring everything together was to develop an internal collaboration portal, and that's a big piece that we did for the data strategy, that other cities didn't necessarily have. You can go online and read the data strategy. A lot of it is geared for an internal audience, but we are also very transparent about what we put in there. We have a tracking system and we are going to be updating the tracking system probably this summer to see where we have made progress in some of the things that we may have started and not yet finished. The idea was that we were going to get a lot done at least by the third year, but we have already completed a lot of these things.

We tried to make it shiny and nice looking and everything, but really, I think for a lot of people, data strategy itself is not a super exciting topic, but what we've found is that what really gets peoples' interest up is really thinking about impact and then thinking about the different projects. So, we went around the City and identified a whole bunch of projects where we are advancing data practices. And some of these are small, like the Pea Patch. And some of them are larger, like the On-street Paid Parking Rate data, which I think is the biggest dataset that we have at the City. It's on Open Data. I think it's the biggest Open Data source that we have. But, was finally just publishing those data and being able to use it to make adjustments and changes in policy on the street. That's a cool thing. In the end, our data strategy is actually a 30-some page document, but over half of it is really looking at specific projects, because that is what we are

emphasizing. That is one of the reasons why Bloomberg Philanthropies liked what we came up with as a template.

I'm going to go through a couple more slides and then I'm going to pass it off to Mark Schmidt, but I wanted to also convey a little bit about how we are continuing to advance these practices. Here are a couple of examples. Just in 2025, this year, an executive order on improving accessibility of the City data information, which is really about doing things in plain language, digital accessibility, making our communication, our outwardfacing products a lot better. And that is in alignment with our data strategy. Another thing that came out recently, as well, at the beginning of the year was our strategy and plan to work on anti-displacement strategies, and part of that is upgrading, updating, and making improvements to our public-facing dashboards around displacement risk, and evaluation, as well. So, the things we established in evaluation we are using for anti-displacement. And a lot of that data that we're using for this is also being published on Open Data and available for anyone. Another thing that we're trying to advance now, another thing that we're trying to do is really think about how we use data for collaboration. For cities for a long time, there is a model called performance stat, which is basically a way to get different people and different functions in the same room, having a strong executive leadership, ask hard questions using data, try to solve problems and get things done or change things or to try to provide accountability or transparency. This process has been around for a long time, and what we're envisioning and reimagining how a process like that happens. One, by democratizing data or broadening the way that data is available for all of the collaborators in the room to have better access, and then speeding up the analysis, so being able to ask and answer questions faster, we can get to more ideas and better solutions. So, that's where I think we are taking a lot of this in a very practical outcome-based way.

I will pass it to Mark.

From chat: Steve Barham 5/13/2025 6:32 PM • Here is the link to the data strategy: <u>https://seattle.gov/mayor/one-seattle-initiatives/one-seattle-data-strategy</u>

Mark Schmidt: Hello, everyone. One of the outcomes of the data strategy and the announcement that the Mayor made in December of 2023 was an executive order and issued to do the four things that you can see here: to encourage all departments to participate, to develop implementation plans, to refine our demographic data standards, and to establish a data governance board. I will talk about a couple of those different things and the work that we're doing to achieve those.

The implementation plans: What we did -- you see the same color theme that was used in some of Steve Barham's slides around those focus areas. We are trying to glean from the different departments in a purposeful may how they use data, how they communicate, how they govern data, etc., so that we could identify some of the best practices within the City, so that we could identify areas of need, so that we could kind of center in on some things that a governance committee might want to work through together. And the way that we did this was that Steve Barham created a questionnaire

that went out to the departments with 30 questions. About half of those questions were self-assessment questions, where departments were asked to rate themselves on a variety of things from a scale of one to five. I will show you an example in a second. And then the second half of those questions were more open-ended around some of their specific plans or some of their capacities or some things that they would like to see come out of this data strategy. Here is an example of some of the data. We asked all of the departments to rate themselves. We use data regularly to make our programs and investments more effective or efficient, one being the low end of the scale; five being the high end of the scale. These are just a couple of the data points. You can see that most departments were somewhere between three and four aggregate score of about 3.7. It gives us a starting point to work with from the data from these different departments. Again, there were about 15 of these questions on the assessment, and then a bunch of open-ended things that gave us some topics for planning. Another thing on that executive order -- the fourth directive of the executive order was that we would stand up this executive data governance board, which reports to the Mayor's IT subcabinet, which as you probably know, department directors that provide direction to the IT department, and provide recommendations around policies or standards that affect the City broadly. So, data quality, data strategy, etc. And the directive is that all departments participate by designating a data governance champion to participate in the process. It's a decision-making body. It is also kind of a collaboration. They are business leaders, they are not necessarily technical leaders in their departments, but they are business leaders that understand the impact of data on their operations, and they come together to make decisions with a focus on product. When we say a focus on product, really the idea is that we're going to identify some specific work products and get some things done. It's not just a forum for discussion. It is a forum for discussion, but we're also trying to achieve some specific things that will move the City forward. As part of that, some members of the board will also be invited, encouraged to participate in some of the working groups that prepare the things that are discussed at the different board meetings. You can see the dates for the 2025 meetings along the top of the slide. We just met last week on May 7, and we'll have another meeting in August, and a final one in October. We are still solidifying our work plan, but are pulling things out of that One Seattle data strategy, looking at how we govern data and stewards, and finalizing a demographic data standard, and looking at workforce development in line with some of the data culture that Steve Barham talked about earlier.

In the meeting late last year, we queried the group, did a white board exercise to encourage people to identify some of their priorities, and the themes for 2025, and you can see them listed there. We would focus on really looking at the data that we have, and who is it that manages the data and gets to make the decisions about that data, highlighting resources that we can share, things that we already doing very well, sharing our knowledge and ideas with one another. Advancing open data -- I'll talk about that a little more in a second. And finishing up that demographic data standard.

Advancing open data: In the policy specifically, here is my whole pitch in two quick bullets. The old policy is outdated, and no longer aligns with how we do business. And we said, as part of the data strategy that we would update that so that we could better

leverage that data as a shared asset. Here is a quick refresher. Some of you probably know some of these things, but in 2009, that's when the federal open data government movement began, Seattle was an early leader. In 2010, we worked with Socrata, a local company down the hill, to launch our first open data platform. We were a very early leader. In 2016, we were one of the first cities across the United States to draft and execute and open data policy, accompanied with an executive order to direct the departments to participate. As Steve Barham hinted, in 2017, and 2020, and 2023, the City chief gold level certifications through Bloomberg, open data is a key component of that. Having a portal, having a process where you are regularly engaging with the public to query interest and value in that data, to advance a key component of that. And again, that is what led us into the City Data Alliance that Steve Barham talked about, and the open data strategy.

Where is our open data portal today? On https://data.seattle.gov/, so today, there are 138 tabular datasets. If you would have asked me that question a year ago, I would have told you there were 300. And you would say, aren't we going in the wrong direction? And the answer is what we really did in this spirit of the One Seattle data strategy is looked at ways that we can make our data more useful. Not just publishing the data for the sake of having a lot of datasets, but aggregating lots of small ones into big ones so that you can visualize data over multiple years, really focusing in on those datasets that get a lot of usage, focusing on improving the quality of those datasets assuring that they're current and regularly updated so that they would be valuable. That was a big effort for us over the last year. Currently, 19 departments, more than 1.6 billion rows of data. Steven Barham talked about the curb dataset. Each of those datasets is 200 million rows for an annual dataset, so there are some really big datasets on the portal. The automation: A little bit more than a third of them are automated. We have some new tools available, where we can pull the data directly from the sources, from SQL databases, from Oracle databases, from Excel spreadsheets, and our Sharepoint environment to completely automate the data flow, which reduces the manual work by departments and increases our data accuracy, because it removes the chance of human error. That's a big push, because again, and we'll talk about the principles of the new policy that we are proposing, timeliness and accuracy are high up there. If the data is not good, the data is not valuable. We are federated with multiple other sites, as well. So we can link to 541 geo-spatial sites to a number of dashboards hosted by police and our court system. and we are also federated wit the US government sites. So, you can search for Seattle climate data on the https://data.gov/ site and bring up some of our datasets. Our current priorities are to align some of the objectives, like the One Seattle strategy. The user experience piece, again, we redesigned the portal front page last year, and recategorized a lot of the datasets, updated a lot of the metadata with the intent of making it easier for people to find high value public data. And then, there is also a focus on stakeholder engagement. What we're trying to do is work closely with departments as part of redesigning the portal. We are getting a little bit more proactive in terms of the way that we respond to some of the suggestions that come into us through the portal for data that we might publish, and looking at ways -- and I'll talk that in a minute -- for how we can engage people across

the departments to decide what is really valuable to publish, so that we're not just pushing data for the sake of pushing data, but that it is useful to people.

So again, the why of the updated policy: The existing policy defines our stance as open by preference, which is a little bit ambiguous. And we're the only ones in the country that do that. The challenge, the way it was written is it says all Seattle data is open by preference. What does that mean? What you will hear is that we're sp9inning it a little bit to say that we are going to acknowledge some things about our data in terms of privacy or compliance requirements or whatever else. Once we get past some of those requirements, we really want to say that we want to be open by default. That will put us again as a leader. Leading jurisdictions are open by default, properly qualified. And I'll talk a little about what that means. The other thing is we just cleaned up the policy a lot. It had a whole lot of procedure detail, and now you will see its principle space.

These are the principles that we're calling out in the new policies. The first one is open by default. It's exciting for some people, and it is scary for other people. Let me read for you guickly the way that we described what open by default means to us. And that is that the City strives to make data openly available for use, reuse, and redistribution by anyone, anytime, anywhere. For this to work, Seattle residents need to feel confident that open data will not compromise their rights to privacy. Similarly, City departments need to know that open data will not compromise security, confidentiality, regulatory legal requirements. Unless subject to these constraints, all data created, collected, and maintained by the City is presumed to be available for publication as open data. So, those last little bits about available for publication, that is not to say we are necessarily going to do it all, and that goes to the second principle, which is published for a purpose. And the fact behind that is simply that publishing more data does not necessarily yield more valuable data, that we're not just going to put a bunch of junk up there that makes it hard to find anything; that we're really going to be thoughtful about that, and we're going to adopt a prioritization process, which I will talk about really quickly. The rest of them are fairly intuitive, timely and comprehensive. Getting less data is relatively recent unless it is maintained and kept current, unless it's comprehensively complete, it's really not as valuable. It need to be accessible and we need to be able to find it. It needs to be open and available to people without license, which is an open data principle. We are putting it out there because we want to be transparent. We want to demonstrate good governance. We want to encourage engagement with the public, with departments, with other people using our data. And of course, we want it to be out there for economic development and innovation, and for whatever creative purpose people might use it.

Last slide is this prioritization idea. Again, publishing what is valuable. What we have adopted in our playbook is just a simple evaluation matrix. We give everything on a scale of zero to three points, zero meaning not applicable at all; one being low; two being medium; three being high. And we rate them asking questions around either strategic importance -- you can see some examples -- whether there is benefit to collaborative partnerships; if there has been expressed external interest in that information, if it is useful internally for improving operational efficiency, if it is good data,

complete and timely. And sometimes it is just a quick review. We already have a dataset that is really easy to publish. We just add up the scores. The higher is better and it is our intent to, as part of engaging with departments is to make decisions together about what we prioritize in terms of some of those datasets, because the work does take resources and departments are busy. And what we really want to do is honor peoples' time and also publish good data.

The last thing I will offer is that we have updated the policy and socializing it internally. It is our hope and intent to present it to our governing board, which again, is the Mayor's IT subcabinet, as a recommendation that we move forward with that policy and probably start implementing and prioritizing things early next year.

So, that's what's up with open data. Any questions about the data strategy or open data?

Robert Kruse: In the discussion about open metadata, ontology, RDF. If you take a look at data.gov, since 2009, they have had RDF as a file type. That allows semantic and AI integrations, which take us from data to knowledge and wisdom.

Mark Schmidt: No. Other than what data and metadata is on the platform or just little bits of -- we have actually created datasets about the datasets, so you can see the general number or the categories most frequent things, but I will say this. One of the projects that we're taking on this year is to look at the way that we tag datasets. So, there are two different ways that datasets are categorized. One is that they are lumped into categories, which is a metadata field, and we use that just for purposes of displaying them on the search portal. But we also allow on our platform tags, and there is no uniform standard way to the tags. One of our goals of this coming year is to actually review the tags that are there and do some thoughtful integration work. We are actually doing a pilot, not using AI -- we thought that maybe we could use an AI tool, but we have some documents published or linked on the open data site for our Office of Police Accountability, closed case summaries. The open dataset is essentially just an index today that just points to a case number and a disposition. And what we would like to do, what we are looking at, is ways to use technology tools to mine the documents to find themes in those documents that we could then glean metadata from, develop an ontology and then glean metadata from using automated tools, align them with those categories, so that then the public can search on excessive use of force, or racism, or whatever might be the allegation in some of those specific cases. We are probably not doing anything sophisticated nearly as you asked, but we are trying to look again at the way that we are actually tagging metadata or tagging our datasets so that we can enhance the search experience for the public, whether that is AI doing it or humans on a search bar.

Robert Kruse: I appreciate that. If there is an interest in Seattle leading in a really important category, which is wide open and something I did speak with the Mayor about, it would be the lead in the area of computational law. And I bring a network of experts in this area. And so, when you look at your legal code, start with the executive

order, but when you go all the way to the top for the Washington RCWs, there is the question that I raised to the Mayor. I said, how do you know that it has been properly encoded by AI? There is a range of AI tools that I just presented to you, which are W3C standards. You can go to https://www.w3.org/ and find the standard on semantic web. It is very important. Paul Allen was funding this kind of work 25 years ago. Europe is way ahead of us in this area, but there are some who want to plug into Seattle from Europe right now that would allow us to really help to raise the bar to another level so all cities can start to have our computational law capabilities, starting with the legal ontology. I could point you to one of those, if you are interested.

Mark Schmidt: I would love to have a follow-up conversation.

Robert Kruse: I'll put my email in the chat, and I am happy to hear that. I'm really glad to see what is happening here. It is hard to do. But taking it to the next level is making sure that people get the semantic meaning of it.

Mark Schmidt: And I did read something, for a long time it is required that open data needs to be machine-readable. That's fine, but when you talk about semantics or the context around the data, or whatever else, also needs to be readable and transferable. Machines need to be able to interpret the context, not just what is in the field, but the context surrounding it with some of the metadata that you're hinting at.

Robert Kruse: I would be happy to share a tool that comes from Geneva, Switzerland, that has already been vetted by (unintelligible) out there, and is directly pointed at privacy. It is an advanced implementation of GDPR and on ontology. It might be instructive to take a look at it, and then a privacy team. Alternately, the data is yours. We own our data. But your hard data could be monetized and turned into money. If you could imagine paying off your car with the data generated by your car, that is indeed a thing. We are kind of on the path of how do people, not just older data, but can make money for their mom, their grandma, and how to do that. If you are interested, I would be happy to invite a guest and teach her some of this capability. Thank you.

From chat: Robert Kruse 5/13/2025 6:48 PM • rkruse@venlogic.com, 206-726-9656

From chat: Robert Kruse 5/13/2025 6:52 PM • We will address how this works using semantic control of privacy consent.

From chat: Robert Kruse 5/13/2025 6:52 PM • We will address how this works using semantic control of privacy consent.

From chat: Robert Kruse 5/13/2025 6:53 PM • Building your City Ontology will solve for this...

New Job: Chief Ontologist

DeiMarlon Scisney: I have a couple of questions. We just talked about ownership and things of that nature, and I'm all for that. But I'm curious, just on the management side internally. Who currently manages access to department data? And I know that the governance plan is in place around that, but I'm curious about what kind of policy is in place to ensure security of data, and then who currently manages access to that?

Mark Schmidt: I guess it's a complicated answer. I think the security of data is either managed by application owners -- ultimately, it's the departments that own the data, or the people who operate the application. In a lot of the applications, themselves, are very secure. We don't do a lot of reviews internally with security folks, privacy folks, and whatever else before things go live with the different applications. One of the goals of the governance exercise, because data flows between systems, or we use data for analytics. Maybe I pull data from one source, but then I put it somewhere else, and use it for analytics. It's really who gets to make the decisions about how the data may or may not be shared. So, I think our systems are very secure. The question now gets to be when we get around sharing it or using it for analytics, who gets to make the decisions. I can say Steve Barham can use my data for six months, understanding the context of the project within which he is using it to ensure that it is consistent with the way that I generated or collected that data. But then after that project, then maybe I say that Steve can't use it anymore. Who gets to make those decisions about the ownership? What we are imagining for some of the governance stuff is purposeful discussions about some of that ownership. We talked a little bit about transportation data. To me, it's intuitive. If we're talking about managing curb spaces and parking policy to raise revenues and blah, blah, blah, all of that data belongs to the transportation department, and in terms of how the data might be formatted or whether or not it is shared, it seems that that ought to belong to the transportation department, because that is integral to their business mission and to operating their department. There is other data that transcends City departments. You say that City employee data belongs to our Human Resources Department. They may operate the application that stores all of their HR data, but a lot of departments across the City are going to have opinions about how that may or may not be used. What we are seeing in those instances where it really transcends City departments, that is where we might look to an executive data governance board for some kind of guidance. The idea is to look at types of data and create these data domains or decision domains and say, 'this is clearly transportation data,' or maybe, 'this is clearly electric utility data.' and the electric utility can make decisions about it. Some data transcends or they're sharing between departments, and that's where departments need to come together and collaborate. And the goal, ultimately, is to look at the major classifications of our data and understand who gets to ask the questions that you just asked. So, who is managing and who is responsible for stewarding it to ensure that it is secure, that we have good management practices in place, that we're managing across its lifecycle, that we're retaining it as long as we should, that we're disposing it when we should no longer keep it, to ensure that we know we can trace the lineage, all of that. The short answer today is the departments are largely doing it on their own. The downside of that is it has created some of the silos that Steve Barham talked about early on. And we're trying to address it by clear ownership, and then broader collaborative conversations.

DeiMarlon Scisney: Got it. You touched on issues around the sharing, and you just went through a long process. I as an AI engineer from Amazon, lunderstand what you just said. And we will have data teams that are doing a lot of that, causing the siloing internally. So I'm curious. Steve Barham had mentioned that there was a workforce component, and the training and certifying. So, I'm curious about what you all are doing to upscale internally and ensure that people are educated on the sharing, or educated on the lifecycle, and things of that nature that is really important when it comes to community data, for example, or demographic information, or whatever that may be.

Mark Schmidt: Steve, if you want to talk, I'll let you talk.

Steve Barham: Sure, really quick. I know we're probably going overtime. The policies are really citywide, so each department doesn't have to come up with their own policies for security, privacy. Where maybe in the departments, specific control policy. One of the things we are doing that was part of what Mark Schmidt mentioned in the executive order, is that we have actually gone to each department, and we have designated a data governance champion. So, every department has a data governance champion. Almost every department has created their own data strategy, and so for capabilities and needs, at least now there is a central person who knows what their capabilities are and what their needs are. And the different policies or different strategies can flow through that person. That's really the extent. We don't go, as part of this, to each department and go through their business, but we are relying on the board, the individual governance champions, and then for each department to have their own strategies.

DeiMarlon Scisney: Got it. Let me clarify. It wasn't about each department getting a policy, because that is outrageous. I know that. It's more so. How are you all educating or ensuring that people are educated. It sounds like all of that is being stored through the data governance chip.

Steve Barham: Across the City, there's a lot going on. Within the Seattle IT department, there are a lot of people and part of their focus is to help train and support people with those kinds of needs. There are communities of practice -- not just the ones we talked about today, but there are also communities of practice where people that are using data for things are coming together. And those are the forums where we know people are doing things. We provide support and cooperation and training to help them figure out what people are doing wrong and what they need across the City. There are probably three or four groups like that, not necessarily formally tied to a structure like a data strategy. But there are a lot of places where that occurs. I would think that the domain is really in Seattle IT for that kind of thing.

Mark Schmidt: At your last meeting, I think Harvey Arnone came and talked a little bit about GIS. Harvey has been named the interim director of a new division called Data Enablement, within Seattle IT. Harvey's group includes GIS, geo-spatial stuff that

includes computer in-design. It includes data, engineering, and analytics. It includes the open data program to kind of pull the data practitioners together and our workplan includes things like articulating and advancing a data governance strategy, aligned with the spirit of the One Seattle data strategy. So, figure on governance models for different departments where we can consult with others and tighten up governance in our own department. It includes an implementation plan and what we, as Seattle IT are going to do as part of the implementation plan, which includes bits and pieces around data literacy and new data services, and make sure that we are offering the right data analytics services to partners across the City, to imagine what data science services might be and how we might support data sciences for the occasional use cases where somebody needs to grab one of those massive 200 million line datasets, to make good, smart decisions about parking policy, because that's what SDOT does with all of that curb data. We are trying to build that out. We have several initiatives that are focused on people and process because a lot of these technology things are people, process, and tools. So, we have another workplan item where we are trying to develop a roadmap and a program over multiple years where we have better enterprise data management tools, so we can begin to help departments manage data across a lifecycle, focusing first on a few foundations of governance, tools that support governance. Once we get a few data quality tools and get some consistent standard ways of building data pipelines with governance built in so that we can scale across the enterprise. To Steve Barham's point, the IT department has several work items to say it's a new area that needs additional focus. It now has an executive that sits around the executive round table and IT department, so data is getting a lot more attention. And we're trying to build out that capability. Hopefully, a year from now, I will be able to give you a much more crisp answer. But your questions are awesome, and I realize that it is of strategic importance to the City and we are working in that direction.

DeiMarlon Scisney: Thank you so much. I just know that when you started. you said something around you are trying to train -- I can't remember exactly what it was, but

Mark Schmidt: The data curriculum will be part of the data strategy implementation plan. Like Steve Barham said, ideally, everybody would be a data analyst to some level, or at least be conversant and understand how data helps to drive decisions and do good evaluation, etc.

From chat: Steve Barham 5/13/2025 7:04 PM • @ Robert NY and London seem to be practicing what you were talking about, will look into this more. I consume our own open data, and this will add a lot of value

DeiMarlon Scisney: I appreciate it. So, thank you for the context.

Omari Stringer: That was a great conversation. Again, thank you to Steve and Mark for the presentation of the One Seattle data strategy. I'm really excited to take you up on that offer to come back and definitely update us on some of the progress. I think DeiMarlon Scisney had some great questions. So, it's definitely something that we will want to keep an eye out for. Open data is one of my favorite things about Seattle. In

fact, I always encourage people to check it out. I am probably one of the millions of users of the fire department's 911 dashboard when I hear fire trucks screaming down the street. So, I really appreciate all of the work you guys are doing. It has strategic value and bringing it throughout the City is no short order, so I appreciate all of your efforts. Thanks to you all, and I think you are going to keep us moving on our agenda. Next up, and I kind of keep it very light for this portion.

RESPONSIBLE AI AND AI GOVERNANCE

Omari Stringer: This is a general primer. I have done a couple of versions of this presentation on Responsible AI and AI Governance. Obviously, AI is a very hot topic across every industry at this point, touching every corner of the economy. But as I have had a career in privacy and compliance, I think AI is the newest tier and foundation for that. I'm going to talk about the intros, the 101s, of AI governance, with the particular focus on the NIST Artificial Intelligence risk management framework. And for those of you who aren't familiar, NIST is the National Institute of Science and Technology of the US government. They set a lot of these standards and measurement, a big scientific body in the US out of the Department of Commerce. They have done the NIST cybersecurity framework, which a lot of organizations, both private and public, use to set a baseline and maturity for cybersecurity. So, it's good practice, a set of controls. They also have the NIST privacy framework which kind of mirrors the cybersecurity framework and the privacy wall with a set of controls or activities around data processing that people can do to ensure privacy or security and ensure confidentiality, integrity and availability.

I want to start off with understanding what AI is and governance. There are a lot of different definitions out there. Basically, you want to think about it as a system that can perform some tasks that would require human intelligence, especially with the development of large language models. We are seeing things that kind of mirror human intelligence in a lot of different fields. Interestingly enough, it feels like a sci-fi thing, but there has been this test in academia called the Turing Test, named after Alan Turing, who was a famous computer scientist. This test basically sets up participants to see if they can judge whether or not they're talking to a person or a computer on the other end, and we have far exceeded that at this point, where a lot of the large language models convincingly fool someone that they are speaking to another human. When we think about traditional software versus AI its about mirroring that type of human intelligence or learning.

So, when we are talking about governance and AI, it tends to be a little bit different than what we just heard about data governance, because the implications of the use of AI are a little bit different and broader. We are making things in an ethical framework rather than just a traditional software framework. So, we're looking at things like bias or accountability that we may not see entirely. They are concepts but the impact is a lot larger. By establishing some of these rules of the road, having good compliance and governance frameworks, we do promote trust in AI technologies. Just as we want to promote accuracy and quality, we can trust that the data is accurate. We want to be

able to trust that the systems that we are using or interacting with have thoughtful considerations built into them, and have safeguards to prevent harm and bias. Thinking about frameworks like the NIST framework helps guide organizations towards some of those responsibilities and practices, because it can be a very nebulous place to start. And for a lot of organizations, they may not have a lot of data scientists but have an emphasis on staff. So, thinking about how do we translate what responsibility, what ethics look like, into a software framework. Again, thinking about responsible AI, it prioritizes these big concepts of fairness, transparency, and accountability. Whether that is implementing some guidelines to mitigate bias or making sure that we are getting stakeholder engagement, so not doing any of this in a silo, not thinking that I am just an Al developer, I'm a machine learning scientist. I want to develop the best algorithm in the world. It's taking a step back and thinking about all of the processes involved, the data involved, and the outputs. How are those going to be used? Who is going to be interacting with this system? What do the end users think? Are there ways that we need to disclose what we're doing. Are we going to be tricking or manipulating anybody. These are things that you want to think about. And it's really important to include other stakeholders than just the people involved in the project to get that outside perspective.

Responsible AI tends to want to maximize the benefits while minimizing the harm. I, myself, am very critical and very skeptical about Ai, but I think at this point, we all have to recognize that it is here to stay. That being said, we do need to think about what are some of the potential benefits of it. There has been a lot of improvement to things like accessibility and other ways to make these systems available to a wider set of audiences, but those things come with trade-offs that we need to consider, as well. So, getting organizations to pivot to adopt that ethical mindset towards AI, and not just say we're a private organization where the only motive is profit, but the motive is how can we build trust with our users so that they will keep coming back, as an example.

These are not theoretical harms. These are real things that have happened when bias in systems shows up in a variety of different way. During the Covid 19 pandemic in 2020, the education department used an algorithm to determine the grades of students because they were unable to sit for the test. They used a variety of factors, but what was found was that it disproportionately downgraded and scored lower students from lower socio-economic status, leading to a lot of criticism, because this is essentially, the SAT for US students. This is the determining their futures. They used an algorithm set up in another way to input or predict their performance. That was a big one that we've seen in the pandemic force all of us online to see that all of the results from that and all of the things online have real world consequences. In criminal justice, there was predictive policing to determine where police were dispatched, where resources are used. And this comes from a real need. I think I want to recognize in all of these use cases, that there is a need and a use case for technology to improve some of these outputs, but when you don't have thoughtful consideration behind it, you end up with some of these outcomes that may be worse and not influencing technology at all. So, for example, this one is pretty famous from 2018 or so, was used for granting bail. So, looking at the probability that someone would reoffend if they were granted bail, and it produced a higher risk score for black defendants rather than white defendants, without

taking into account things like the severity of the crime and other circumstances. So again, that's where we see a lot of uproar about the usage of AI or algorithms in the criminal justice system in particular. And finally, in the housing case, and this one is pretty common, almost all of us have sat in some kind of online application or other, and you get the little spinning wheel that tells you if you are approved or rejected. That was an AI system. So, already we're seeing these with impacts to economic or housing availability. But we have seen and we have known for a while that some credit scoring models and other systems that are used to determine validity or credit worthiness in a variety of instances, minority applicants are more likely to be denied. What's interesting about these is that there is often no recourse. So, you get that spinning wheel and you've been rejected, you can get a reason why you've been rejected, but there is often not a human appeal element. Some of this touches on the privacy, where we've seen legislation like the GPR in Europe or the CCPA in California mandate some rules around automated decision-making where things have a significant impact on a human, there needs to be a human reviewer, or the ability to appeal that automated decision. So, you see where keeping that ethical value of having a human in the loop, as they call it in machine learning, you want to be able to make sure that significant decisions have some way for a human to review that, or say, 'oh, the machine did make a mistake,' or there are extenuating circumstances in which we would reconsider the decision. Those are some real examples. There are far more examples of not just harms, but just weird things in addition to harm. There is a recent case in the UK of a department wrongly flagging people for possible fraud and error. I'll get on my soapbox here, but you've also seen things related to this coming to the US, with DOGE as an example of putting too much faith in AI or algorithms, and saying that we can massively reduce fraud or other government waste, but these systems aren't always accurate. So, that's something that you just need to keep in mind, a healthy skepticism when people claim large amounts efficiency is gained at what cost, and is it accurate? Another one was a manipulation tactic of researchers from Switzerland using AI to falsely interact with users on Reddit, and did not disclose that as a part of their research, which I think is an ethical violation. But then you have that conflict of was that a violation Reddit's terms of service, as well, and does Reddit have recourse to go after those researchers. However, that being said, also coming from a psychology background, sometimes you need to manipulate people or to have some kind of false premise in order to get a truthful response from people. So, I do understand the need for deception in some cases. But in research when you are cleared by the IRB, the Institutional Review Board for their ethical practices, deception is one of the highest bars because of the research and is considered very unethical to deceive people. That was one of the really interesting ones, and I think we can expect to see more things like that about as AI hits the mainstream, academia and research. What does that mean for that field? And then, finally the last one, which is a very recent one, was an victim impact statement. It was a road rage murder case, where the family actually fed kind of a script in, and had AI recreate the victim in order to give his impact statement. I believe this was in Arizona. This is the first time that a victim has testified in court. And that one just kind of gives me the heebie-jeebies. I think people have all sorts of feelings about it. I think now that I have to put in my will, 'Do not recreate me in AI,' as an example. These are things where it is now becoming not just a fringe thing or not just a researcher. These are real actual applications that we are

seeing. So, think about when this technology gets into the wrong hands, what it can also do. Now we have to be worried about creating visuals or audio of things he did not say. That's just an example that these technologies have these uses out there. But it is in the usage of them that can make something good or bad, depending on where you stand with your ethics or your morals.

So, to set some context or background for this framework, this framework enables organizations to do four things: to govern their risk, to map their risk, measure it, and manage it. This is what I've done in my own career to help a wide audience level set on what are we doing with AI, to help establish those policies, processes, good practices, and procedures. We are saying that AI risk is real, it's here, we don't have to be scared; here is what we can do about it. And also build that organizational culture that's built to prioritize identification and management of risk. Thinking about risk is not something that we need to hide and be scared of, it's something that we need to talk about to proactively manage it, and have a vocabulary to do so. By using something like a risk framework, you can have a common language with common definitions for both the really technical folks who are building these systems, and the policy groups who are wanting to help govern these systems, as well, who may not be as technical. It also is very outcome-focused. It is what gives you clear actions, a dialog and understanding a clear set of activities to manage that risk. Again, it is wanting to promote those positive impacts and minimize the risk, but it also has flexibility to what you give to your particular landscape, your particular priorities, and your risk tolerance. For example, a bank may have more strict regulatory and legal, and security requirements, as well, versus something like a cotton candy retailer, or something like that. Or an ecommerce company may have different requirements than a behavior or mental health company. So, thinking about the different environments that industries sit in, you want to have something that can be terrible and is not too restrictive on any of those. It also cannot lie with, if your organization has existing values or principles. A lot of companies have privacy principles that we just saw with the data strategy, moving towards a principlebased approach, AI is kind of taking the same path. It's a lot easier to have that big principle that you want to follow rather than really a prescriptive line item. And also, AI is changing so rapidly, so quickly. Almost every day, something is coming out. So, you want to have something that is going to be flexible over the longer term in areas where there is not legislation or regulation, excluding the EU, because the EU is ahead of where we are in a lot of places, but as I work for a multi-national company, we do have to comply with the EU AI Act, but we also want to be ready to prepare for any future USbased legislation that might come out. And finally, this framework helps align with industry with best practices, so whether you are in a team like this in Seattle, a Fortune 500 company, or a research institution, you can all learn from each other if you are aligned on the same framework.

So, the NIST AI risk framework that I have been talking about, the govern, mathmeasured functions, but it really takes a wholistic perspective regarding risk, and it thinks about impacts and harms to people, to organizations, and to eco-systems. So, it's not really just thinking about one specific thing. A lot of security things are really specific or really technical. This is really trying to take a wholistic approach, thinking about the

whole AI lifecycle, thinking about the different ways to adapt to and address those risks. It is extensible and flexible and you can use it with other frameworks like ISO or OECD or all of these other different acronyms that I love throwing out there, and really pulls it into a broader enterprise risk management framework. So, it is law- and regulationagnostic, but it will help you fit in and make that first step into complying with the new laws. Thinking about the government function in the middle there, and not thinking about the culture but thinking about the policies, the procedures, it is kind of about the people behind the technology, but not function thinking about the context in which you are using AI and how you relate this to those contexts and identifying those and naming what you are thinking about. Measurement is about once you've identified those risks, how we are assessing those, how we are analyzing those. Are we seeing trends of the risk? Is one particular business unit creating a lot of risk that is not being addressed? Are we seeing particular trends in seasonality, or responses to customers? And are we tracking those risks? Are we saying that this is a conception, this is where we introduce some mitigation. Are we seeing the residual risk go down or are we just letting them fly all over the place? And then, finally, are we prioritizing which ones we are going to go after and say, we have a risk tolerance; this one that is going to bring down our company is a higher risk and we get to go after that, versus this one that's going to affect a few users or edge cases. So, thinking about what is your projected impact and what is your likelihood of that risk happening, and using some kind of framework to map those out.

Thinking about risk as they emerge and how they are different, this risk management framework offers a path to minimize those negative impacts and maximizing those opportunities. So, organizations will have to decide for themselves what their tolerance is, what their priority is for these risks, and how do we want to organize and integrate these into existing responses. And they do a lovely job, again, of thinking about the potential harms not just to the organization, but in those three categories of people and whether or not the individual's civil liberties, physical harm in some cases, where there is thinking about things like autonomous vehicles, or economic opportunity nd thinking about credit scores, those kinds of things for individual harm. Is the group or community harmed? So, thinking about bias in this bucket. Are we discriminating against a particular population or a subgroup of people? Are we saying that teachers are just not going to get to use this algorithm while taking on a particular class of folks? Or societal, thinking about democratic participation or educational access. Again, that kind of goes back into that DOGE bucket. If we have concerns about the trustworthiness of the systems that our government is using, would we be less likely to participate in some of those processes? Thinking about harms for organizations, harm to their business operations, harm from breaches or monetary loss, or harm to reputation. These are some of the ways that we think about privacy or security risk, and so you will see those come up again in AI, as well. And then finally, harm to the eco-system, which I was kind of surprised by, but really happy that NIST included this. A lot of people talk about the enviro impact that AI would use. I joke with my friends that every time I use Chat GPT, I'm draining a town's water supply. And it's kind of a joke, but it is real that these systems and data centers use an incredible amount of electricity and water, so we need to think about the elements and the resources, the rare earth minerals that are used,

and the pursuit of all of the raw materials that are used to build these systems, the global supply chain, the financial system. I think we've seen ups and downs with the trade war situation tariffs, but also thinking about how important chips have been in this entire conversation, and global discourse about what countries can receive our highest end chips. And I think about something like DeepSeek, which was reportedly developed using less powerful GPUs because we placed restrictions on Chine, we start to add that geo-political component to this. So it gets really broad really quickly. So, thinking about those eco-system harms if I deploy this AI, now we're getting into that nationalistic area, and that's a potential harm there; harm to natural resources, the environment, and planet. So again, having a comprehensive framework gives you the language, gives you the tools to think about all of these things. It's really something that I haven't seen working in compliance, but I'm really happy because AI is so unique that we have a unique framework out there to address some of these issues.

One of the last things I will talk about is characteristics of trustworthy AI. And because trustworthiness is a social concept -- you think about a person or a friend and you trust them or you don't. It's a spectrum, but it is only as strong as its biggest characteristics. On the screen here in the graph, you can see a couple of different characteristics, like safety, secured, resilient, explainable, interpretable, privacy-enhanced, fair, with harmful bias managed, and then over that, accountability and transparency. And then underpinning all of that is deliverability and reliability. So, you want to look for approaches that enhance the trustworthiness, which is all of these characteristics as a whole, and reducing those negative risks. These are all different and distinct, but they influence each other. So, for example, you could have a highly secure system, but it is very unfair in its outcomes. You could have an accurate system but it is opaque and uninterpretable and you can't really figure out how it got there. Or you could have an inaccurate system that really doesn't work, but it is very secure and privacy-enhanced. Thinking about these trade-offs and the different levers that you have to pull in order to get something that averages out to something that is desirable. You have to acknowledge that there will be trade-offs, and that again goes back to the priorities of that organization and what they are willing to go for. There are a couple of things also that you will hear in responsible AI or AI governance. I want to provide some more context on the concepts, because some people will use them interchangeably, and there are slight differences. Thinking about validity and reliability, validation is confirmation from usually objective evidence that your requirements for intended use have been fulfilled. So, is it doing what we said it would do? Reliability is thinking about can it perform as required without failure for a given time and under given conditions. So, is it going to work again and again as we need it to? Accuracy and robustness, they kind of contribute to the overall validity and trustworthiness of the system, but they can be in tension with each other. Accuracy is defined as the closeness of observations, computations or estimates to the true values or values accepted as being true. Thinking about how tall is the Empire State Building, that's where we can get to an accurate measurement. Robustness, or generalized ability is referred to as the ability for a system to maintain its level of performance under a variety of circumstances. So, if I am thinking about a voice assistance, can it maintain its ability to keep calm under pressure. If I am interrupting it and constantly asking new questions, can it keep up with

me? Safety is thinking about preserving human life, house, property and preventing those things from damage. So, again, not all will have that concern, but for some, like autonomous vehicles, that's going to be probably one of our highest concerns. Security and resiliency, again, are we protecting the confidentiality, the integrity, availability, the protection mechanisms preventing outside access, preventing or changing outputs or outcomes. And resiliency is thinking about can it withstand those impacts or unexpected changes. Explainability and interpretability are the two that get confused the most. Explainability is the representation of the mechanisms underlying the AI system's operations and really thinking about using this or that algorithm, this and that model, to combine and make the system that will give this particular outcome. Where interpretability means understanding the output of an AI system in the context of their functions. Thinking about why did the system come up with a particular result. Because traditional software, which is deterministic, like when you are writing code, it is ones and zeros. And AI was called a non-deterministic system, which means if you ask it a question 50 times, you may get 50 different answers; versus something that is programmed to have the same answer every time. So, understanding why it came up with a result and how it came up with the result, that's thinking about why we are getting the answers that we're thinking about. privacy-enhanced is pretty straight-forward. These are the norms and practices that safeguard human autonomy, identity and dignity. It can have a pretty big concept of thinking about, in the context of AI, for example, can I get someone's personal information out of the system without putting in personal information. In fairness, just thinking about where equity, equality, and those other concepts come into play by addressing some of those issues, such as bias and discrimination. But these can be really complex and difficult to define, because this changes, based on the cultural context. This change is based on geographical context; this change is about organizational context. So, it can be tricky, but that's why it is important to make sure that all of these things are thought about before the deployment of the system. You will see a lot of these phrases when you look at a company's AI principles, or here is our approach to responsible AI. These are going to be things that you are going to see, and I think NIST has tried to distill those into core values, where you may see some algorithms kind of relate to these, but these are what you need to be looking at across the lifecycle.

This is one of the last we will cover.. This is a very big diagram, but think about what are the different actors and key dimensions across the AI lifecycle. So, thinking about the context in which it is being used, the data and input that's being used in a model, the model itself; trying to solve for the output; the application context again, so how we're using it in operations and monitoring; and then where the people that are operating or using are impacted by these. This leaves out what testing, evaluation, verification, validation, or TEVE steps you should do at any particular stage to make sure that you are accounting for some of these risks, as well as who are the people who are involved in these particular things. So, thinking about whether this is going to be an end user that is impacted, or a domain expert, a data scientist, model engineers, other integrators, those types of folks. They all have a role to play. So, really thinking about how this is multi-faceted, multi-disciplinary, and it takes a wide variety of people to come together to run these systems in an appropriate manner. This is kind of laying out the ways that

you can think about, having a development lifecycle that matches up to this. Or it can be standardized. So, when I say to someone, my company, we need to talk about the deployment and usage phase, they can say, got it; we're thinking about this particular context. Here are the stats that I would expect from you or someone in that stack.

So, I'm going to stop there. I ran through the high level view, but I encourage you to get all of the publicly available info. The beautiful thing about this is that it is taxpayer funded, and we are seeing a lot of different organizations adopt this. In my research, Google DeepMind I think has all merged into Google product teams, but they are one of the research arms of the GPT, the Generative Pre-Trained Transformers model, they an assessment using the NIST AI framework, and came up with an analysis that organizations can do to map themselves up to what's going on, even at the highest level. So, thinking about Google as an industry leader, they are also adopting these, and we're seeing other companies like Workday, Microsoft, others. A l9t of the key players are taking a look at the framework. So, I highly encourage everyone to check it out. I think, again, as someone who has used this professionally, and seeing the assault that's happening on our institutions and federal agencies, this is one that I think also needs support. It's not a popular one, but they do a lot of good work behind the scenes that are keeping a lot of organizations up and running. I'll pause there to see if anyone has any questions. I'm just pulling up the chat.

From chat: Robert Kruse 5/13/2025 7:29 PM • In this briefing by Omari, he did not say "ontology" once. That's a real concern. 5/13/2025 7:30 PM • EU has spent \$B of Euros over past 25 years across all major universities, is well taught and applied in thousands of projects. Our local Universities are simply behind.

Omari Stringer: Robert Kruse, I see you're concerned that I'm not mentioning ontology. I got my master's degree from the University of Washington a couple of years ago. Certainly, it was not a topic then, and I don't know that it is a topic now. I think whether or not our universities are behind, all of us are playing catch-up with something that's coming fast, and we can never really keep up with the pace of technology, but we have to look at the tools that we have available to us. I will definitely take a look at ontology and figure out how that fits into the AI governance framework, but I do think that I agree with you that the US in particular has a lot in terms of regulation and education about how we can responsibly use these psychologies for sure. Anyone else with comments, questions?

All right. In that case, I will keep us rolling. I know the next item on the agenda was our CTAB workplan items. I did have a chance to sync with Phillip Meng today. He would prefer if we table that to the next meeting so we can have an in-depth conversation. We might have a little bit more planned time for the workplan items at the next CTAB meeting. So, I will roll into committee updates. I think I will also take this moment before I call on some of the CTAB members, I do want to give a shout-out to Coleman Entringer. I believe that this will be his last CTAB meeting, so I wanted to say thank you, Coleman, for your participation and really leaning in over the last few years, especially with the Digital Equity Committee. We hope that you won't be a stranger, and hope to

see your familiar face at some point. I wanted to acknowledge your contributions to CTAB, and thank you very much for your membership.

COMMITTEE UPDATES

DIGITAL EQUITY COMMITTEE

Coleman Entringer: Thanks so much, Omari. I really appreciate that. And just for the general update for the group as well, besides my personal absence, I will be moving away from Seattle. So, please reach out to me if you would like to get involved in the Digital Equity Committee. I would love to definitely be a resource, going into the future, but would definitely like to get some more people connected so that we can continue to do the great work that we've been doing.

For updates from last session, the committee has completed our hand-out documents for the affordable and accessible telephone and internet access forum. I think that, given that we have a few absences, we can review that in full at the next meeting. But I've sent it to Phillip Meng and the board. Hopefully, you can send that to the rest of the members for review, and then send out to. through our communications channels in the future. That is the update for the DEI Committee.

Omari Stringer: Awesome. Thank you. I believe DEI worked with the Outreach Committee. Are there any updates from them? I'll take that as a no. Hailey, on the Digital Wellbeing Committee, do you have any updates for the larger group?

DIGITAL WELLBEING COMMITTEE

Hailey Dickson: I don't really have any updates for this session. We are working on scheduling a community listening session in the next few weeks, and then settle on some topic areas that we want to align on. I will reach out over email if anyone is interested in being involved. Please let me know and I will make sure you get the invitation.

Omari Stringer: Thanks. I think for the Privacy Committee, we have not -- I have an open action to get those scheduled. I have had some personal leave of absence over the last few months, so getting back into the swing of things, so there will be more coming with that. I do want to acknowledge and trying to be honest here, I don't know if it is open or closed, but there is a comment period for surveillance technologies upcoming or just recently passed.

Vinh Tang: Yes. It was an update for SPD tracking devices. The public comment period has closed. It was through the last week of April. The original SPD tracking devices and GPS from 2023, the Seattle Police Department is seeking a Washington State commerce grant totaling \$250,000 to pilot over a two-year period a technology called Star Chase. Captain (unintelligible) presented to the committee about the technology.

Omari Stringer: Thank you for that update. We definitely want to follow that. In the meantime, I don't know when the Council vote would be. The comment period ended on the 28th. I think the Public Safety Committee met this morning. It's something folks want to keep an eye on, keep tabs on that. So that brings us to our last segment of the agenda, which is public comment and any announcements. I will turn it over to folks on the call for any final public comment. I don't have any particular announcements at this time, so this is an open call for any public comments.

PUBLIC COMMENT

From chat: Coleman Entringer 5/13/2025 7:37 PM • My email is <u>coleentringer@gmail.com</u> if anyone would like to get involved in the DEI committee!

Omari Stringer: All right. That being said, last call for anyone with a public comment. I'm looking in the chat to see if there is anything there. It doesn't look so. With that, I wanted to thank you all for attending. Thank you again to Mark Schmidt and Steve Barham for their presentation, and to the folks who had those great questions and comments in the chat. We appreciate that, as well. This will conclude our May CTAB meeting. Thank you.

ADJOURNMENT