EPISODE 1074

[INTRODUCTION]

[00:00:00] JM: Geospatial analytics tools are used to render visualizations for a vast array of applications. Data sources such as satellites and cellular data can gather locations and location data, and that data can be superimposed over a map. A map-based visualization can allow the end user to make decisions based on what they see.

ArcGIS is one of the most widely used geospatial analytics platforms. It's created by Esri, the Environmental Systems Research Institute, which was started in 1969. Today, Esri products have 40% of the global market share of geospatial analytics software.

Max Payson is a solutions engineer at Esri and he joins the show to talk about applications of ArcGIS and the landscape of geospatial information systems more broadly.

If you have an idea for an episode, whether it's about a company or project that you're working on, you can go to softwaredaily.com and submit a topic. We're always looking for good ideas. We're always looking for new companies to cover and new open source projects that are exciting and have some traction. You can also support the show by becoming a subscriber and you can get ad-free episodes, all that at softwaredaily.com.

[SPONSOR MESSAGE]

[00:01:23] JM: If you listen to this show, you are probably a software engineer or a data scientist. If you want to develop skills to build machine learning models, check out Springboard. Springboard is an online education program that gives you hands-on experience with creating and deploying machine learning models into production, and every student who goes through Springboard is paired with a mentor, a machine learning expert who gives that student one-on-one mentorship support over video.

The Springboard program offers a job guarantee in its career tracks, meaning that you do not have to pay until you secure a job in machine learning. If you're curious about transitioning into

machine learning, go to softwareengineeringdaily.com/springboard. Listeners can get $500 in scholarship if they use the code AI Springboard. This scholarship is for 20 students who enroll by going to softwareengineeringdaily.com/springboard and enter the code AI springboard. It takes about 10 minutes to apply. It's free and it's awarded on a first-come first-served basis. If you're interested in transitioning into machine learning, go to softwareengineeringdaily.com/springboard.

Anyone who is interested and likes the idea of building and deploying machine learning models, deep learning models, you might like Springboard. Go to softwareengineeringdaily.com/springboard, and thank you to Springboard for being a sponsor.

[INTERVIEW]

[00:03:00] JM: Max, welcome to the show.

[00:03:02] MP: Hey Jeff, excited to be here. Thanks for the opportunity.

[00:03:06] JM: You work at Esri, which makes ArcGIS. ArcGIS, being a geographic information system. Explain what a geographic information system is.

[00:03:17] MP: It's a system for working with all aspects of geographic data, from upstream, actually collecting that data, to storing it, managing it, and downstream analyzing it and communicating the outputs of that analysis. It's a comprehensive system built up of many components that kind of get orchestrated together into a GIS.

I think another way that I would describe it too is almost as a framework. You can imagine that geographic data is very heterogeneous. You have satellite imagery. You have row networks you want to route against, and GIS is really good at abstracting the differences between those different datasets and bringing it together so you can visually match the data together, or computationally match the data together, or pull it into work flow seamlessly.

[00:04:03] JM: Is it a database? Does it hold my actual data or is it just creating a materialized view like a business intelligence tool for geospatial information?

[00:04:16] MP: I think it's all of those. We really do build technology across the entire stack. The way I would probably describe it is that we build the system to fit three patterns. One is a system of record. Another is a system of insight, and third one is a system of engagement, which is a little bit of IT parlance, but I think it gives some good context to how we architect the system.

You can imagine on the system of record side, we have to build databases to store geographic data. On the insight side, we have to build out tools to interact with and analyze that data and then engagement applications that users can interact with to understand the data.

[00:04:57] JM: The data in our GIS, is it like I bring my own data and I use it to stand up an instance of ArcGIS, or does ArcGIS come with the data in it already?

[00:05:12] MP: Originally, back when ArcGIS was first developed, it was entirely bring your own data to the point that you would open up the application and you would just see a blank screen, which is kind of a funny concept to think about. But over time, users have asked for more datasets to be provided by default. So there are datasets that Esri curates or sources from providers that we make available into this ecosystem for users to manage and work with their geospatial data.

[00:05:41] JM: As far as the query process, so if I think of an ArcGIS instance loaded and let's say I'm managing a port, like I'm somebody who just is managing the supply chain of a part and the containerships that are coming in. What's a way that I would use a GIS?

[00:06:03] MP: Yeah. You can imagine that there are a lot of workflows that are around it that are grounded in location for the port. There is understanding where the ships are in a given moment. So you have the sensors on the ships, like AIS signals that emit where the ships location is, and those signals can get streamed into the GIS so that people operating the port are able to see a snapshot of where the ships are. That's one sort of operational workflow.

Another workflow is you might want to go out and dredge the port. So you go underneath the water or into the water and start cutting out the grounds that ships can fit next to the docks. In that case, you need to understand the bathymetric contours of the surface underneath the water

and you need to understand where are the priorities for dredging, and you would also use the GIS to collect the data about the dredging and then figure out the appropriate spots to perform that dredging. So workflow-base is probably how I'd describe it.

[00:07:06] JM: Great. Okay. Well, let's go through some other workflows that might illustrate how people use ArcGIS. Esri works with Mobileye, which is an autonomous vehicle technology company. Can you tell me about the use case there?

[00:07:23] MP: Yeah. One of the primary users of ArcGIS are state and local governments. So state and local governments use ArcGIS as a way to maintain their authoritative data about a municipality. So you can think about parcel boundaries where buildings are built as well as stop sign locations and components of the road, right?

Where Mobileye fits is they're, like you said, an autonomous vehicle company. They are collecting all these datasets at the – Or they're collecting injury at the edge to build out driver assistance and autonomous vehicle workflows. But there's also this use case of taking that data and making into local governments or other stakeholders to enrich their workflows as well.

If you're a local government, for example, you need to maintain your stop sign inventory, and sometimes that inventory gets out of date or it can be expensive to collect. If you have these cars, they're driving through with Mobileye devices, they're able to detect stop signs. Then we've partnered with them to bring the stop sign locations into the GIS so that local governments have access to fresher and fresher data about the locations of their assets. That's one specific workflow, but the high-level intent here is to work with Mobileye to figure out what data is being emitted from their cars and being able to bring that into ArcGIS whether for asset management or in some cases even analytics about broader patterns.

[00:08:52] JM: Your work is as a solutions engineer at Esri. Can you give some examples of companies that you work with?

[00:09:00] MP: Yeah. Mobileye is a good first example of that, but my role as a solutions engineer is to help emerging tech companies build out products with ArcGIS. That spans from helping them figure out what components of this GIS system could benefit their stack and what

they’re building into their product as well as how they can integrate their technology into the ArcGIS ecosystem and the user community. There are a lot of emerging technology companies where that fits in and locations becoming I would say increasingly pervasive. There's Mobileye.

Another good example is this company Spatial.ai that is taking social media feeds and mining demographic data from the social media feeds. You can see where do people like to drink coffee, for example, and we pull that data into ArcGIS and make it available so that users like Starbucks or someone who wants to determine a location can figure out where there are areas that have demographic profiles of the people they’re trying to sell to. So Spatial.ai is another example.

We have companies like FernLeaf that are building out resiliency toolkits in the event of a flood. How do local governments and engineering firms effectively respond to those? A whole plethora of maybe 450 more companies that I engage with.

[00:10:23] JM: Maps are fundamental to how a user is going to use a geospatial information system. Can you tell me about the different kinds of mapping systems that exist across the ArcGIS user base?

[00:10:40] MP: I'll probably start on the developer side. We make APIs and SDKs that developers use to render maps into their applications. We have a JavaScript API that renders maps, and WebGL as well as native SDKs. The first way to kind of create maps within the system is to use those lower-level developer tools to create maps in your own applications, and you can render the data, you can style the data and you can create map-based interactions like filters and queries.

Now, within Esri, we use those same tools to also build out additional map authoring environments. One is what we call ArcGIS Online. It's a managed SaaS mapping and location platform, where as a user or a developer, you can upload your data and then use mapmaking application to style that data and create maps and visualizations against it.

Also, the third aspect I think is we also have a heavier desktop client, which is called ArcGIS Pro, and that's used for advanced cartography. If you're making this beautiful topographic map

of a national park, for example, you might use ArcGIS Pro to have these really advanced and fine-grained controls over the map so you can print it out in a very compelling and almost artistic way.

[00:12:00] JM: How do the maps get customized? If I'm a user, I've got a very specific vision for how I want the map to render, what kinds of use case I – the kind of use case I have? How am I customizing that map?

[00:12:17] MP: It's another good question. A lot of it is often driven by the underlying data. You can imagine there is a class of geographic data called Vector Data or feature data, which is abstracted points, lines and polygons, each with associate attributes and you use those attributes to control how the data is visualized.

If you're looking at COVID-19 response efforts, for example, and you want to visualize based on the counts of occurrences of COVID-19 or the coronavirus, then you'd be able to use that attribute of the location to determine how it gets styled in the map. We have a couple different tools to make that process easier. We have what's called smart mapping, which interrogates the statistical nature of the data and makes smart default representations of that in the map.

[00:13:10] JM: Can you go a little bit deeper on the smart mapping? I think about a bunch of data that I'm going to be putting into a map. That map can display interpolated statistics of the data that I'm putting in. I'm sure these interpolations can improve with machine learning. Tell me about the smart mapping.

[00:13:32] MP: We'll break it up in two ways. One is on the visualization side. How can we quickly make visualizations of our data in a way that's compelling and informative? That's where smart mapping fits in. Can we quickly create a representation of the data that highlights outliers in a given numerical count, for example? That's where smart mapping comes in.

We also have this whole division, I would say, of spatial statistics that is building out algorithms through Python APIs and notebooks that allow you to interrogate the data, run it again some process and produce an output. There is an ecosystem, maybe a thousand plus tools that are built for these analytical workflows. Does that answer your question?

[00:14:22] JM: Yeah, I think so. Can you give me anymore examples of how machine learning goes into ArcGIS?

[00:14:28] MP: Yeah. I think maybe one way to distinguish between the visualization and the analytics side is a really common way to represent density and a map is with a heatmap. So you open up a map and you see this sort of blurred representation of data, and that's a good visual representation that shows basically these aggregated pixel values fading away from where the data is located.

On the analytical side, I'm just trying to get into the more statistical algorithms and machine learning algorithms, we can actually quantify where the densities of data are. So we can look at the spatial distribution of that data. Identify clusters where things are dense together and categorize those as hotspots that can then be visualized.

We can also do that over time. We're not only categorizing hotspots in a given snapshot, but also looking at where there are emerging hotspots or declining hotspots. You can imagine that has some use cases in fighting crime, for example. You want to understand if there is an increase activity in crime in a given municipality and be able to allocate law enforcement to respond to that more effectively moving forward.

[SPONSOR MESSAGE]

[00:15:42] JM: Over the last few months, I've started hearing about Retool. Every business needs internal tools, but if we're being honest, I don't know of many engineers who really enjoy building internal tools. It can be hard to get engineering resources to build back-office applications and it's definitely hard to get engineers excited about maintaining those back-office applications. Companies like a Doordash, and Brex, and Amazon use Retool to build custom internal tools faster.

The idea is that internal tools mostly look the same. They're made out of tables, and dropdowns, and buttons, and text inputs. Retool gives you a drag-and-drop interface so engineers can build these internal UIs in hours, not days, and they can spend more time building features that

customers will see. Retool connects to any database and API. For example, if you are pulling data from Postgres, you just write a SQL query. You drag a table on to the canvas.

If you want to try out Retool, you can go to retool.com/sedaily. That's R-E-T-O-O-L.com/sedaily, and you can even host Retool on-premise if you want to keep it ultra-secure. I've heard a lot of good things about Retool from engineers who I respect. So check it out at retool.com/sedaily.

[INTERVIEW CONTINUED]

[00:17:18] JM: If I've got a map that I'm using really intensely and I'm putting tons and tons of data into it, what kinds of scalability challenges might be encountered in the rendering and the generation of that map?

[00:17:36] MP: I think one good way to talk about that topic is to look at the history of how maps have been displayed on the web. In the early days, maybe in the early 2000's, maybe in the late 1990s, the approach was that an application requests a map from a server. The server generates an image of that map and returns the results to the client and they just display that image. You can imagine that that starts to have some scalability issues though where every user is looking at unique sets of the map, and so every time the server has to go into the database, get the relevant data and cook that into an image.

What Google Maps did in the mid-2000s was they had created this notion of tiles, and tiles are pre-computed images that are consistently applied so that when someone requests data, they just get the appropriate set of tiles, and this allows them to be cached and served out without going back into the database. So it's much more dynamic.

The considerations around that, though, are you now have to pre-compute these tiles, and so you assume that the data doesn't change that frequently, or users always want to see the same representation of the data. What Esri is doing is building out this notion of feature services, which are a way to both maintain a reference to that authoritative data within the database, but also scale it out through smart caching and some other optimizations as users request the data.

Again, a tangible example of that is in response to COVID-19, John Hopkins University. I don't know if you've seen their dashboards or not, but they're using Esri's software behind-the-scenes to create those dashboards. The data powering those dashboards gets billions of requests in a given day, and that was in March. So it's probably grown since then. There's this unique engineering challenge of having to respond effectively to billions of requests in a day as well as give always the most up-to-date data for the different consuming clients who want to interact with it.

[00:19:45] JM: Can you give me more of an overview of the other kinds of tools that exists across the Esri platform in addition to mapping? Maybe take me through – You could take me through a customer use case and describe how the tools for that customer fit in.

[00:20:02] MP: Yeah. I think an interesting use case would be we worked with a company, SafeGraph. They create derived datasets from human movement patterns. So you can see aggregated and anonymized who goes to a given location or where they're coming from. This is used in site selection, for example, where a retailer wants to know what's the profile of people that are visiting a given store.

The way that we've started to build a solution around this is we're able to pull in SafeGraph's data as a service and basically join where these people are coming from to demographic data. This is done in a desktop tool, for example, or with a Python script that takes the SafeGraph data, looks at a given point of interest, looks at the corresponding geometries, where those people are coming from. Joins that data using a demographic dataset and a geo-enrichment service to the demographic data and is able to create a visualization from it, from the output. Did that make sense to you? I feel like I could've explained that a bit more coherently.

[00:21:14] JM: No. That that makes sense. So what's the implementation process for that? When you're working with them, when you're working with SafeGraph, how are you figuring out what exactly to build for them?

[00:21:26] MP: Yeah. The implementation process is looking at what tools do the users have available to them to make sense of that data? Then how do we format the data in a way that it can be leveraged within those tools? There are two tools that are really common in this use

case that Esri provides. One is called ArcGIS Insights, which is an interactive charting, almost BI-like tool. Then there's another tool called Business Analyst, which is used for generating reports and summaries about demographics of people visiting a given location.

The question is how do we format and preprocess that data in a way that business analysts and insights can create reports and inform some sort of investment or site selection decision? That preprocessing step involves, in this case, with SafeGraph, using a Python API to read in the data and output into a format that ArcGIS can consume, as well as looking at the locations of the dataset where consumers are coming from and joining that to demographic data that Esri provides so that's enriched and more actionable within those insights and business analyst tools.

[00:22:45] JM: Okay. Can you tell me more generally, do people who are using ArcGIS, do they deploy it on-prem or is it mostly for cloud-based applications? What's the typical user?

[00:23:00] MP: We have all of the deployments, which is kind of interesting. It's increasingly moving towards the clown. We still have a lot of users who deploy it on-premise behind the firewall, but definitely in the trends and the new normal is the cloud for us. We have two options. One is we have licensed server software that you can deploy into your cloud environment. We also have our own cloud- managed hosted software as a service that you can just interact with through the applications without having to worry about any of the IT governance or cloud deployment considerations.

[00:23:41] JM: This brings to mind the fact that the company has been around for like 50 years. Do you have much insight into what the stack looks like? What the software stack looks like and what, particularly, the legacy pieces look like?

[00:23:58] MP: Yeah. I mean, that's a fair question. It's a 50-year-old company. So do we have components written in COBOL? The answer is no. Our core engines are written in C and Java, and then I believe our server product is also written in Java as well. Then we have different applications that write for their native platforms. So for example, our JavaScript apps use WebGL to render the data and they use, in some cases, even WebAssembly to use compiled C core engine in the browser, which is really, really cool, I think.

But the history of it is pretty interesting, and that we have been around for 50 years, and almost every 10 years, at least historically, we fundamentally reengineer the platform. When it first got started, users were interacting with the GIS through a command line. Then it went to desktop, than a client/server architecture, and now we're in this web-based architecture where people can seamlessly interact between different services emitted through the platform and we're actually in one of those reengineering efforts right now where we're retaking the core server products and reengineering it into microservices that can be deployed into containers and orchestrated through Kubernetes so it can be more seamlessly deployed in this cloud native fashion.

[00:25:19] JM: Have you taken part in that re-architecture or is that a different part of the company?

[00:25:23] MP: I wish. I wish I could say I did, but it's a different part of the company that is building out that server and sort of core product.

[00:25:32] JM: What are the other elements of the platform that are being developed? What are the newer features that are coming out?

[00:25:40] MP: I think one of the really interesting ones is on scalable analytics, and this is both real-time analytics as well as big batch analytics to drive insight from data. We have a new offering called ArcGIS for IoT. That's a fully managed IoT offering. It's built on Kafka and Spark and some other technology components that we deploy in the cloud and then can consume billions of sensor input feeds a second and rationalize over those with spatial semantics. You can do things like geo-fencing, or track how long a given feed has stayed in a location, and then output those into this mapping system to communicate current location of assets or clusters of activity.

[00:26:30] JM: Do you have an example of when that will be useful?

[00:26:33] MP: Yeah. A use case that I believe came up recently is for Department of Transportation planning. Department of Transportation, for example, has a bunch of new feeds

that they want to connect to to have a picture of what the transportation system looks like. One of those feeds might be Waze alerts, for examples. We have a connector into Waze that says there is a reported traffic incident here and we use this IoT platform as a way to quickly communicate that into the Department of Transportation. Also, looking at the current location of buses, for example.

As a bus is driving along, it emits its location and the Department of Transportation not only has a picture of where those buses are, but if the driver is even deviating from their normal route. As that location is coming in in real-time, we can see is that bus on the route supposed to be on or is it actually someplace else, and communicate that through a software automatically at scale to the Department of Transportation.

[00:27:42] JM: The data flow in that situation, can you take me through that? You've got like a real-world device and its communicating data over cellular and then it's getting buffered on ArcGIS's servers and then getting processed? Just take me through the data flow there.

[00:27:59] MP: Exactly. Exactly. You have some sensor that is connected to the Internet into a cellular network and it emits data, and then ArcGIS can connect to that stream of data or the sensor can just post it directly to ArcGIS and then ArcGis, exactly like you said, it buffers all those input feeds and processes them in real-time and then outputs them into a database within ArcGIS where users can then later visualize the historical records of that data or perform computations against it as well.

[00:28:36] JM: Cool. The internal systems at ArcGIS, do you know if they use like cloud provider technology, or is it all stuff that's written in-house?

[00:28:49] MP: That's a good question. I'd say it's a combination. In our ArcGIS online, that software is a service that we provide, it's deployed in a multi-cloud environment. We use both AWS and Azure under the hood, but we also write a lot of our – Like I mentioned, we also use some open source technologies like Apache, Kafka, and Apache Spark for some of our – That we then build on top of for some of our software as well. But then we also write a lot of it in-house. We have our own software, for example, for managing imagery data at scale and being able to reference imagery data, let's say in an S3 bucket, and dynamically process that and

serve it out to a connected application so that the users can work with. Even if that source imagery data is petabyte to scale in the S3 bucket, we know how to index it and then operate against it in a way that an application only sees the subset and visualization of that petabyte scale imagery that they need.

[00:29:52] JM: Let's say you make a solution for a company, like Mobileye or like SafeGraph. Does the user or the company that has purchased that application, did they have the ability to change the structure of the application? Can they kind of rewrite how the GIS application they've made works?

[00:30:15] MP: Definitely. Definitely. I would say, is from an end-user perspective, the schemas and ways to work with this geographic data are really abstract. To make that tangible, if you're working with Vector Data, which is those points, lines, and polygons, and associate attributes, that's really the only structure that ArcGIS enforces, and you can use that very flexibly to represent all different types of geographic data. So you can configure the system in a way to work with your dataset. Then from a developer's perspective, you're also able to use those APIs and SDKs that I talked about earlier to build your own workflows or extensions against that core system.

[00:31:02] JM: Can you take me through another example? Like let's say I am a city planner. If I'm a city planner, how would a geospatial information system be useful to me?

[00:31:14] MP: I think the first part on city planning is you have to have a picture of what the current state of the city is, right? This is getting back to that Mobileye example where cities need to know where their stop signs are and they need to know where their bike lanes are. Sort of the first use case in the city planning example is having the different departments collect and maintain that data, and a lot of that is stored in a geospatial database, and a lot of those geospatial databases are created and managed by ArcGIS, right?

Now, once you have the current state, then the question is how do you start to create scenarios or new plans. There are two aspects that. One is the tools to create the plans, making those seamless and easy. We build out tools, for example, to help you automatically create proposals based on various inputs like zoning regulations. So you don't have to sketch out, as a city

planner, your exact proposal by hand. You can say, “Given these constraints, show me what a hypothetical proposal could be.”

Another component on the planning side is the analytics. What are the effects of this proposal? Is it adhering to my metrics? Is it in achieving my objectives? I think a good analytical example that I learned about when I first joined the company is the City of Boston has a shadow bank where they’re only allowed to have so much shadow cast in the public spaces. When they’re evaluating new plans, they use ArcGIS to look at the proposed new building and whether or not that draws too much shadow from the shadow bank. That’s sort of this cool 3D analytical workflow. That is another example of using ArcGIS as an analytical component. The last part is just being able to communicate those plans. Making sure that a city planner is able to seamlessly through applications share that data and the plans to the city council or to the public for feedback and other stakeholders.

[00:33:16] JM: The shadow system you described, that a system for visually replicating how shadows would be cast from 3D buildings?

[00:33:27] MP: Yeah. Yeah, exactly. It's really cool. You can imagine there is a set of tools that we provide to first create those buildings. I’m happy talk about that more in depth if it’d be interesting. But once you have the representation of those buildings, then you can use ray-tracing algorithms and other computations to trace out where the sun is and how those intersects with the buildings and cast shadows on to the surrounding environment and use that quantitatively to look at things like a shadow bank.

[00:33:56] JM: I know there are some drone technology companies that are using ArcGIS. Can you give me a description for how drone companies could use it?

[00:34:06] MP: Yeah. One of the earlier questions, a lot of the GIS is built around collecting data. Drones are really great feed for data collection and creating new geographic data within the GIS. We have drone companies that we work with or partner with to produce the raw data, and then ArcGIS can process that raw data into consumable forms of geographic data.

An example of that would be a digital elevation model or a train model that gives you the profile of a given landscape from the raw imagery collected from a drone. We have tools that process that raw drone imagery, create this digital elevation model and then are able to consume that digital elevation model in different applications and services depending on the workflow. If you want to do just a simple elevation analysis, you can use that output elevation model to do that.

[00:35:11] JM: Go ahead.

[00:35:11] MP: Yeah. So then the other part that I was just going to add is in addition to this actual data collection workflow, there's also this notion of managing drones and actually creating the flights and making sure that those flights are to compliance, and so we have partners in the ecosystem, as well as end-users that are using ArcGIS as a system to manage and record flight corridors, for example, and then use that to plan drone operations that feed that collection.

[00:35:42] JM: Let's revisit what's actually going on here. If I am a drone company and I'm using a visual map representation, the mapping data that goes into that, is that produced by satellite data? How are you gathering the overall picture of the maps that have information overlaid on them?

[00:36:09] MP: Yeah. A lot of it is gathered from satellite data. Esri does provide some datasets by default, and one of them is what we call a satellite base map. You can open up a map and see this geographic context of imagery coming from satellites and other data feeds, and Esri does some of the work to stitch that together and make it globally available so that when a user opens up a map in any location, they can see the underlying satellite imagery.

That can be used to kind of direct, for example, if you identify the building, you can identify that building in the imagery and then say, “This is what I want to fly my drone,” and use that to create the plan around the building.

I would also add that we provide tools for working with the imagery as well. It's not just us creating this one satellite base map, but we also have users that are collecting their own imagery, for example, with drones, processing that, turning it into a new imagery layer and then using that again to select where they want to fly a drone or some other workflow.

Let me give a quick example of that. So in disaster response use case, when a natural disaster has hit, you want to know where are damage buildings and where are damaged roads, for example. The partner or the user may go out and fly a drone and the use ArcGIS to process that imagery, and now you have this nice satellite or drone-based imagery that you can use to flag damaged buildings and damaged roads.

[SPONSOR MESSAGE]

[00:37:57] JM: When I'm building a new product, G2i is the company that I call on to help me find a developer who can build the first version of my product. G2i is a hiring platform run by engineers that matches you with React, React Native, GraphQL and mobile engineers who you can trust. Whether you are a new company building your first product, like me, or an established company that wants additional engineering help, G2i has the talent that you need to accomplish your goals.

Go to softwareengineeringdaily.com/g2i to learn more about what G2i has to offer. We've also done several shows with the people who run G2i, Gabe Greenberg, and the rest of his team. These are engineers who know about the React ecosystem, about the mobile ecosystem, about GraphQL, React Native. They know their stuff and they run a great organization.

In my personal experience, G2i has linked me up with experienced engineers that can fit my budget, and the G2i staff are friendly and easy to work with. They know how product development works. They can help you find the perfect engineer for your stack, and you can go to softwareengineeringdaily.com/g2i to learn more about G2i.

Thank you to G2i for being a great supporter of Software Engineering Daily both as listeners and also as people who have contributed code that have helped me out in my projects. So if you want to get some additional help for your engineering projects, go to softwareengineeringdaily.com/g2i.

[INTERVIEW CONTINUED]

[00:39:46] JM: Can you tell me more about the backend for what goes into an individual maps? If it's loading on my web browser, what's going on in the backend?

[00:39:59] MP: The backends differ a little bit depending on the geographic data that you're working with. There's some – What a GIS does is kind of abstracts those different backends together into the system, but to double-click on the most common one, I would say Vector Data, right? What's happening on the backend?

At the at the base level, the data is stored in a relational database. So you have a column that represents the point, the line, or the polygon geometry and then you have the associate attributes, which are just additional columns in the database. Then on top of that, there is a server for allowing users to interact with the data in the database. This is going back to a little bit around the scaling considerations that we are talking about earlier. We've built what we call a feature service, which is really good at disseminating the data in that database at scale while still maintaining a reference to it so people have the most accurate and up-to-date information.

That server tier is doing things like dynamic caching. When you request for data within a given map, it reaches into the database if it needs to, creates tiled representations of that data and responds, caches them at the server. Optionally catches them at CDN tear, and then the client is then able to consume the outputs of that server response and the cached data. Did that answer resonate or do you have additional questions on that?

[00:41:30] JM: Can you go deeper into the caching infrastructure?

[00:41:33] MP: Sure. You can imagine in the map space that we can create repeatable queries in a way that we're able to catch them. When you're looking at a given map, you can subdivided it into tiles and then those tiles can be dynamically created from the database. You make a request to the server that says, "Give me the data for this screen, or for this map view that I'm looking at." The server checks and sees if it's already responded to requests for similar geographic areas. If it has, then it just responds with the similarly structured requests and the client then receives them back. If it hasn't, then it goes into the database, pull them out and then sends them. Does that answer help clarify a little bit or would you like me to go –

[00:42:27] JM: Yeah. Yeah, no. That's totally fine. I think it's worth talking about the landscape of other systems that can be used to build maps. You have Mapbox, you have Google Maps. I'm sure there're some others. Can you tell me about this landscape and how those different application platform slot into different use cases?

[00:42:50] MP: Yeah, there's deftly been a proliferation of, I would say, location technologies. The way I would probably frame it is that a lot of them are either different components within the location stack, or they are different verticalized relations. If you look at Mapbox and Google, these are tools for rendering map applications into an application. There is less development on the server side or on the database side for how to manage the data that gets rendered into those applications.

I think on the ecosystem and how I view that as interacting with Esri, is we build a lot of those components across the entire stack to make this comprehensive GIS, and then there are other providers out there that also engineer components in that stack.

[00:43:43] JM: Some of the use cases for maps are like for embedded maps, like in mobile applications. Can you tell me about how an embedded mobile application, embedded mapping mobile application, might work for ArcGIS, or do you have any examples?

[00:44:03] MP: Yeah. When you say embedded applications, you mean in a car or a worker going out into a device or the actual geometries loaded into a drone for some quick decisions? I guess maybe expand on that.

[00:44:16] JM: Oh! I didn't know that the drone use case. I was thinking just about embedded maps in a mobile application, like on my smartphone, but I didn't know. People put ArcGIS software on drones?

[00:44:28] MP: No. I probably should have referenced that is an example, but that is something that's come up a little bit, is we have – So we have these SDKs for building native applications, right? You can use these SDKs to build a mobile application that pulls data in from the ArcGIS servers into that app and optionally creates offline packages of them in the mobile app. You can

do things like go out and collect data from the field completely offline and then synchronize those edits back into the GIS system.

One of those native SDKs that we provide is also based on cue, and so that is often used in embedded workflows. So one of the things that we're having discussions about is how can we use some of these SDKs in interfaces that aren't necessarily within a mobile application on a phone. That's why it was – I think the idea of bringing a geometry that's stored within ArcGIS on to a drone to determine within the drone, "Is this a place where I can fly?" is an interesting concept, but not something that we've engineered a solution for specifically.

[00:45:46] JM: Okay. So what about like mobile applications? Do people use ArcGIS on mobile apps to have some field worker be enabled with something?

[00:45:58] MP: Yeah, that's a really common use case for us. A common industry for this would be utilities, for example, where you need to go out and inspect the status of utility pull or you need to respond to an outage, right? These are all location-based workflows that are in the field. Esri, one, provides the SDKs to build these types of applications, but it also builds out-of-the-box apps that can be used to orchestrate field-based workflows, and that has different components. One is looking at the workforce and helping that workforce navigate to the locations that they need to go to. Another is having an individual field inspector, for example, collect the data in an online or offline scenario and then report that back into the authoritative database.

We have field applications as well as SDKs that are built to handle those database replications and syncs as the data comes online and offline. Then the last part is also being able to look at the feeds of data coming in from the field in either a web client or a desktop client to see what data these field workers are collecting.

[00:47:15] JM: Interesting. The field worker applications, that could be like if I'm somebody that maintains trails or if I'm somebody that goes out and fixes telephone poles, you could have domain-specific applications for those kinds of things.

[00:47:30] MP: Exactly. Exactly. So we try and build a couple standardized applications that have configurable data models depending on the use case. Stop sign inventory versus utility pole inspection, right? There are different data models that are able to be consumed in these out-of-the-box applications. But for developers and partners, there are also very specific workflows and reporting or industry solutions that would benefit from not just a configurable data model, but a fully out-of-the-box experience. We have partners that build these out-of-the-box experiences for reporting and data collection as well.

[00:48:07] JM: What are you working on at the company today?

[00:48:10] MP: One topic that I think has been surprisingly trending is this notion of risk-based routing. There are a bunch of companies that have expressed interest in this recently. This is the idea of taking novel datasets or new datasets that are coming off of connected cars like Mobileye, for example, and using those datasets to, one, create a model that predicts risk for individual road segments and then making that risk actionable by allowing fleet operators or insurance companies to route around the risk, or taken into account when generating the risk, or understanding whether or not their drivers are following risky roads. A really one area where I've spent a surprising amount of time in the past couple months is helping these companies take their risk predictions and put them in a consumable application and routable network that allows them to generate these risk-adverse routes in a way that can be applied to those industries.

[00:49:16] JM: Cool. What have been the difficulties in implementing those?

[00:49:22] MP: One of the things even I was surprised by a little bit was that we implemented a proof of concept for this in about a day. One of the tools that we provide is a routing service. In addition to allowing a developer to call out, "Give me directions from point A to point B," it also has much more advanced parameters to solve for problems like given a hundred different pickup locations and a thousand drop-off locations, optimize the entire workforce for that, right?

One of these advanced configurable parameters is this ability to dynamically overlay or adjust the network. What we did was we took these risk scores and overlaid them on to the network through just an API parameter that allowed the solver generating the routes to take the risk into account and give us less more risk-adverse routes as an output.

I think it's a really cool use case and [inaudible 00:50:24] as a proof of concept, the way to take this into production would be to actually start joining the risks into the network itself so that you're solving for the route across the entire optimization process and not just overlaying a single bit of risk information on to the network from an API call.

[00:50:45] JM: Got it. I did want to ask you if you've seen any strange impact on your work or the people you're working with, the customers, from the changes due to COVID-19.

[00:51:02] MP: Yeah. I think, strange impact is an interesting way to phrase that. I think there's definitely been a lot of impact whether or not that's strange. It definitely abnormal, I guess. As a company, we have dedicated a lot of resources to helping our users respond effectively to COVID-19. Our software is used by FEMA, the World Health Organization, Census and others that are boots on the ground engaged in helping respond to the pandemic and the disaster.

As a company, we're trying to help them however we can, whether that's through granting software at no cost or providing our expertise to help them utilize the software or help their partners utilize the software to collect data on COVID-19. Collecting data in terms of what hospitals have access to PPE or where are the cases occurring. Then once they have the data, communicating it out as well as analyzing it to inform some insight in the response process. In terms of strange behavior, I think the strangest behavior is that I feel as a company and as a broader community, we're really focusing our efforts to respond and it's been exciting to be a part of that.

[00:52:24] JM: Very cool. Well, Max, is there anything else you'd like to add about the company or project you're working on?

[00:52:32] MP: I think there's a lot of exciting technology just from my perspective that Esri is working on right now from the reengineering and the Kubernetes and containers to running these advanced analytical operations in Spark as well as making a lot of the location services like that routing network available to developers. There's really a lot of exciting technology. I think it's going to evolve rapidly within the next year as well. If you're interested, feel free to

reach out. You can find me on LinkedIn or check out esri.com, and I think it's an exciting time in the location space.

[00:53:08] JM: Yeah. I realize now, I forgot to ask you about the mapping accuracy. We talked a little bit about – You and I talked offline about the fact that you listen to that episode about Facebook and Facebook mapping accuracy. Do you have anything to add about how you maintain accurate maps across ArcGIS?

[00:53:28] MP: Yeah, I think. In the Facebook episode, they were talking about accuracy of maps. Making sure the data is accurate in the right locations so you can find the appropriate store and navigate to that store or discover it, right? I think what's interesting in sort of our dimension of accuracy is that we also have to consider the coordinate systems and the way that that data is actually recorded or the reference system that's used to record that data.

The traditional way that we think about recording geographic data is in latitude and longitude, but when we think about accuracy of doing calculations against that data, you can't use latitude and longitude to calculate area, for example. We have these use cases where users need to have very accurate representations and not just their data stored in a database, but also the underlying reference system for the globe that makes it positionally accurate in the real world.

So we've built out this underlying engine that allows developers and users to work with many different representations of the globe so that it's accurate in real space and also can be analyzed in a way that's geographically meaningful.

[00:54:52] JM: Awesome. Well, Max, thanks for adding that bit, and it's been a real pleasure talking to you.

[00:54:59] MP: Jeff, yeah, thanks again for the time, and been really nice to connect.

[END OF INTERVIEW]

[00:55:11] JM: The Uptake is a new show about all things tech and community. It's hosted by Anna Chu who travels the world of technology uncovering people's journeys, and each episode

showcases communities around the world helping listeners understand the value in community activity and exploring the different ways that they can stay up-to-date and explore different ways of doing things with technology. Every episode has a focus topic, guest perspectives from Microsoft MVPs and community news. There're also updates on events and conferences, virtual conferences these days and more. You can find it by just searching for the Uptake in your podcasts.

[END]