EPISODE 1456

[INTRODUCTION]

[00:00:00] KP: Modern organizations eventually face data governance challenges. Keeping track of where data came from, what system updated it, and what ways they can update it are just some of the issues you'll face. Large organizations have additional challenges like training, onboarding, and capturing the institutional knowledge that leaves with the departure of key team members. As a team grows, these challenges often grow faster for unprepared organizations. Select Star helps companies unlock the full context of their data. Their solution automatically catalogs and documents your database tables and BI dashboards.

In this episode, I interview Shinji Kim about the functionality of Select Star and how companies have achieved successful adoption.

[INTERVIEW]

[00:00:43] KP: Shinji, welcome to Software Engineering Daily.

[00:00:46] SK: Thanks, Kyle. Nice to meet you. Good to be here again.

[00:00:50] KP: Absolutely. Well, for listeners who didn't catch your earlier appearance, could you give us the quick summary, what does Select Star do?

[00:00:58] SK: Sure. Select Star is an automated data discovery and governance platform. What we do in a nutshell is scan the metadata, the data warehouse BI tool or different applications and join it in a way where anyone can use to find and understand the data they have access to. So usually, instead of just getting the list of people’s dashboard, or databases and schema, we will also tell you for each data set or dashboard, where the data is coming from, who's using this inside the company today, what are the dashboards that are built on top of this, and how are other analysts using the data?

So it provides a lot of auto generated context, and also, we also put a documentation on top of it. That is basically a base portal, like a data portal, or like a hub that a lot of people can use to additional documentation, define metrics, and also utilize it on top of any other data stack.

[00:02:09] KP: So when I've had roles that were really a sort of analyst positions at heart, despite whatever the title might be, where I was going to spend most of my time in the database, looking things at things, preparing reports, that kind of stuff. I've always had some sort of IDE be it SQL Server Management Studio, or the Snowflake Portal or whatever. There's some sort of formal workbench I'm in. How do you imagine your tool sitting alongside whatever the analyst is in besides that?

[00:02:36] SK: Yeah, that's a great question. Today, Select Star is primarily used as an independent tool. A lot of the data team use Select Star as the starting point, as if it's their Google for data. So this is a place where you will find and learn more about either the table stakes or how the tables are being used before you decide to actually create. If these are tables that you already know and you are acquiring every day, then you may not use Select Star. But if you have hundreds and thousands of tables, or if you need to dive into a data model that you haven't used before, Select Star is a great place to start. We are starting to add more support and features to enable more like normal data, consumers be more familiar with just the BI tool or they want to stick with or their like an IDE. So we do have an API available for teams to be able to pull some to different tools, and we do have a few integrations and plugins that works. So that Select Star data can be displayed other tools as well.

[00:03:56] KP: One of the hardest parts of an analyst life is the early period of time when they've just started a company. As you described, there could be hundreds of tables, and there's just a lot to take in. How does Select Star help that professional?

[00:04:11] SK: You're talking about for data engineers or data analysts or?

[00:04:16] KP: Well, I guess I had an analyst in my mind, but if it's applicable to both roles, let's hear both routines.

[00:04:23] SK: Sure. So our primary user base today, our data engineers, and data scientists and data analysts, and most of the time, we work with data platform teams, that are tasked to better environment manager for the data consumers primarily that the analytics team. First and foremost, data platform teams, once they set up Select Star, they get to understand, which are all the tables that are being used the most or not being used anymore for our domain top users who should be kind of like the owners to document the data set. So with that insight, as they start inviting the data analyst team, the

analyst teams can start putting either documentation tagging or ownership of data. And also, they would use the Select Star primarily to search and like finding and understanding the data better.

[00:05:22] KP: In my own experience, I know one of the hardest parts of the job is when you first join a company and you're unfamiliar with the schema, as you said, a data warehouse could have hundreds of tables, there's a big learning curve there. How can Select Star help a new employee get their head around the data?

[00:05:40] SK: Yeah, for sure. Data discoverability is really where one of our focus is. I mean, our major focus is today, and where that manifests itself is to help new analysts or people that are starting to learn more about data, to be able to be more proficient on the text. The way to do this is first and foremost, by providing a really simple user interface, regardless of how technical you are, you should be able to easily search and filter through the information that Select Star displays. And as much as the interface is really simple, because there can be a lot of deeper, detailed data, a lot of that is something that analysts and data scientists can dig into, to understand the dataset better.

In fact, some of our customers are having select star as one of their data scientists or data analyst onboarding package. So from like the week one or day one of that data team member, they get to use Select Star to get on-boarded to the data that they work with.

[00:06:57] KP: When using your user interface, what type of searches can users do against the data?

[00:07:04] SK: Search. They can just type any keyword, and that will search through what's been indexed from Select Star regarding the name of the asset, as well as description will also index through any tags, or that can also be eventually like, also you users and whatnot. So anything that matches the keyword will show up. One thing that we do really well is combining that search result of the relevance with the keyword, with popularity.

One of the things that Select Star does underneath when it connects to your data warehouse, is looking at how the data has been used inside the company in the last 90 days. And part of that is looking at the Select queries executed and by whom. A lot of that comes down to hence how many unique users have acquired this data set how many times. Based on that, we have a relative score that compares every single table, every single column, and every single dashboard to assign this popularity score. Based on the popularity score, we will rank sort the search result, which is one of the aspects that we consider for

our sector. So in Select Star search, you will be able to see everything that matches the keyword, but it's also very often that you will be able to you – you're going to get like hundreds of results if you just put, let's say, like an order or user secured. But because we will use the popularity ranking by the one that everyone is using, or the ones that are most official as the top search results.

[00:08:52] KP: I really liked the idea of looking at the actual search logs to come up with your popularity metric. But I'm also curious, do you have to do any statistical analysis or post processing on that? Because the most popular word is “the” and that's rather boring. Do tables turn out to in a more obvious way when you look at raw popularity? Or is there some munging to be done there?

[00:09:16] SK: The popularity it's not by the usage of Select Star, it's actually coming from SQL queries and the dashboards. So “the” would probably want a shift. So what we do underneath with the query history is that for each successful query, we will look at which table and which column has been mentioned. How many times, by whom, when? And based on that we will run stats to calculate the popularity score.

[00:09:46] KP: Obviously, you can index table names and column names and their types and things like that, things all database users know exist. You'd mentioned that there are tags as well. Could you expand on what other types of metadata you can index?

[00:10:00] SK: So, in Select Star, we allow users to create their own tags. And currently the tags can be a category tags, which are going to be collections of data. So this is where you can use to tag whether this dataset is related to like sales, marketing, finance, operations, so on and so forth. We also have a tag type called a status, where you can mark any table or column as PII sensitive, gold, silver, bronze, or certified or analysts the probe type of status of the data.

These tags that are Select Star specific, and we do provide a way to sync these tags, and which columns or tables are tied back to Snowflake. This is one of the integrations that we have with Snowflake. We do have a new feature coming up where we will be able to sync tags along with DVTs. But today, most of the tags are currently being used within the Select Star context. And obviously, this is also available as an API, so if the customer wants to replicate this in other systems, they can use it programmatically.

[00:11:16] KP: Can you expand on any of the ways you're seeing different organizations use these or features you have in the tags that can help people or better organize the way they leverage the tagging system?

[00:11:25] SK: So the part that we really recommend our customers is to use tags as they're part of tables, columns, dashboards, and or documents. We have all tags where anyone can see the full view of how the data is laid out within the organization. And from there, if you were to go into that specific category tag, like finance, or marketing or sales, they'll be able to find out, not just like tables or data sets, but also documentation or dashboards or users will have its own description, owners, and some explanation through the icon that it has. This is a way that we see how having it in “folders”, with tags, you can basically apply multiple tags in the same asset, and I think that's really neat way to organize data overall.

[00:12:31] KP: You provide a suggested taxonomy or do I come up with my own categories and subfolders and whatnot?

[00:12:38] SK: We give some suggestions, but what we find is that most of the time, customers already have their own way of organizing their data. A lot of them also already do this in their database through different schemas and databases. We try to just ensure that they are aware of the structure of the tags, and they can match it together organizing data internally.

[00:13:04] KP: So many companies are concerned about privacy and PII data and for good reason. One strategy is then don't let any vendors in, keep your data kind of locked up. What can you do for someone who's very security conscious about their data?

[00:13:20] SK: Yeah, so currently, our SOC 2 Type compliant, so we have an internal measures and all the data is encrypted at rest and in transit. We've gone through a number of pen tests and auditing to ensure that we have the most secure way possible to treat our customer data. And by default, our customer, we access from Select Star is metadata and logs only. Meaning by default, the access that we get, or the service account that our customers will create to connect to Select Star has only access to the system level metadata, not the actual values of the data.

So with that, by default, most of the time, do not have any access to sensitive data or PII. I would say like less than 1% of a chance, if we end up seeing any values that might come through, that may be

sensitive, that have a like a PII just through the logs or SQL queries, that is something that we stripped off or any fields that are PII. This actually happens in memory before we actually parse through the rest of the query. But whenever we match any tags that are marked as PII, when we process it, we will take that value out first. So within Select Star, you want to search through any values or any sensitive values within Select Star. Does that make sense?

[00:15:04] KP: Absolutely. I'm curious if you've seen any patterns in the way different organizations adopt the tagging principles. I could imagine one company has strict rules about how you tag things, and others are more freeform. What are typical deployments like?

[00:15:19] SK: Yeah, that's a good question. I would say it really depends on the organization. But more often, what we recommend, and also a lot of customers do today is afford the data platform team, or the data analyst team to come up with like a high-level taxonomy of how they would categorize either the status or the domain level that they want to specify. Once that is in and they have so many examples, as their data analysts, or data scientists get more used to Select Star, they start either creating sub tags, or they will just use those tags to different tables, columns and dashboards. This is how we generally, so that there is a good high-level framework that you are working with, at the same time, you are giving freedom for the actual data consumers and the main users to put more tags and add any specific tags that they need.

[00:16:30] KP: There have been a number of occasions in my career where I've left a role and kind of taken notice of how much institutional knowledge I was taking with me. That I was the only one who knew that the development team changed the definition of a certain column, and you had to adjust for it, things like that. How can Select Star be used to socialize some of these hidden wisdoms, or skeletons in the closet, if you will?

[00:16:56] SK: Yeah, this is something that we're trying to do. There's so much we can do, obviously, just from the data logs and metadata. But today, one of the things that we bring out automatically is column level lineage. So end to end lineage for Select Star will show from your route tables and the downstream effect all the way to the dashboard level. And this will also be shown not just between one table to dashboard, we will also show how the transformation has happened along the way, if there are any transient tables in between, or if there are different data models that has been created in between, whether that's on the database side or the BI side.

This really brings a lot of visibility of whenever there is a change that might be made on a table or a column, you get to see right away, we have like an area where we share all the assets that's going to get impacted – for our domain top users have those assets, and who created those assets and their own popularity, and when's the last time it was upgraded. What this gives you is that, let's say if you're thinking about replicating a table, you can look at whether the downstream assets are actually being used today or not. It may be safe to deprecate it, even though there may be like 10 different dashboards, if the dashboards are themselves I haven't been looked at for the last six months.

So that type of analysis is all baked in, and I will say that it's one part that really brings out this context. The other ways that we try to do this really easily is by providing a way to leverage some of these relationships between the data assets that we have uncovered, along with lineage and ways that we determine what we call related tables and similar tables. So in Select Star today, if you update column description in one place, you can propagate the same description down to either the downstream lineage or upstream lineage, or even tables that look exactly the same, either duplicate a table or tables that look very similar.

What this means is you can write documentation once and updated without having to copy and paste documentation. Other parts of aware we are trying to make this easier is by bringing an adding public documentation that already has like, let's say, like the definition of the main public SaaS API's and how that would – if we see like a matching table column name, we will give you a suggested documentation. This suggests that documentation also works through lineage and similar tables. Even if you don't have any documentation that's been updated, if there is somewhat of a related documentation somewhere in the metadata of Select Star, you'll be able to see that without having to wonder about or asking anyone else.

Regarding the transition of the context of data, one part that we recommend is setting the ownership as a team. So today, we allow you to set business owners and technical owners, and that those can be – so that it's not really just geared towards the individuals. If it's a team level, team members get like notifications, and they are responsible for like either maintaining the documentation or maintaining the quality of data. That's like another way that we're trying to preserve the context. But at the end of the day, there's so much automatic context that we gather from which data sets are being queried, who's querying this, and where are they using this data on? Regarding what is really the purpose of this query, and what is the business context that's specific to the dashboard is something that we recommend our customers to add as an annotation in the form of tags and documentation. And

anything that you add, we are trying to basically either replicate it or make it super visible, so that there is a single source of truth that you can refer to and everyone can discover it really easily so that they don't have to do the same work again, even if you like leave the company, there is some context left in Select Star.

[00:21:41] KP: I know Select Star is a solution that sits on top of my data warehouse, could you expand a little bit on what technologies you integrate within support?

[00:21:49] SK: Sure. So today, we support every major cloud data warehouses, including Snowflake, BigQuery, Redshift, and Postgres. We also support direct integrations with cloud based BI tools, largely also with Tableau server. So that's not all that can be behind the firewall, but we can integrate Tableau, Looker, Mode, and Site Sense. We have other integrations that are coming up very soon, including Power BI, Sigma, and MetaBase, as well.

[00:22:25] KP: So to the best of my knowledge, there's no like unified agreed schema for how we describe metadata and databases. In particular, I've done it on SQL Server and Postgres and MySQL, and I believe they all work slightly different if I want to query. Could you talk about some of the challenges of these integrations?

[00:22:45] SK: Yes, it is a challenge. That's all I can say. Everything in our house, and also the IT are very different in a way how they store metadata. Also beyond metadata, because we require query logs and activity logs. Getting that data also is slightly different per database, or any data source that we integrate with. For example, for Snowflake, we just need to access to their Snowflake database account usage schema that includes primarily just metadata. And underneath a usage, there is a table called query history where we will get the logs from.

Whereas if we were to integrate with either Redshift or Postgres, first of all, the admin on the customer side would have to enable, like audit log and like user activity log, and then they would need to also like, make sure that those logs are available in the S3 bucket that we have access to. So that's kind of like just an example of the hoops that you have to jump through to get the right metadata.

On the BI side, sometimes we get, in terms of the activities – I guess, let's talk about the metadata first. On the BI side, like for example, what Tableau has like the structure of workbooks and data sources, and within workbooks, there are dashboards and sheets. Whereas like for Looker, there's a separate

data model for like a Looker ML model with explorers and dashboards and looks. And then for mode, it primarily the queries and mode reports. So these are all different types of metadata that we need to consume through and also a match in the sense where it looks consistent within Select Star. Plus it varies between the tool regarding how we could get the understanding of the popularity and lineage. So that really comes from whether the BI tool also supports sharing the information around views of the dashboard, versus whether we need to get these types of information from the queries that the dashboard generates whenever it runs.

[00:25:01] KP: My mind when I'm learning about Select Star goes right to how this would help my onboarding experience. But I'm also curious as you've been able to observe different customers stay with the product for a while, what's the experience like a few months in. Nce maybe my low hanging fruit questions are answered, what sort of insights can I start to get then? Or do I typically get in my day to day workflow once I become an experienced user?

[00:25:25] SK: Yeah. I will say it's really just getting it as a day to day workflow, and helping others. So is like a knowledge base of your data, right? Initially, we have kind of like our own like chart of reaching the data discoverability, like full data discoverability of an organization, which starts from having the data platform team to first like incorporate Select Star and ensure that all the fiscal metadata is showing correctly and is in there, whereas the next step would be for data analysts teams and data scientists team to start utilizing Select Star and adding some high-level tags and some documentation. And then, the next level will be kind of like making sure that ops users and other quote “citizen data analysts” to be more comfortable utilizing Select Star instead of always going back – going to the data team for support, for them to start finding information.

During this second and third phase is where we start recommending the users to add more tags, ownership and documentation. Also, in this phase, you can start adding documentation. We have a metrics documentation, which is like a more metrics definition, as well as like high-level concept documentation that you can start using, which will get to a point where the last phase where you can start onboarding more business users on Select Star.

So throughout the cycle, I would say, well, we see consistently after the first two to three months, and then another two to three months after, and then another is somewhat new set of users starting to get a acclimated to Select Star, and starting to get value out of it in different ways. And kind of starting to contribute their knowledge or their contexts into the platform. This happens first as like a data

dictionary, like tables and columns, documentation to a metrics and concept documentation. Also, we have a discussion, which is like comments on any comments, questions, answers, a lot of that, and that is all searchable within Select Star. So I would say that's like where the full cycle of communication, answers, documentation and update of both metadata and the context being happening, which can be applicable, not just for the technical team, but also for the business teams.

[00:28:01] KP: I'd love to expand on documentation. I think everyone agrees it's kind of important. But we've all had our own frustrating experiences with it. To what degree is the documentation auto generated? And what are some of the ways you've seen people be successful with your tool and documentation?

[00:28:16] SK: Yeah, documentation is something that we all want, that we all benefit from. It's not something that we want to do. So the part that we are – I mean, I will say, we are continuously trying to make this experience easier. But today, what we do is first and foremost, bringing out any existing documentation to Select Star, so we will automatically load anything that's been already defined in your database, your BI tool, your DVT report, or DVT models. Or even if you have it in a different tool, you can also update the file to update them all together. And within the UI, it's very easy to just like one click and update and that will be saved.

On top of that, the part that I mentioned earlier, related to propagating the description is something that a lot of our customers are starting to use. This has actually just got released like couple of weeks ago. This allows you to basically update or write once and have it propagated everywhere. Let's say you have description in your raw table that can be replicated to like thousands of other columns that uses that specific data, within it, we determine whether the data has been inherited directly versus whether the data has been aggregated or transformed. So the propagation happens only for the ones, the columns that are using the data as is. We actually do this the same way for tags propagation. If you have a PII tag and you want to find all the columns that's using the same data, like the actual like **[inaudible 00:30:07]** it's basically an example.

Last one, at least, we have our own, what we use as "suggested documentation". So if we observe – you didn't necessarily propagate the documentation, but if we observe an existing documentation that can be a match, then we will show it in gray, instead of black, which is like a divide documentation, that user can also edit. But the suggested documentation will be done. If let's say you have entered your Looker ML model, but not on the table side, then you will still see like the documentation that came

from your Looker Ml, because we will look at the lineage and see that the documentation already exists there.

Similarly, the part that we are starting to add a lot more and this actually just to happen, along with five trends documentation. So any, like raw tables that were loaded on five trends that are like SaaS tool that everybody uses, like Salesforce, Stripe, Google ads, Facebook ads, a lot of those documentation is already in us as a “public reference” all docs. And if we match the table and column level, we will actually just like fill in the documentation for you as a suggestive doc.

So these are ways that we are trying to make this a lot easier to have somewhat of the fuller data dictionary for your fiscal metadata. On top of that, as you start creating data metrics definition and concept documentation, within any descriptions of Select Star, as well as anything in the discussion, like comments, you can mention user’s table columns, or dashboards, any assets. And from there, you will also create like a back link on the resource so that you can let's say have like a one large documentation that explains all the core concepts, and anything that you mentioned there, you will see the backlink on that table. So even if you have like one document with like 10 different table mentioned, all those tables will have that document as a link. This allows you to also not having to document many places in different ways, and you can manage, basically a fewer number of documentation that has relevance to multiple physical data assets.

[00:32:35] KP: Very neat. Well, data lineage is one of the particular features that most caught my attention when I was researching Select Star. I've been in a lot of meetings where there was some report pulled up and people said, “How did we get to that exact number? And did it account for all of these other things?” Could you describe a little bit about the user experience? When I have a question about a number, how can I trace it back?

[00:32:56] SK: Yeah, that is definitely a common use case of data lineage. We have, I will say, many use cases of data lineage. One is, when an engineer is thinking about changing a table or structure of the table, they get to see what the downstream impact will look like, who's going to get impacted, which tables or dashboards are going to break. So being able to find that right away, also down to the column level, what happens more often than you think is either somebody changing a column type or column name that may break a bunch of dashboards.

The second example with a second use case will be the vice versa, the part that you just mentioned, I'm an analyst, and I was just asked, “Why is this number wrong?” Why does this number look like this? I might not have even created the dashboard. Because I'm kind of “in the data team support channel”, I should and I’m an on call support, I may have to answer that question. So now, if without lineage, you will have to go like find that query within the dashboard, and look at like what type of tables and columns it's querying. Understand, is this the query that's wrong? Or is this the data that hasn't been updated most recently? Or somebody has made a change and most recently?

That can be like multiple different questions. And from here, column level data lineage will show directly like which chart is querying which specific column? If that's the correct one, then you can just go find whether that table is querying has been updated most recently, and whatnot. So that is definitely one part that some of our customers really utilize Select Star for.

Last but not least, another traditional, I would say also, usage of data lineage has always been a more on the migration side. We've also seen it with some of our customers doing migration of their BI tool or their data warehouse, where they can see the full data model in one go so that they can make sure that everything is migrated correctly into their new data set. So yeah, those are the areas that we've been seeing where the niche can be really useful. And for Select Star, we really like try to use lineage first along with popularity, because like when you are propagating tags or descriptions or even like own hours and whatnot like these are – when you look at it with popularity, it adds more context of whether it makes sense for **[inaudible 00:35:31]**, and also like notification of other users. Being able to see top users along with that is also the other part of it. When you know all the owners assigned, but if you can find out all the top users, this is like great way to also notify everyone that has, let's say, touch the table in the last 30 days, or seen the dashboard in the last 30 days to let them know that you are planning to like deprecate the dashboard, or the freaking table.

The other use case that I mentioned before, that's also specific to Select Star is propagation of information, like tags, documentation, and ownership.

[00:36:09] KP: In this lineage, could you expand a little bit more on how you're able to track it. I mean, if I have some ETL tool out there doing some data transformations, I don't know that you have visibility into that, what's the level to which you can probe to get information about how data has moved around?

[00:36:25] SK: So today, this is primarily happening through parsing through the SQL queries, data models, and anything that's more of infrastructure of YAML files or JSON files, et cetera. So, we plan to Spark jobs, or also integrate it into like ETL systems like Airflow and whatnot. But even without those so far, for major cloud data warehouses and BI tools, we were able to store the full lineage in detail. So the way that we look at the lineage is by first of all getting into the types of queries that generates different tables. These are all DDL or DML queries, create table as Select, merge, update, or insert that includes a Select query. So that's like where the lineage first gets generated.

Sometimes we definitely see ETL jobs that's like creating a table, that's like on there, like a temp schema, or it's called temp. And then once that all the data is loaded, and it swaps with the main table. So we have a few traits that we look out for, in order to make that between the upstream of the temp table as well as along with the actual table, so that you can get the full page but not with the – all the temp table thing you don't necessarily want to see, or it's like a duplicate of the tables. Any of these jobs, you can also define them as like a “service account” and those queries won't affect the popularity as much. Does that answer your question?

[00:38:13] KP: Definitely. Well, I'm wondering also, if you're curious at all about column names and things like that, it's always occurred to me that avg_order_value, pretty much anyone who works with databases knows or has some expectation about what's in that column. Maybe not every column is quite as transparent. Is there anything natural language processing can do to help you help users understand their data?

[00:38:38] SK: That’s a question. Right now, we don't have like an NLP specific help on the search front. But we do have some fuzzy search enabled, so that anything that I guess like fuzzy searches, but by NLP, maybe that's where I should make this more fancy. If you are misspelling something, or if there are similar words, then we will match those keywords by default. So those are already baked in. I do also think that there are many different ways that the customers really put like a prefix or postfix is to indicate what type of data this is, or what aggregation or grain this data is using. So I think that's where also search really can help on filtering through the different names.

One of the views that we have is kind of like seeing all your columns within the database in one place. That's like one thing that some of our customers use to basically define, or understand how different column names have been already created, and we’ll also detect which columns have joins or they’re

associated with it. So even if you may not have primary key, foreign key define, you'll be able to know which might be the join key you might want to use.

[00:40:12] KP: Can you talk a little bit, maybe tease some of the roadmap, what's the future for Select Star?

[00:40:18] SK: Oh, there are lots of things to do, and so many possibilities with the metadata, and also as we are increasing more integrations down the road. So in the first quarter, we focused on helping users to put high level documentation and having this expanded relationship between the physical metadata and the domain level documentation. And going into Q2, we are going to start introducing more advanced features like fine-grained access control, audit logs, and some analytics behind Select Star, so that more customers can implement different workflows of either like getting access or having some kind of like approval flow for their Select Star instance.

We've started a lot of the use cases as companies just wanting to make sure that all of their data systems are available to everyone. But today, more and more customers are starting to use Select Star in a more advanced way, and this is like one of the major, hence features that's coming up. Q3, Q4, rest of the year, a lot of other integrations and also ways to enhance our end user experience currently are scheduled. There's a lot of out here.

[00:41:41] KP: Absolutely. Well, for listeners who want to learn a little bit more, or see if it's right for their organization, where can they go online to check out Select Star?

[00:41:50] SK: Today, just select selectstar.com. We have a 14-day trial. It is fully self-service, so you can just connect your data sources and get going right away. Most customers actually finished their ad hoc within minutes, even if you have lots of data, it will take just couple hours to load everything and we email you once the data is fully loaded. Yeah, it's fairly easy to use, I would say. So that's where you can start. Obviously, I'm available anytime through the intercom or also over email. We also have a Slack group that we invite our users to ask questions too. So those are areas that people can also ask questions and reach out to us as well.

[00:42:38] KP: Sounds good. Well, Shinji, thank you so much for taking the time to come on Software Engineering Daily.

[00:42:43] SK: Thanks so much, Kyle, for having me here. I really enjoyed talking to you today.

[END]