Double Click

Palantir

Q2 2021

Series No. 01

April 14, 2021

↳English Transcript

This transcript is provided as a courtesy and is intended to be viewed, and is subject to, the accompanying oral presentation and related materials, including any legal disclaimers.

Contents

SECTION

Introduction

SPEAKER

Shyam Sankar, Chief Operating Officer

Welcome to Palantir's Double Click. I'm Shyam Sankar, Palantir's Chief Operating Officer. I'm excited to kick off the first in a series of Double Click events with you today. We're going to show you how our software is addressing our customers' hardest problems and helping them with their biggest opportunities. And we're going to show you how it's doing so in unique ways, ways only Palantir can, leveraging out-of-the-box Archetypes on top of Foundry.

Today, we're building on our inaugural Demo Day back in January. That event brought exceptional interest in our platform from around the globe. But in particular incredible demand in the US, where we just almost can't keep up. Today, we'll be taking a deep dive into life sciences and industrials.

In the last year, as scientists accelerated research at unprecedented rates, as manufacturing ramped, global supply chains were tested, Foundry was there. It was there to accelerate outcomes, to manage shocks across both these industries with customers like Merck Group, 3M, the NIH.

Most recently, after partnering to distribution 7.5 billion PPE items, we've been proud to work with NHS England on their vaccine distribution program. The NHS ordered, allocated, tracked, and delivered every single one of the 29 million and counting vaccines that they've administered, and the hundreds of SKUs that are needed for each and every one of those jabs in Foundry.

We have thousands of users, from GPs at the front lines to MPs on the phone lines, drawing on the same source of truth. Nearly 2,500 vaccination sites are managed using Foundry. It helps the NHS navigate supply fluctuations, manage cohorting, all while protecting patient privacy as a first-order concern using our novel purpose-based access control technology. And we congratulate them on the tens of millions of vaccinations to date — a truly heroic feat.

Introduction Cont.

Before I turn it over to our engineers to take us through some of our life sciences and industrials work, I wanted to touch on a few core Palantir concepts. Our two platforms, Gotham and Foundry, they're operating systems for the modern enterprise.

Gotham provides an end-to-end solution, from space to mud, that integrates every single sensor and every single shooter for US and Allied defense around the world. It is a single platform that helps you understand, decide, and act.

Foundry is, by analogy, that same operating system for the enterprise. It helps you understand, decide, and act. It isn't just about analytics, it's about decisions. It's about making better decisions. It's not just about being more efficient, it's about winning — winning by generating sustainable alpha by outlearning the competition.

And only Foundry can do this. It is a completely unique offering.

The team will walk you through our life sciences and industrial Archetypes, which deliver value through out-of-the-box, end-to-end workflows immediately. Our Archetypes in every industry, not just in life sciences and industrials, are a growing and continuously improving set of capabilities that are enabling our customers to rapidly get increasing value from Foundry.

SECTION

Life Sciences

SPEAKER

Kathleen McMahon and Ben Amor, Forward Deployed Engineers

Overview

We're going to show you three ways our clients use Foundry's capabilities to accelerate clinical research and improve patient outcomes. That frankly, no other platform can do.

Firstly, it allows them to be totally compliant with the stringent governance requirements around healthcare data, whilst still allowing researchers to access that data, and uncover valuable insights.

Secondly, how it enables them to do research in ways that are completely reproducible and transparent.

And finally, how it allows them to share and collaborate securely across their organization, and across user profiles in ways that accelerate research.

Today we are going to show you the governance, the analytics, and the knowledge sharing through an end-to-end example of the development of a prognostic model for lung cancer.

Part 1 → Governance

Life sciences organizations have more data about human disease than ever before - from clinical observations to genomic imaging, and wearable sensor data. We can all feel how important it is to have strict governance and access controls on this data. Less intuitive is how to actually meet the corresponding regulatory requirements. Foundry gives them full control of their data access and usage.

For any given dataset, governance administrators can see every place that's been used in one click. And more critically they can immediately understand what purpose the data was used for, because the original access intent is recorded automatically alongside every analysis that was performed.

Life Sciences Cont.

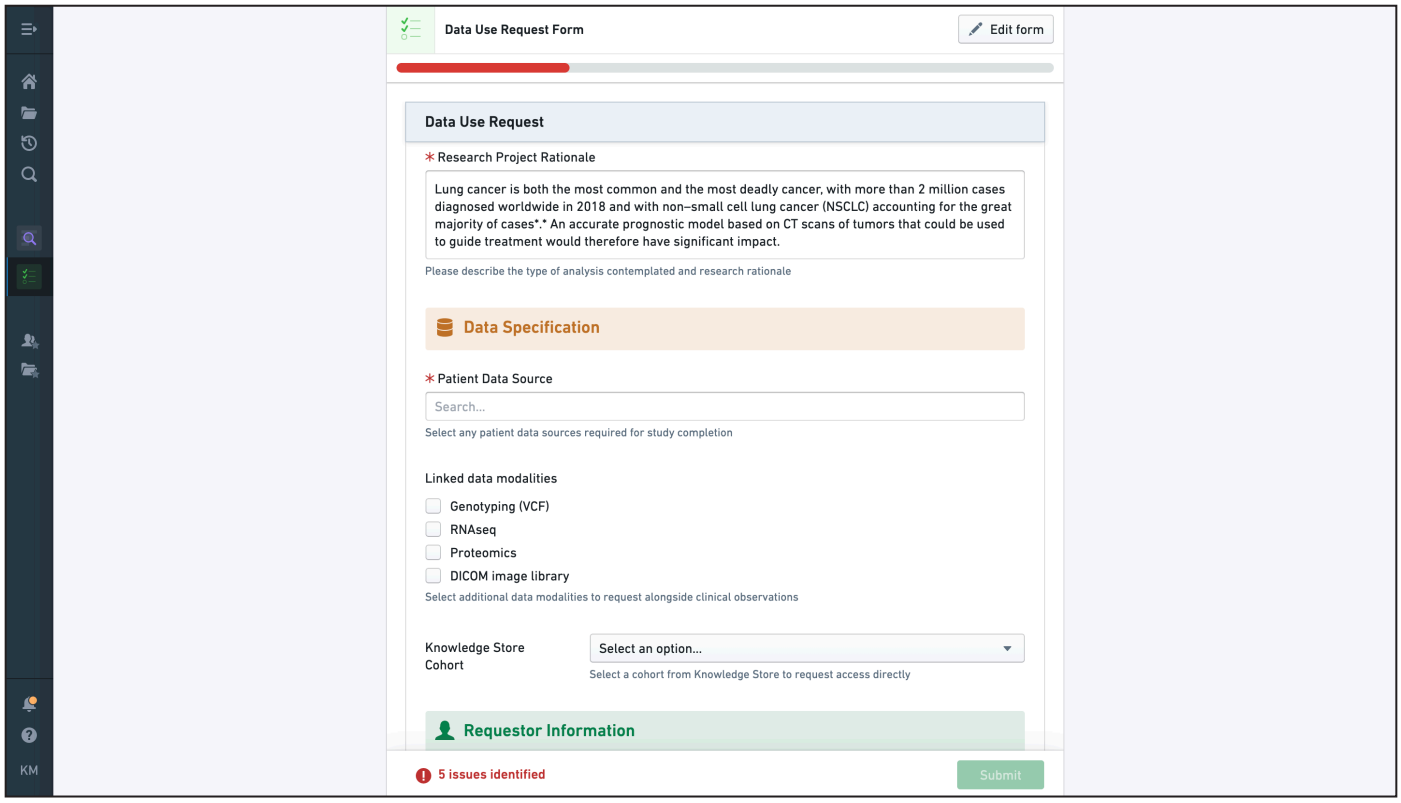

The result is increased control for governance teams and reduced risk of unauthorized repurposing of highly sensitive data.

The full framework can be deployed out-of-the box with our Purpose-Based Access Control Archetype.

Let's walk through how it works for a project like our prognostic model.

As a researcher, the first thing I'll need to do for my modeling study is submit a data use request. Having access to Foundry does not mean having access to all of the integrated data; users have to request data for the specific splices of data that they need. The request is evaluated alongside the purpose of their research access. This ensures the access is necessary and proportionate to the research outcome.

Life Sciences Cont.

In Foundry, I'm prompted to state the project's purpose and which data I need. For the prognostic model, I'll want to use both real-world evidence and closed clinical trials as training sets. I'll also need linked CT imaging data for both.

Based on the sensitivity of the data requested and the proposed analysis, this request will be reviewed and approved by the relevant set of administrators.

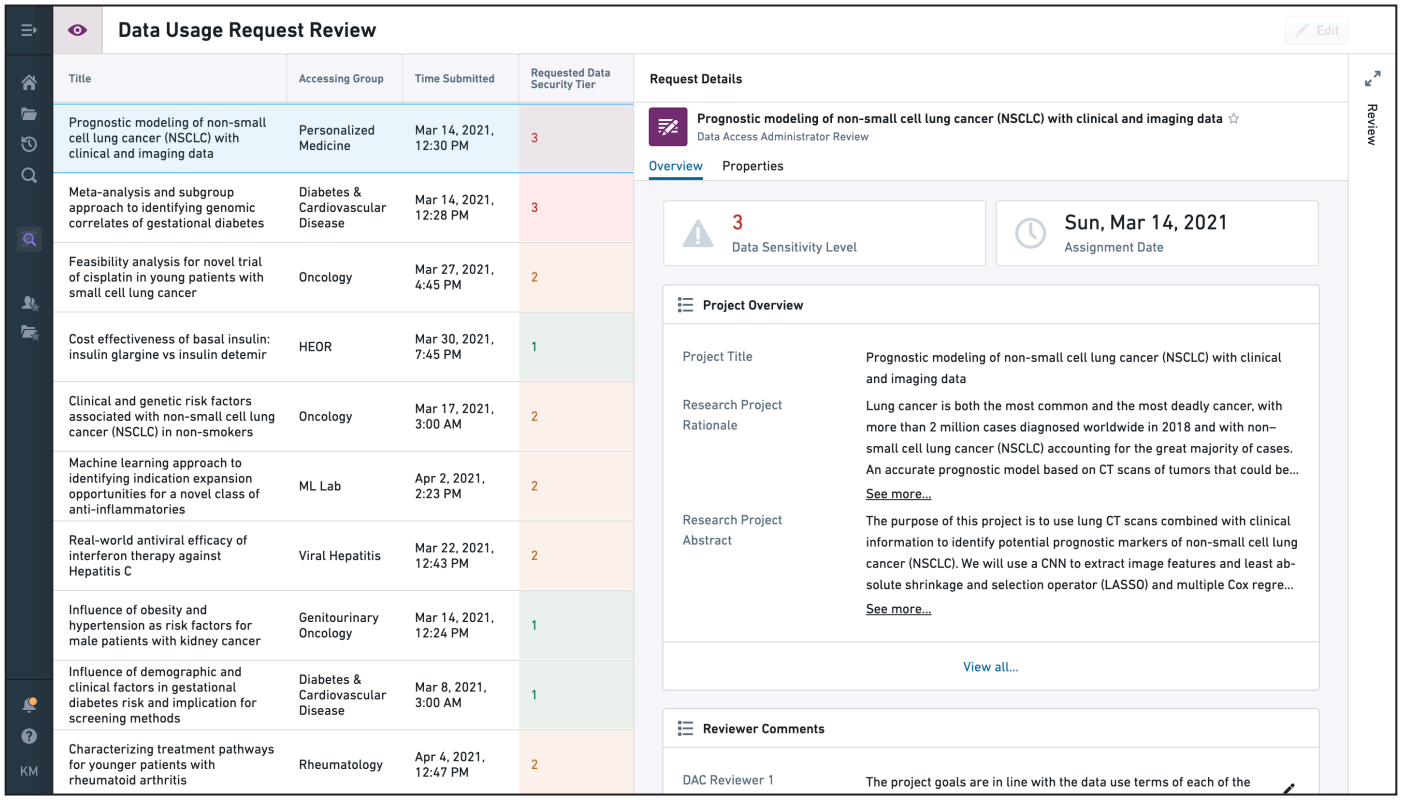

Organizations can choose how to grant data access. Some distribute their governance responsibilities to data owners, while others choose to centralize them to a data use committee. Foundry's programmable governance capabilities can support both, or anything in between.

Once my request has been approved, a new private project workspace is created and the exact slices of approved data are made available within it. This can be configured down to the level of individual rows and columns; and only those users who've been approved to work on this project will have access to the workspace.

The original request is preserved, and all work done within this project is fully transparent. At any point, the governance team can understand not just who has access to what data, but also why they were given access — with all the context that went into that decision.

This lineage is what makes Foundry so unique for sensitive data. Good governance isn't about a single snapshot decision, but rather control over the ultimate uses of any data.

Part 2 → Analytics

Our researcher now has a secure place to work with collaborators, so it's time to start developing the prognostic model.

Foundry accelerates studies by allowing researchers to use the tools they are comfortable with, while making it easier to collaborate and iterate.

We'll walk through three examples of how Foundry enables this collaboration and iteration.

Every study begins with protocol writing and feasibility testing in order to identify a cohort of patients that is both medically relevant and large enough for statistical analysis.

Foundry's cohorting app shortens this process by bringing in medical experts, data owners and biostatisticians to a common interface and building a library for reusable criteria.

For this study, I want to build our prognostic model for patients with lung cancer using a real-world data source.

I've already filtered down to patients between 35 and 55 in the US with both imaging and clinical data available, and see that I have 16,000 patients matching that description.

Life Sciences Cont.

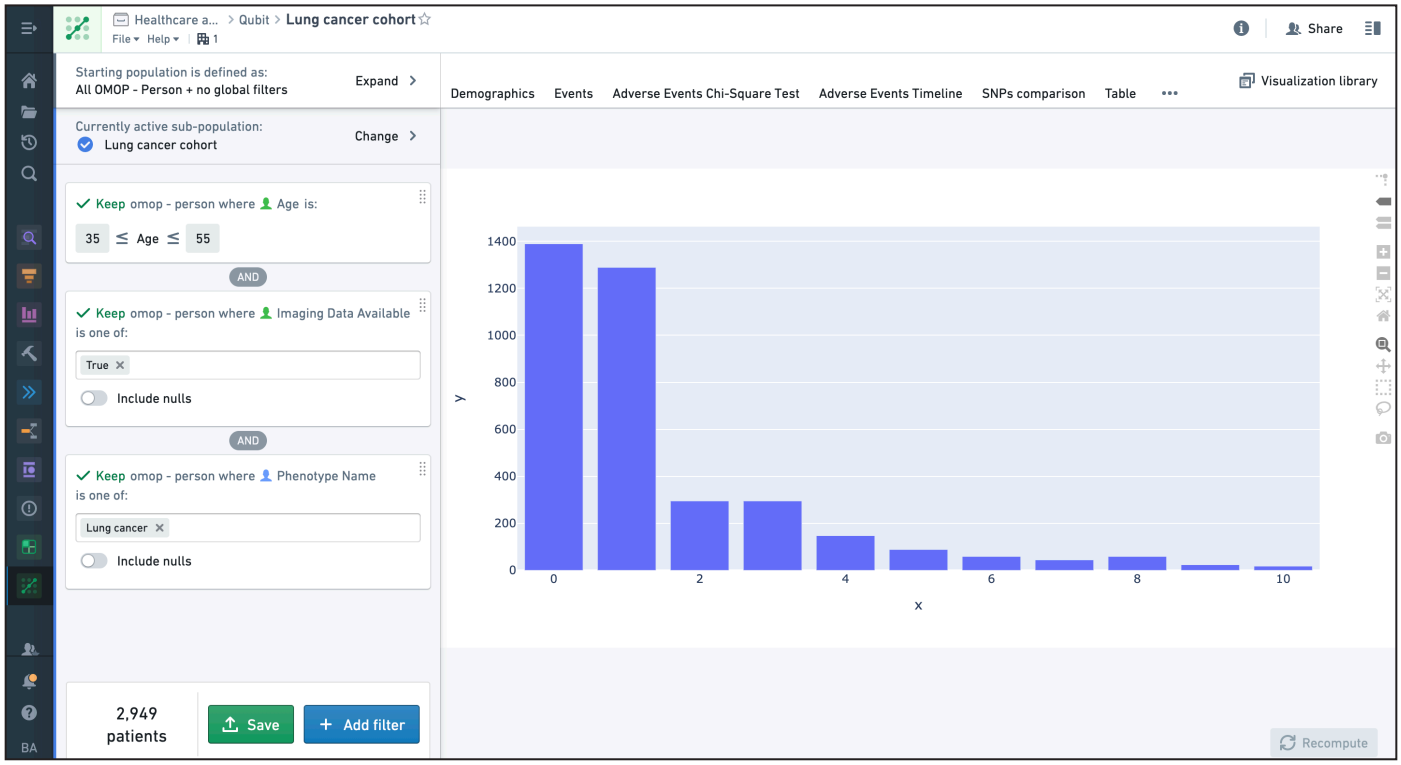

If that's not enough patients, I can easily adjust the criteria to expand the population, without cycling through various data experts to check counts.

Traditionally, updating inclusion criteria can often take a week or more, meaning months can go by before analysis begins.

Next, I need to filter to filter to patients with lung cancer. Instead of spending weeks compiling the necessary code sets and logic to define disease state, I search our phenotype library and pull from the experience of other experts.

In this case, I'll use the OHDSI-defined phenotype for lung cancer to check patient counts.

Because all my data is already standardized in OMOP format, I can automatically import and apply any code lists and cohorts defined by the open-source community. Any user can add to this library and administrators can review and promote new definitions for future use.

Life Sciences Cont.

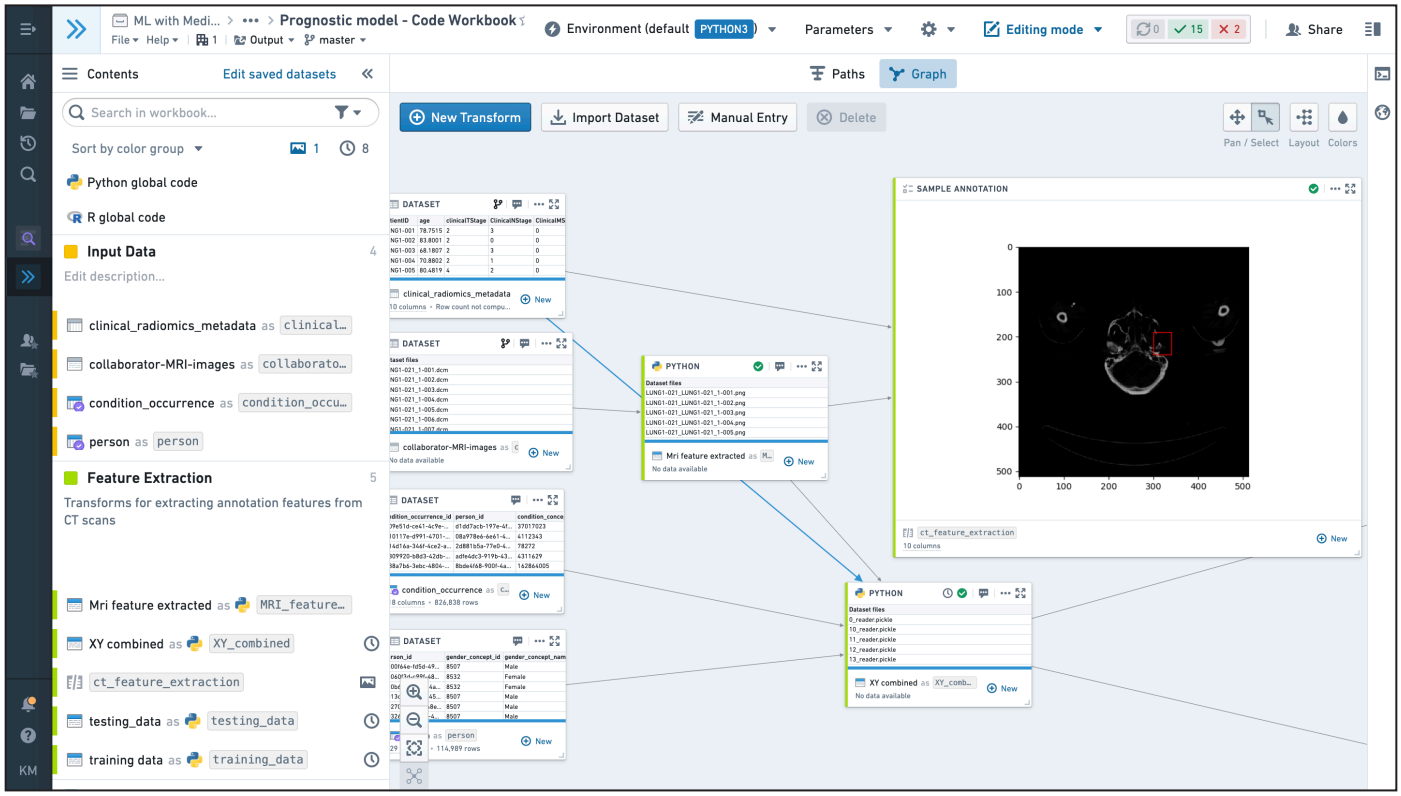

Now that we have 3,000 patients of interest, I'm ready to pass this over to a colleague to build the predictive model.

A second type of collaboration is between data scientists developing a model. We realize data scientists and biostatisticians often use custom environments for development. Our goal is to preserve that flexibility, while simultaneously ensuring that their analyses and models are accessible, auditable, and shareable.

We've built Foundry to be as open as possible for data science. Our native code workbook tool allows data scientists to jump between R, Python & SQL. And with a couple of clicks, any open-source libraries or packages can be added to a workbook.

So as for our model, all of the clinical, demographic & imaging data we need is ready to go. Our data engineers have already cleaned & linked it, and the governance team sorted permissions. As a data scientist, I get to jump right in.

Life Sciences Cont.

We can use any open source or proprietary machine learning libraries to extract features from the images. Then we'll train a predictive model using these in combination with features extracted from the clinical data, such as age or medical history.

But what if we want to use other tools for model development? We have options. Foundry has deep integrations with RStudio, and also allows clients to connect to custom HPC clusters to develop with the help of specialized bioinformatics tools.

Let's suppose another member of our team wanted to develop in RStudio instead of Foundry code workbook. She can easily reference data in from Foundry and start training a model. Once she's happy with a version of the model, she can push it back into Foundry to iterate with the rest of the team.

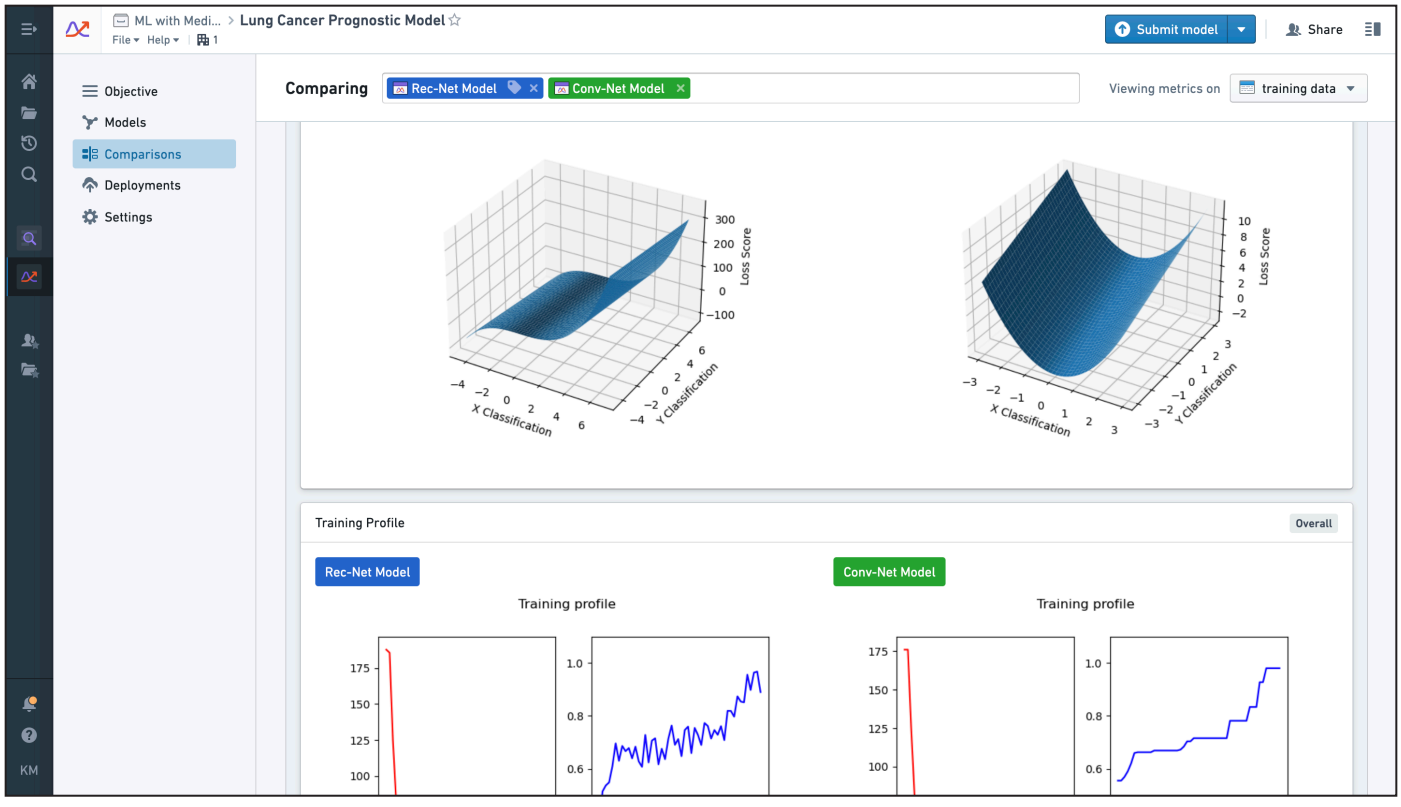

Life Sciences Cont.

This collaboration all takes place in Foundry's Model Management archetype. This archetype provides a framework for model governance, and reproducibility of results.

We'll show what this looks like by opening up our team's homepage, the lung cancer prognostic modeling objective. We can see a couple version of our model: the one we developed in code workbook, another from RStudio, and a couple more from our team. From here, we can keep track of exactly how each version was developed, update the model when source data updates, compare performance metrics, and tag versions for release into staging or production environments.

Secure collaboration is especially important with highly sensitive patient data. If we're working across teams with different data access, we can use Foundry's access controls to federate further training. I might run and refine my colleague's model on data only I have access to, but then share an improved model version with a broader group — they can benefit from the patterns in the data, without actually accessing sensitive information.

If I'm happy with my model's performance, I might move to a third type of collaboration: sharing my findings for operational use. Impactful research depends on translating results into new drug discovery or improved clinical care. I can deploy the production version of my model for use - either within the Foundry environment or to an external environment such as an EHR system.

At pharma companies, a prognostic model like ours might be used to refine inclusion criteria for new trials, or identify patients for trial enrollment.

Whereas at a provider, the model could be deployed into an external system such as an EHR, where they can inform clinical decision making.

Part 3 → Data

Models are only as good as the quality of the data they are built on. And developing a model like ours requires pulling high-quality data from multiple, disparate sources from across the patient journey.

Foundry's interoperable architecture allows me to do this by making sure that the fractured source systems can be easily unified.

Each organization has a unique IT ecosystem to capture its biomedical data, and that the data itself comes in a wide variety of formats.

Foundry provides the connective tissue between these systems.

Out of the box connections for any data storage system ensure that data of all formats can be easily integrated into a single analysis platform.

In this example, I can see that our organization has already set up connections to an internal genotyping lab, third party claims providers, and a data warehouse.

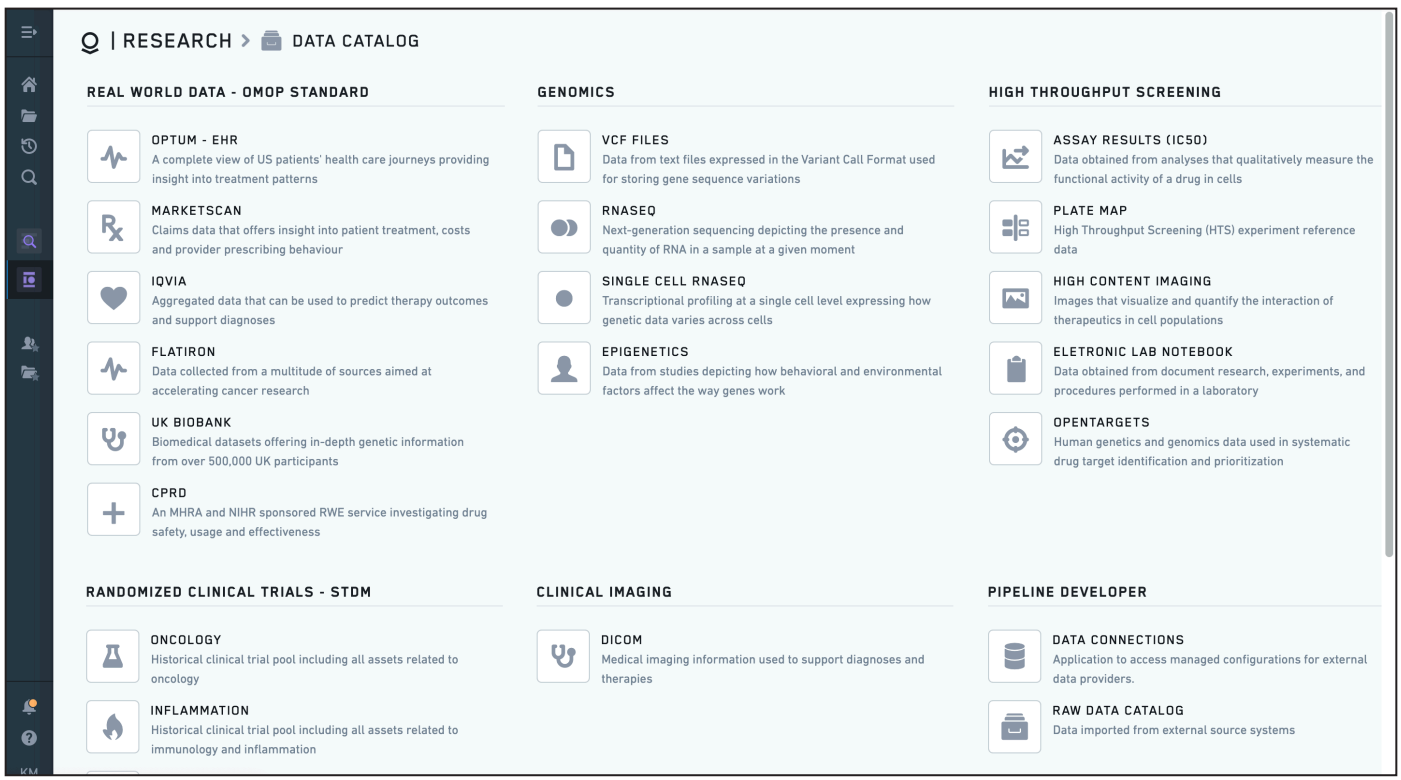

Life Sciences Cont.

We also have connections to HTS machine for use in discovery workflows, and ERP & CRM for downstream work on commercial or sales use cases.

Using Foundry, the MIH has ingested 170 million HTS experiments for potential new drugs and combined these with other data such as gin expression levels, toxicity, and high content images to discover new drug candidates for diseases such as malaria, multiple myeloma, and prostate cancer.

As an engineer, I can automatically check for data updates and set up alerts for when these deviate from their expected quality.

Data of any modality can be rapidly centralized and cataloged, so the engineering team can spend more of their time on the second piece of the puzzle: how to standardize this data across all patients.

Because not only do I need to virtually reassemble a single patient's journey, I also need to do so in a way that allows me to compare patterns across many patients at once.

Keeping with our engineer persona, I'll switch now to a data lineage view to show how Foundry accelerates this harmonization.

Each of the boxes is a dataset, and as we go from the left to right our data is transformed from raw ingests through to cleaned and harmonized analysis-ready datasets.

Life Sciences Cont.

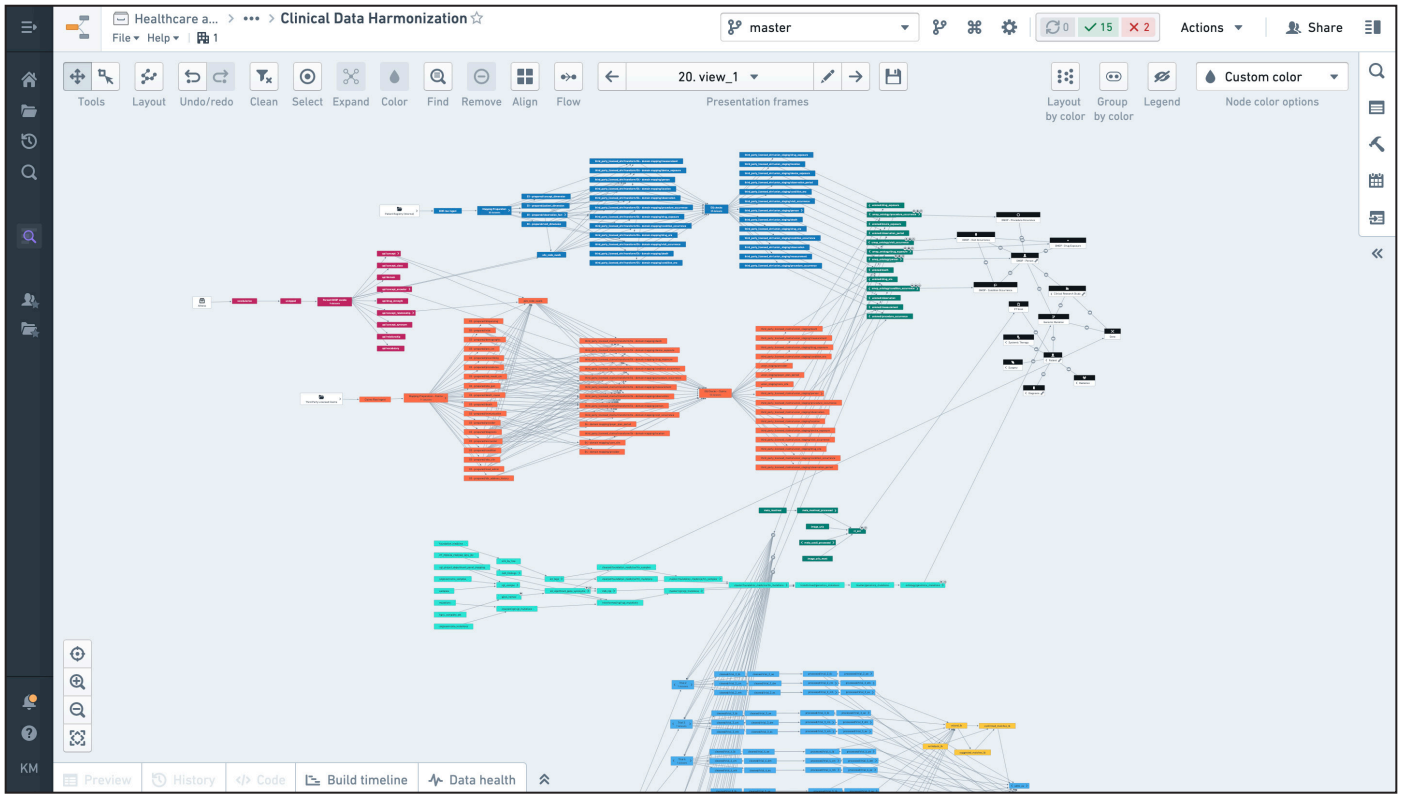

As you can imagine, our clients want their data science talent to focus on analysis, not data wrangling. As a data management team, our goal is to provide clean, reliable, analysis-ready datasets that use standard medical vocabularies and data models.

This makes it easy for researchers to pivot across sources and pull features from different modalities as projects evolve.

However, data standardization projects can take years. We've built Foundry Archetypes to cut this process down to weeks, for both real world data and historic clinical trials data.

I'll start with an example from our real-world evidence work. Let's zoom in on our graph, we have two data sources being translated into the OMOP common data model, in blue and orange.

Life Sciences Cont.

Imagine we have just licensed a third source, EHR claims. We want this data in the same OMOP format as the others so the researchers can easily develop models on any of the three sources.

In just a few clicks, I can use a Pipeline Archetype to deploy a pre-built data pipeline that converts this new data source into the common data model. All I need to do is point the Archetype towards my input datasets. Then, I'm given a set of predefined health checks and build schedules. I can customize these, or just review the deployment plan and execute.

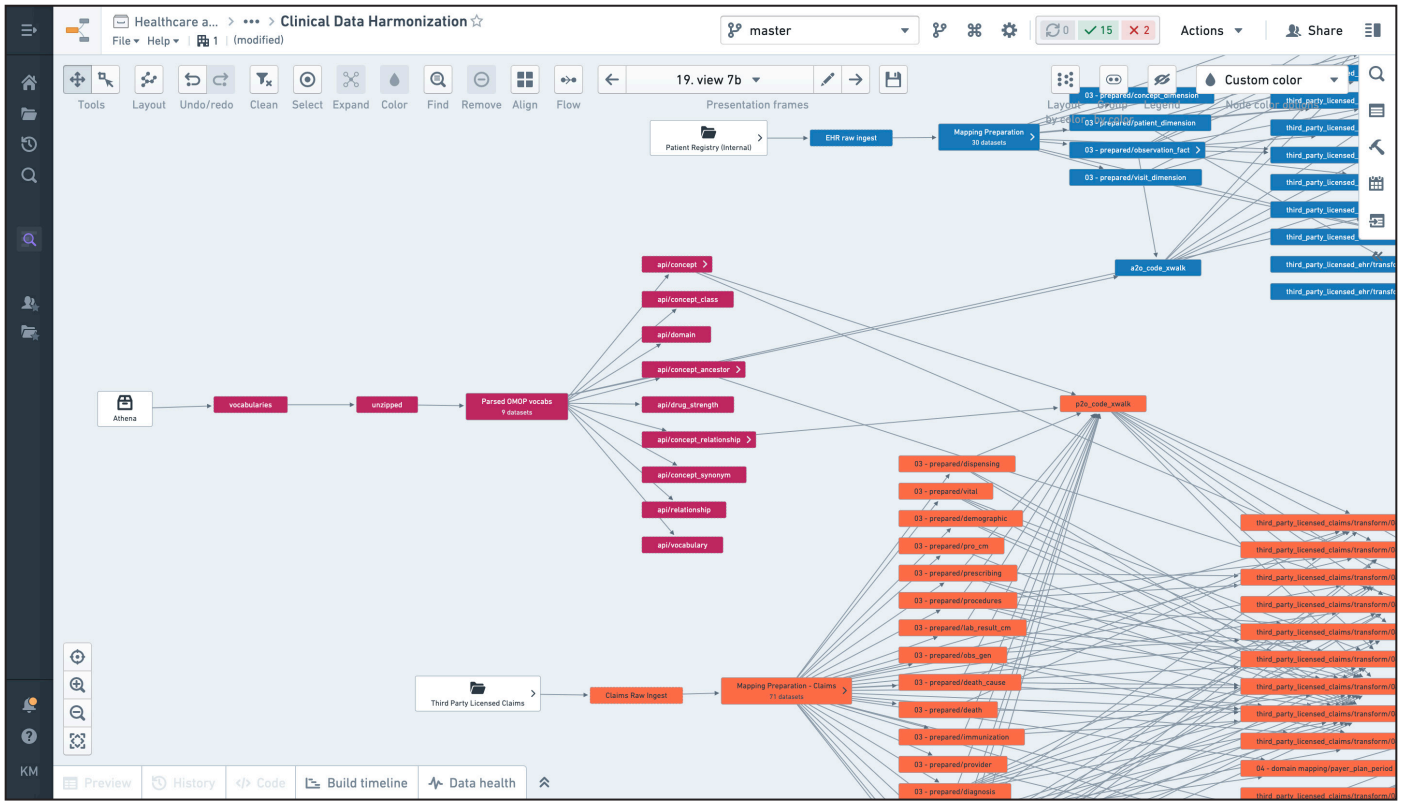

On my data management graph, I've added a full pipeline without writing a single line of code. It's in green. I've mapped source data to medical vocabularies like RxNorm and SNOMED, and transformed my source tables into ready-to-analyze OMOP tables.

Researchers in my organization now have all the real-world data at their fingertips, without having to translate between raw schemas.

Life Sciences Cont.

Another area in which Foundry accelerates standardization is with historic clinical trial data. Closed trials are rich high-quality sources of multi-modal patient data. But each trial is run and recorded just a little bit differently, leaving R&D orgs with hundreds of data silos rather than a re-usable pool of patient data.

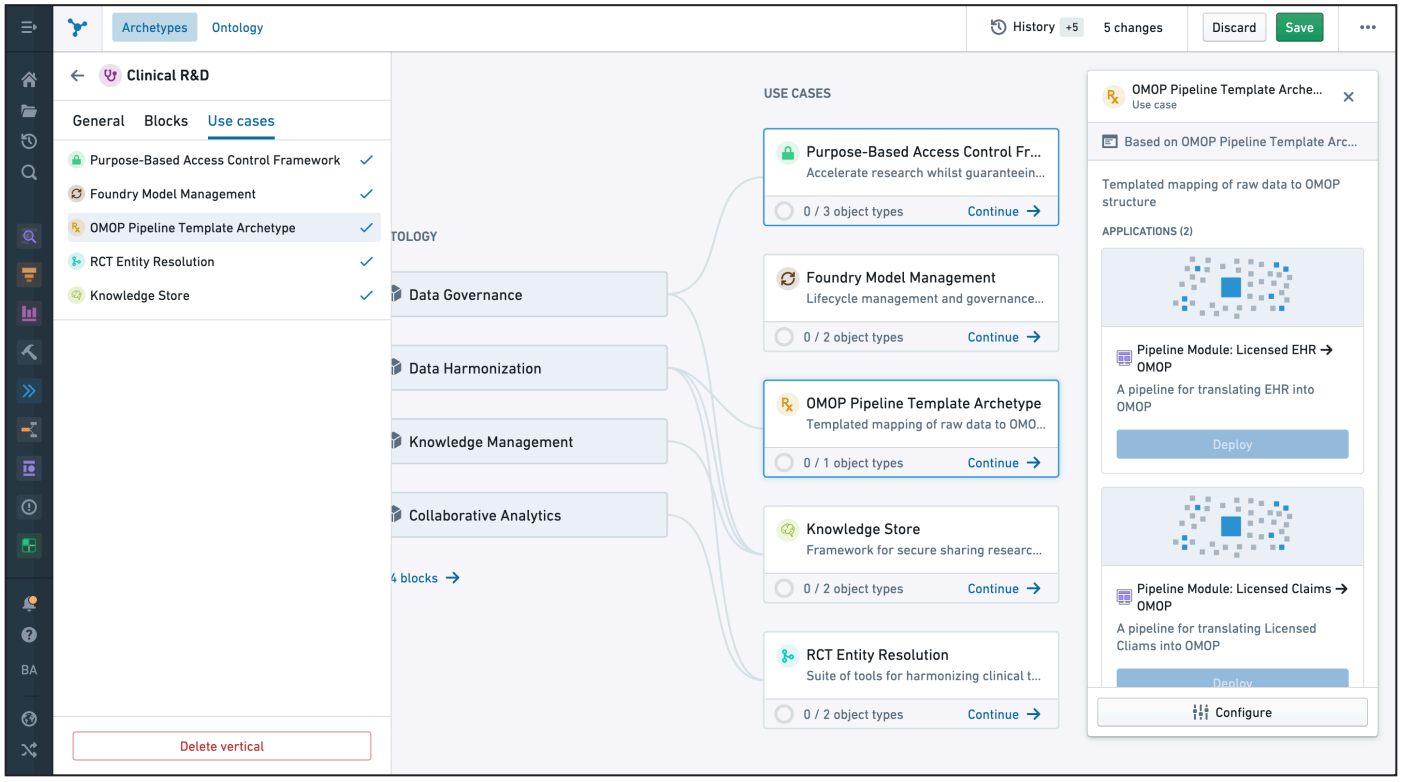

Cross trial-harmonization is the key to unlocking additional value from this data: once patients are comparable across trials, researchers can use much larger pools to generate hypotheses about indication expansion, biomarker strategy, or disease progression.

Here, in my data foundation, I've imported 10 trials from a source system. Each has 20 component datasets, with dozens of variables in each one.

Foundry's Entity Resolution Archetype guides the harmonization effort between these — it automatically surfaces suggested domain & variable matches, but then lets subject matter experts review and confirm the logical matches.

The full mapping provenance - the matches, confirmations, and transformations - are saved and transparent for any future users.

Confirmed rules are automatically applied to new trials added to the pipeline, ensuring I only have to do this harmonization once over consistent datasets.

So, I've now built a 360-degree view of each patient, and I've put their disparate data points in the context of their overall journey. Even more importantly, each patient story is now told in the same language — the standardization archetypes allow me to compare symptoms in one patient directly to patients from a different source.

Part 4 → Knowledge Management

We've looked backwards at how Foundry can be used to prepare data for research. Now let's look forward, to how this model could be securely shared and built upon by others within my organization.

Instead of starting from scratch, researchers can discover past projects and get a head start from the cohorts, datasets and logic created by their colleagues Even seemingly uninteresting results might lead to the next breakthrough down the line.

Models that have been shelved for no results might generate new hypothesis when run against updated data sources. Once the Knowledge Store Archetype is deployed, all projects are set up with a homepage for collaborators.

Life Sciences Cont.

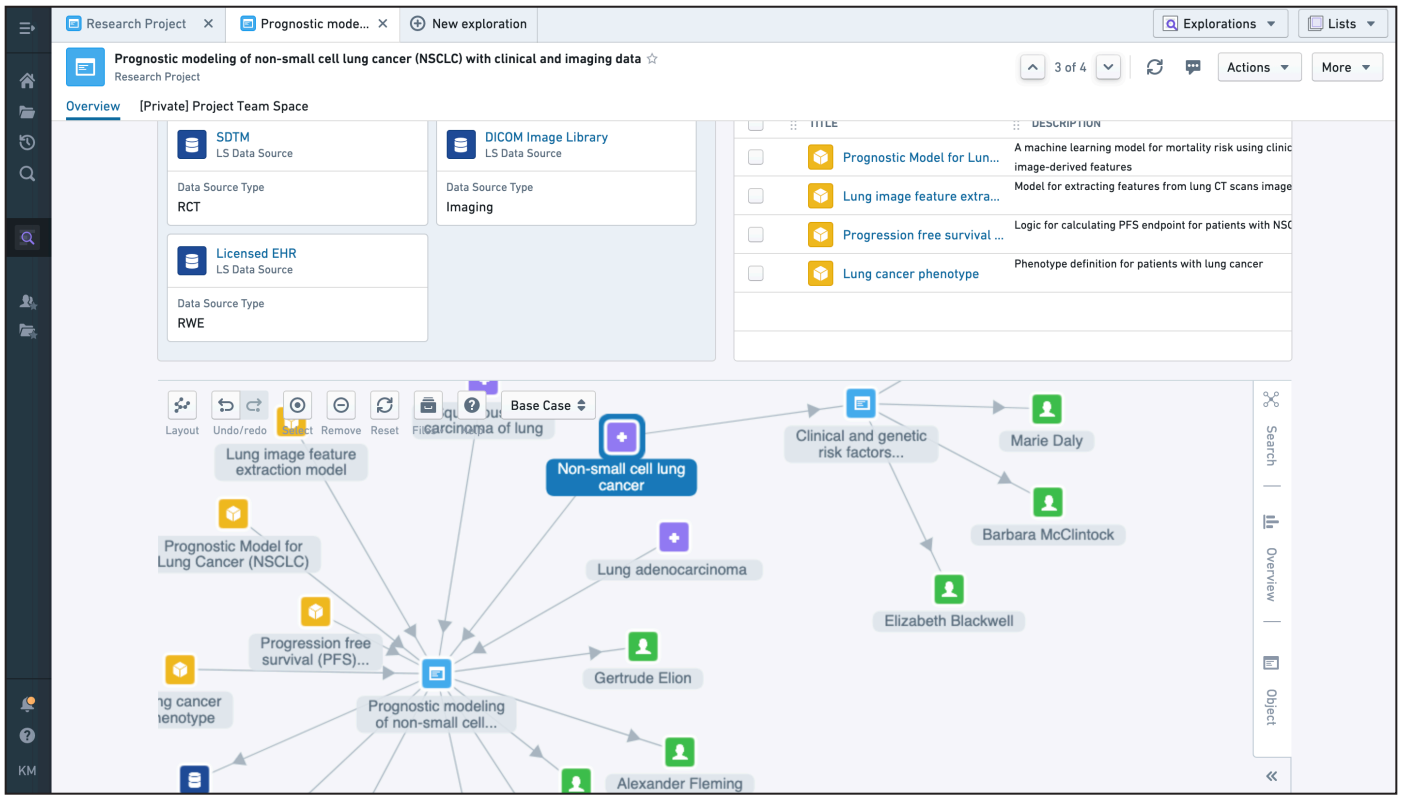

In our example, this shows the objective of my study, data I used, and links to the models I developed. We want to make sharing as lost cost as possible. My study is automatically registered as a new node on my institution's knowledge graph, and linked to metadata that might help future researchers discover it.

I can use the same framework to search for ongoing research that might accelerate my own. I remain in control of just how much I want to share. By default, anything submitted to the knowledge store inherits the permissions of the raw data used to derive it.

Other users are made aware of our project, but can only dig into my work if they have the necessary approvals in place, or the output has been designated "safe" to be shared by an administrator.

Life Sciences Cont.

In this case, I'll nominate my model for wider production usage. Once approved by administrators, it will show up as a reusable resource from my study. These published Knowledge Objects are available to discover in the Knowledge Store application. Here we can see the prognostic model we just published, now available for others in my organization to use.

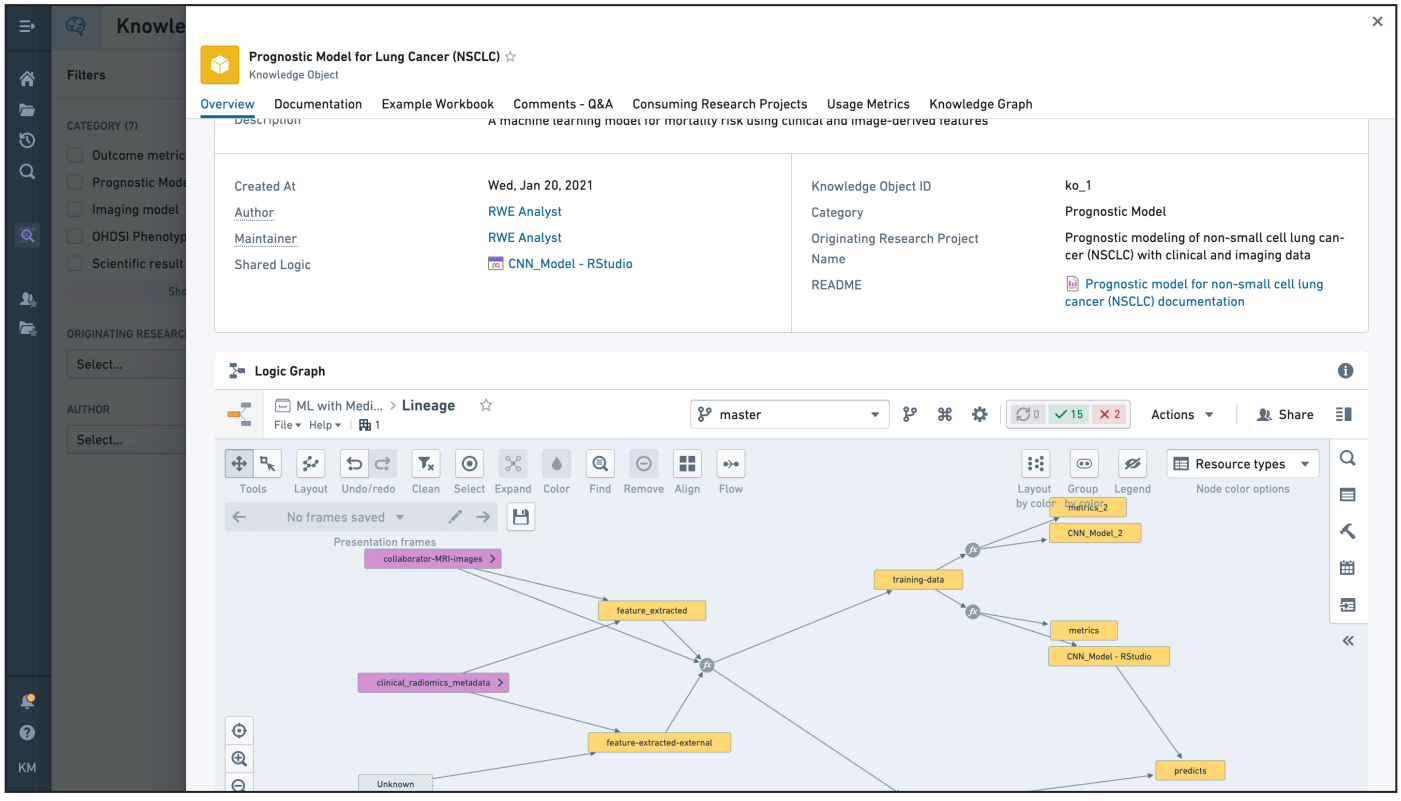

Reproducibility in scientific research is a serious crisis. High profile retractions of COVID observational research demonstrate the importance of getting this right.

Foundry helps address this issue with automatic provenance tracking. Any research result, whether a dataset, model or visualization, can be traced step-by-step back to the exact version of the raw data that was used to produce it.

In a similar manner, Foundry helps organizations track attribution. Whenever any artifact from the knowledge store is re-used, this usage is tracked and the originating author is acknowledged.

Life Sciences Cont.

The combination of robust security and automatic provenance gives our clients confidence to collaborate in new ways, from working with third party analysts to private public research initiatives.

Conclusion

And this is a good place to leave our demo. As we showed we think Foundry is unique in its ability to address some of the most complex parts of clinical research; data governance, results reproducibility and secure collaboration across studies. We help our partners accelerate from one study to the next, and ultimately translate research into improved patient outcomes.

National Institutes of Health

SPEAKER

Dr. Joni Rutter and Dr. Ken Gersing

Hello, I'm Joni Rutter the Acting Director of the National Center for Advancing Translational Sciences, or NCATS.

NCATS was created 10 years ago to remove costly and time-consuming bottlenecks in the translational research process. These bottlenecks may be scientific, they may be operational or they may be administrative.

Our driving hope is that we will bring more treatments to more people more quickly. To shorten this timeline, we focus on what is common across diseases and we take advantage of multiplier effects.

We take on translational research projects at any stage of preclinical or clinical development and we collaborate to advance them along the pipeline. Through our study of translational processes, we can overcome the bottlenecks and can get treatments in the hands of providers and their patients faster—that's translational science.

For COVID-19, researchers needed science to move faster than we have, faster than ever before. We needed to carry out basic, translational, and clinical research, as well as clinical trials, and implementation science all at once.

National Institutes of Health Cont.

A major speed bump along the way though, was accessing clinical health data – and making those data meaningful, open and accessible. Electronic health records, or EHRs, they're the largest source of clinical data. But in the US, we don't have a standard way to collect and manage those clinical data. Therefore, there is no standard way to use patient EHRs for research, to help make or inform public health decisions using near real-time data. And this is where the National COVID Cohort Collaborative, or N3C, comes in.

My name is Ken Gersing, I'm the Director of Informatics at NCATS' Division of Clinical Innovation.

One of the most amazing things about N3C, is you know we can talk about statistics we have we have 5 billion rows of data, we have a million COVID patient EHR records starting prior to COVID but what really is amazing about N3C is in the public health world, often data is splintered, and it exists in siloes all over the country. And what we've been able to do at N3C, really for the first time, is create a nationwide dataset that everyone has been willing to give so we can address this pandemic.

In science, you want to know the provenance of data, you want to know everybody who has touched it, every change that has happened to it, but reproducible science is hard because often the methods are just described as words. But in Palantir we can actually look at the keystroke of the actual thing that got us from A to B. And so we could just basically print it out or even more, have another computer pick it up and run the same thing, so reproducibility becomes the push of a button, where before it was almost impossible. It's indispensable what we are doing.

We often say the N3C wasn't shovel ready. It was ready. It's translational science in action and it's also inherently "NCATSian".

NCATS has been working on a solution to make health data sharable and had resources in place that could quickly pivot to build the N3C and its data enclave in a matter of weeks. The modularity and elasticity of the cloud enabled the environment for experimenting in rapid fashion. We could scale quickly when we needed to. So we leveraged that existing technology to have answers at our fingertips.

National Institutes of Health Cont.

When the crisis hit, we had many, 86 centers who were giving us their data. But they all speak different languages. We said to the centers, we can never get you to speak one language, you send us the language you talk in, and N3C will take every one of those languages and we will harmonize all of those languages and by using Palantir we were able to put it all into a single repository and have scientists look across these different languages but they don't even know they're doing it. From their point of view, it's one continuous dataset. It's really quite amazing.

It's all stored in a single repository that's secure, and very important because this is electronic health record data. The enclave is actually hosted at NCATS in our Gov Cloud instance and a Palantir instance where the community comes to use that dataset.

We have about a thousand investigators that are using the data every day. We clustered people together in what we call domain teams. A domain team is a group of investigators interested in a certain aspect of COVID, and they work together to solve problems. And they work together to solve problems. What's really wonderful is that discoveries that are made or tools that are developed or algorithms that are used are shared across all of the domain teams. We call these shareable objects "knowledge objects". If you think about it, it's almost like a grocery store.

You can grab a knowledge object off the shelf and reuse it. So the speed of the science is really accelerated because instead of every time reinventing the wheel, you pick the parts off the shelf that you need to use, you modify them as you want, but they're reusable. So we don't have to keep designing it each time.

Lots of people have large datasets. But if you make them into knowledge, or you make them usable to create knowledge, that's the secret sauce. That's what Palantir gives us that, that ease of use.

Today, the N3C is a secure, national resource of real-world electronic health record data from COVID-19 tested patients as well as controls, that is speeding COVID-19 research and ultimately improving patient care.

National Institutes of Health Cont.

We are working with students and researchers from all types of scientific disciplines and teaching them how to use the N3C Data Enclave to identify and answer critical research questions. The global pandemic created a heightened focus on this work, and we must now consider how to apply this approach to all diseases.

SECTION

Industrials

SPEAKER

Sam Woodward and Liam Mawe, Forward Deployed Engineers

Overview

Hi, I'm Sam Woodward. I'm a computer scientist and I work on IoT at Palantir. I've spent the last few years working across dozens of our Industrials customers and believe that our platform is capable of handling data in a way no other platform can. We've worked alongside our customers to understand the analytical computations they need so that we can build them directly into our world-class time series infrastructure. We've worked on the features that allow them to build and run models across the full history of their sensor data - without sacrificing performance. And we've built the tools to combine that sensor data with other data in the platform to fully contextualize their operations and scale models across the fleet.

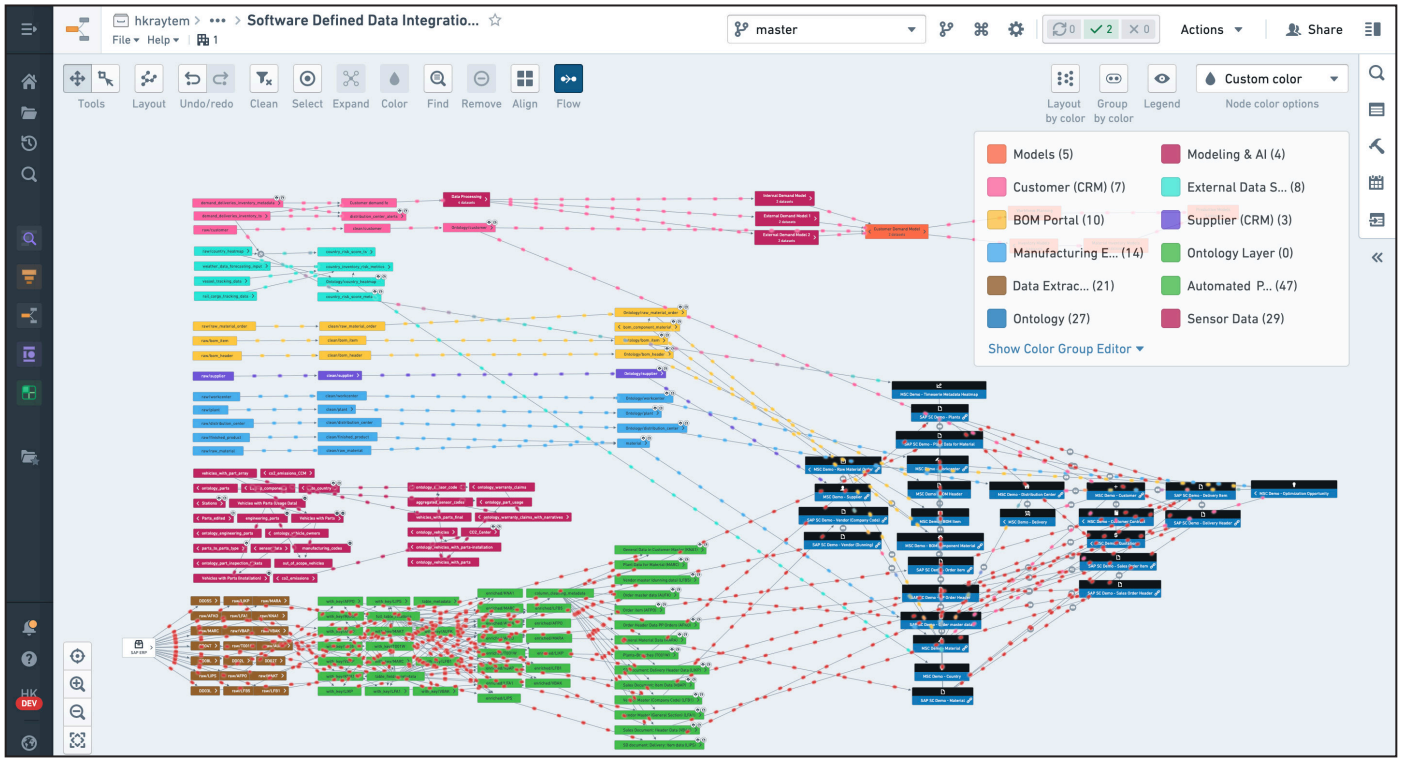

Industrials Cont.

Working side by side with companies like 3M, FCA and Airbus, we've actually uncovered some of the common challenges that these businesses face, both in terms of their data landscapes and business operations. Because we're a group of highly passionate engineers, we've spent years working on these problems.

We've actually been able to bake these hard-earned learnings back into the platform, so now you can apply these to each customer's specific context in a matter of weeks.

Part of this is what we call software-defined data integration, taking what used to take months or even years of painful manual work and you need specialist knowledge of the underlying data sources and actually replacing that with a few clicks.

- Understanding 1000s of underlying tables with non-intuitive column-names and mappings in systems like ERPs or CRMs
- Managing low-latency and data storage at petabyte scale
- Ensuring seamless write-back of data to underlying data sources meaning we can instantly deploy models into day-to-day workflows
- And key to all of that, ensuring you have the world's best data security, access controls and provenance baked into every corner of the software

Industrials Cont.

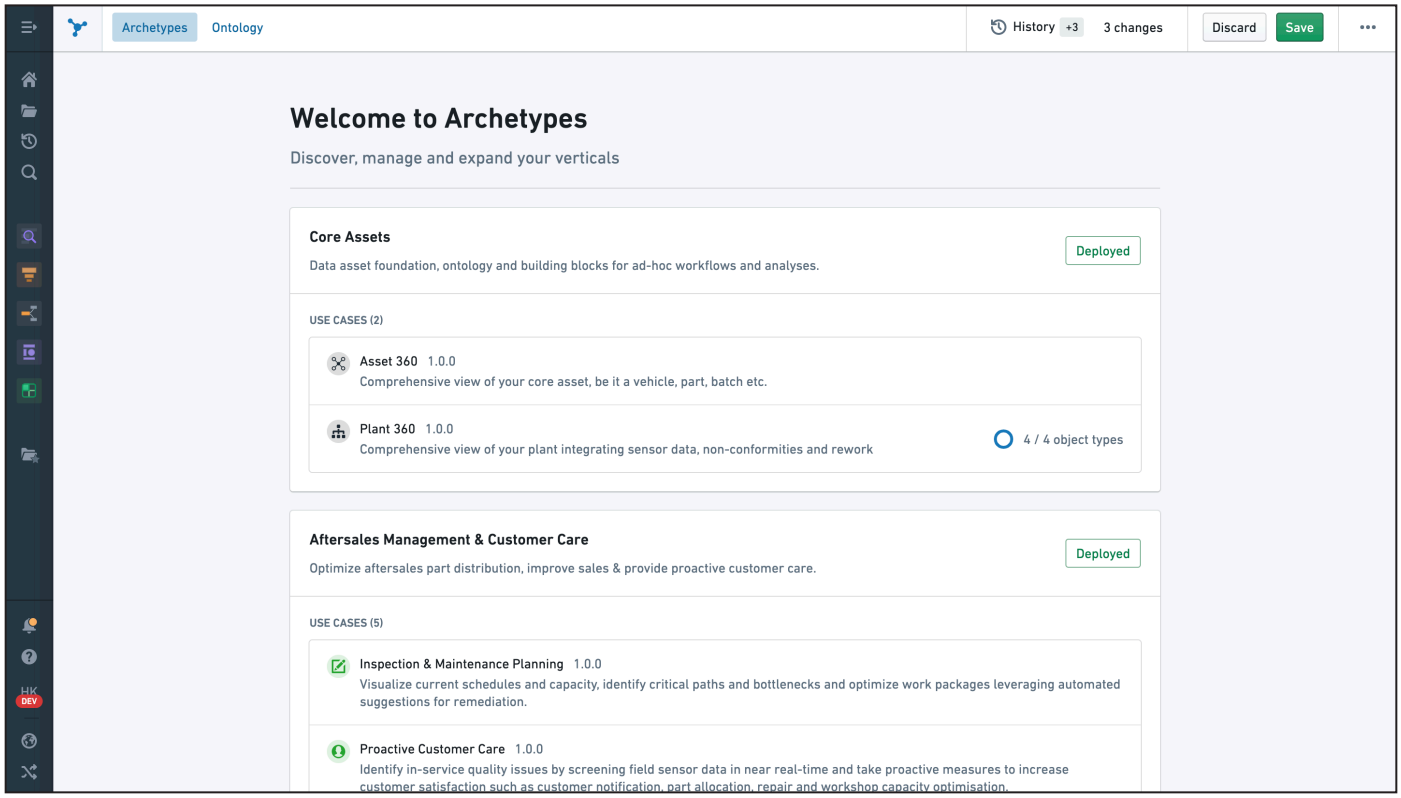

It's actually because we have built this technical foundation, that we can now apply that to business operations in weeks. For example, manufacturers can trace quality issues or better manage complex product portfolios, global logistics networks can optimize responses to things like trade regionalization, or energy companies can optimize production using massive scale sensor data from hundreds of assets across the globe.

Atop this technical foundation, we're excited to bring you our Industrials archetypes. Our catalog of Archetypes spans functions, ranging from production optimization to back office. Our customers can use these Archetypes to impact decisions across the business quickly, without rebuilding a solution from scratch for each emerging need.

Industrials Cont.

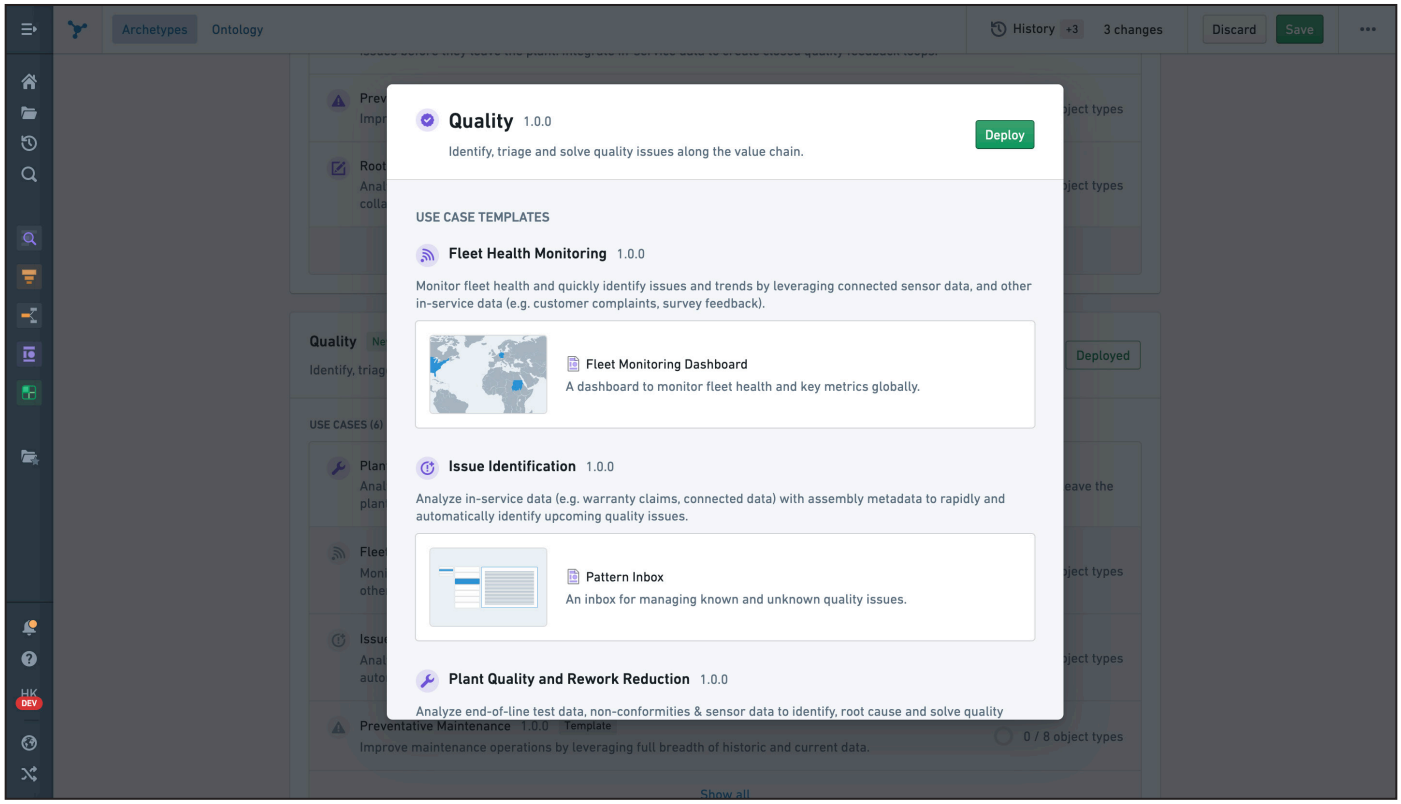

Today we're going to show you two Foundry Archetypes that can help transform operations — starting with the Quality Archetype.

Part 1 → Quality

Take battery manufacturing. Usage of lithium ion batteries is expected to continue climbing, and manufacturers face a huge challenge in increasing quality while continuing to reduce costs. Bathroom manufacturing is an extremely precise, multi-step process, and deviations can be introduced at any point. There have been countless point-based solutions available for years that, in theory, meaningfully improve this process; but in practice they have really struggled to deal with the breadth and flexibility you need to tackle here.

We're going to show you how we've cracked this code taking these really complex processes and delivering outcomes in a matter of weeks with the Quality Archetype.

Industrials Cont.

This Archetype will set up the necessary data integrations, generate the connections between data sources, and deploy the tooling to run relevant quality analyses. As one example, the Archetype will deploy our IoT data connector, which ingests high-scale, granular sensor data. This sensor data - whether it's batch data or streaming data - can then be used alongside other structured data sources for a full picture of an end-to-end process.

I'll start by taking a look at my global view of plant health, with any open production-related alerts. I can see that I've got one critical alert open, so I can inspect the details related to the alert to start to understand what's going on. To drill deeper, I can investigate my recent batches to identify where my deviations stem from.

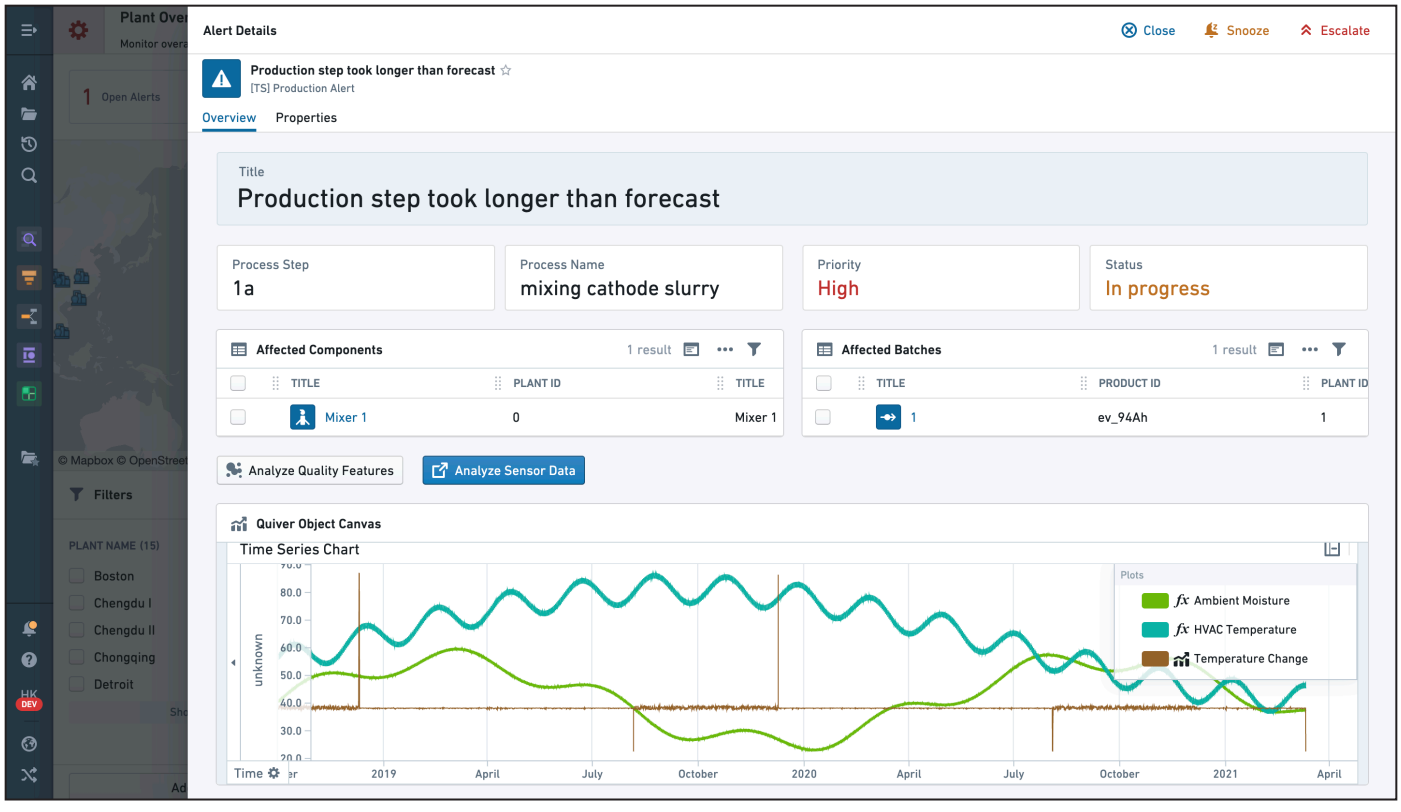

Industrials Cont.

I'll take a look at sensor data and features of our batches to compare the difference that led to batteries passing quality checks versus those that eventually had to be scrapped. I can generate a predefined view that plots relevant quality features for my failing step to understand where my batches differed early in the process. Once I determine where I'm seeing meaningful differences, I can add my batches from all time to highlight any patterns that could help me replicate more good batches.

There are several things that could be driving these deviations. To better understand whether I have a problem with my raw materials or my process equipment, I can drill into one production line in particular, and plot my successful batches over time. I can see that I've got periodic degradation, which may indicate an issue with the maintenance for this equipment. To further diagnose, I can choose to overlay some relevant information about my equipment, including recent maintenance data. The power of overlaying data from different systems with my quality and sensor data is game-changing in terms of developing a real understanding of my process. I know that my mixer requires regular work, but it appears that I'm suffering from an increase in rejection rates before I perform it. This may indicate that our model for how often to perform maintenance and repairs is inadequate for consistent quality.

To understand how to improve my model, I can use the sensor readings from the mixer itself. I can quickly pull up all available sensor readings and clean up my noisy sensors to prep my model inputs. I can then select a training interval and my target (my good batches) and quickly create a model that I can tweak in real-time to understand how I can better anticipate upcoming failures.

What's technically unique here is the ability to easily incorporate different data sources into my model and configure it in a point-and-click fashion to get visual feedback instantly.

My model is only useful if I can incorporate it into my everyday operations, so from here, I can export my model to refine and compare against my existing maintenance schedule and incorporate any of those changes into my plant operations.

Industrials Cont.

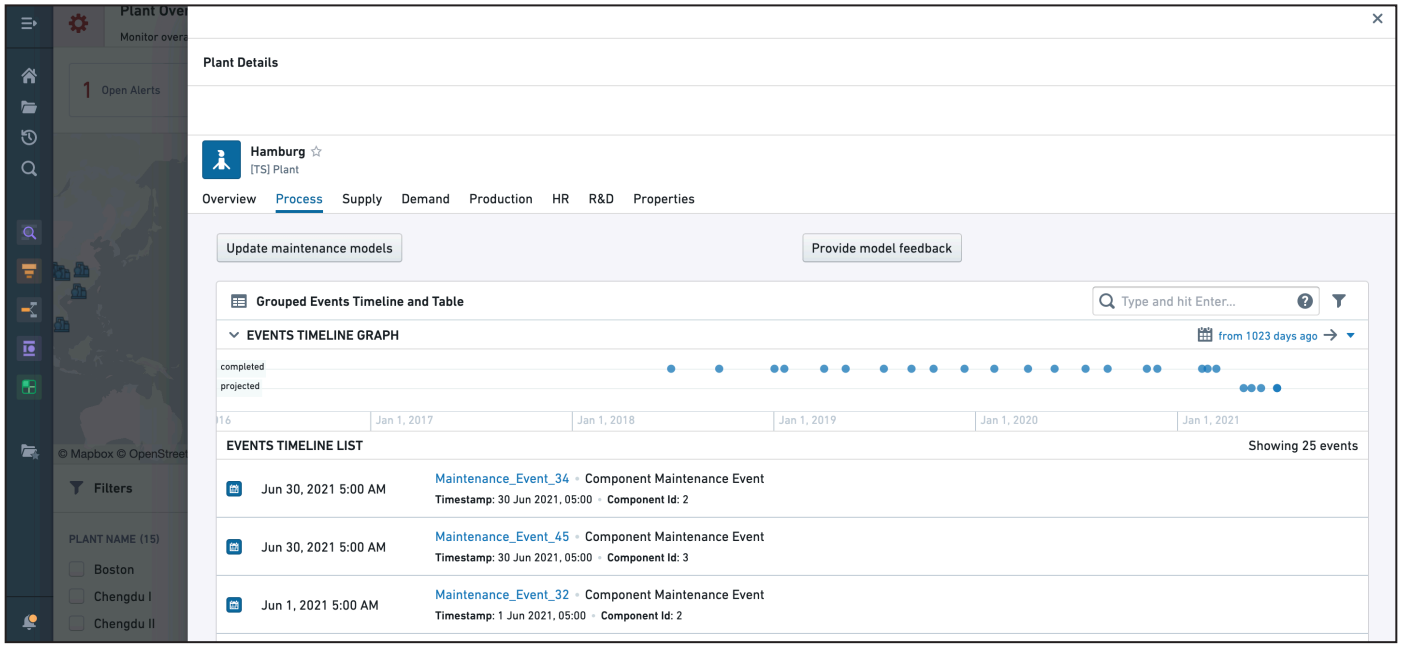

The Quality Archetype enables flexibility to tackle many different angles of quality improvement that hasn't previously existed in the space.

This seamless integration of maintenance, quality, and sensor data is what allows me to quickly understand the full context of my process and inform my next changes to operations. In this case, I can better estimate when I need to perform maintenance, auto-generate orders for parts and schedule it during a safe and cost-efficient period. But I can use the same techniques and tooling to also make sure I'm maintaining high levels of quality while tweaking process parameters, changing my raw materials, or increasing my throughput.

Part 2 → Supply Chain

To understand some of the other ways our customers use Archetypes, let's take a look at another critical Industrials function: the supply chain.

Supply chains are under constant threat of disruptions and keeping a supply chain healthy is a complex balancing act; increasing the on-time-in-full-performance in one part of the supply chain might cause operating costs to rise in another part.

Industrials Cont.

For years, other solutions have claimed that they can solve this complex balancing act. And whilst this may be the case when all the data and decision making is in one place, in practice, this is never true as supply chains run in silos, with fragmented data and functions.

The past year has clearly shown what happens when supply chains are pushed to their limits without the right solution - which is why when it mattered most, more customers than ever turned to Foundry.

As a Supply Chain Manager, I have to react to everyday disruptions as well as the bottlenecks and delays due to macro shocks such as COVID and the semiconductor shortage that stress assembly lines and after sales distribution. On top of responding to these incoming disruptions, I also want to understand how to be better prepared for future disruptions and challenges; rebalancing risk concentrations, preparing for the next product launch and accelerating sustainability initiatives.

Let's take a look at the Supply Chain Archetype; just like with the quality archetype, this will set-up the integrations and connections between data sources, and deploy operational applications that will help me run my Supply Chain decision making.

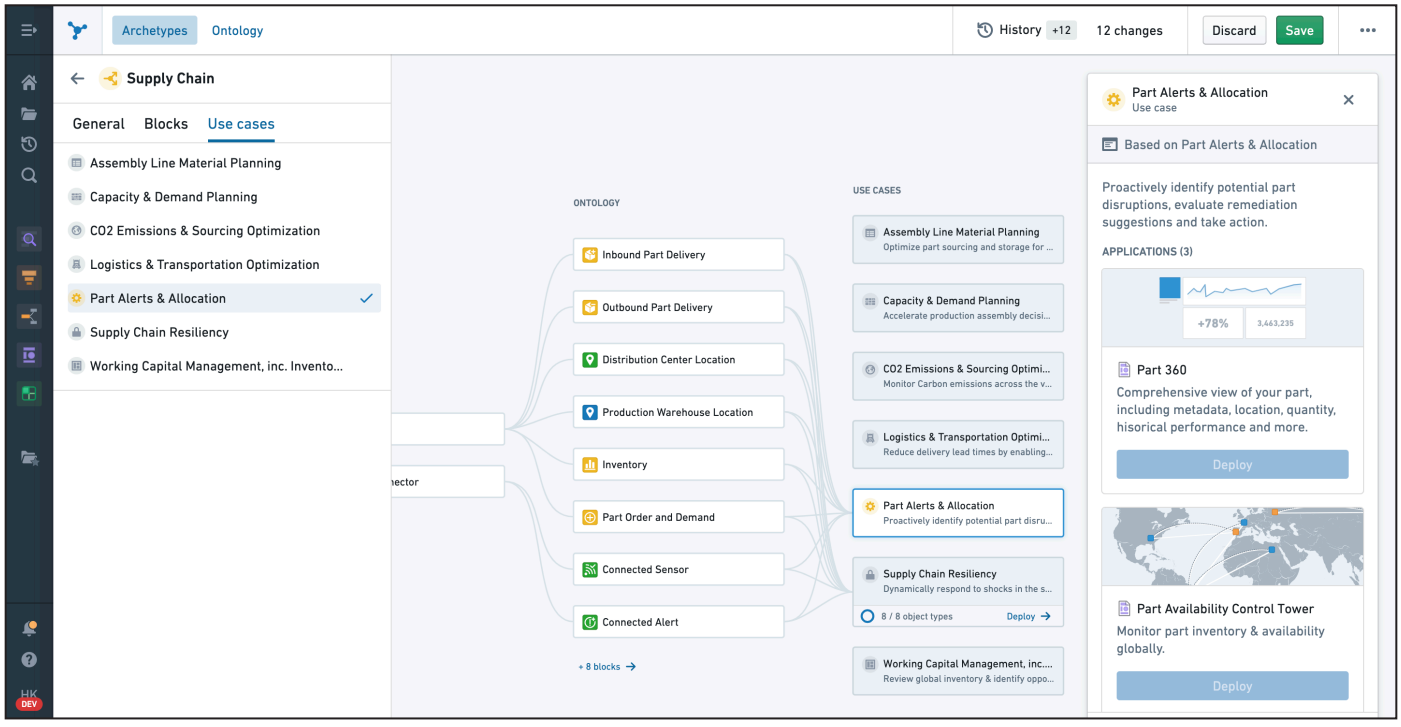

Industrials Cont.

This time let's highlight another example of software defined data integration through Palantir's ERP Suite.

Gathering data from an ERP system for specific workflows is manually intensive. This process that would take months or even years can now be done in a few hours with Palantir's ERP suite. You can instantly explore the Objects, models and workflows you want to build and once you've confirmed your selection the software automatically takes care of the rest - identifying which tables are needed, creating the syncs and required mappings, and building pipelines out of the box.

Now that I've deployed the Supply Chain Archetype, let's zoom in on just one small piece of the supply chain - managing spare parts for Electric Vehicle Batteries. Suppliers, supply parts that are held in warehouses across the country, to meet the demand of dealers, customers and plants.

Let's take this simplified view and show you what's actually happening: behind the scenes Foundry has taken that underlying complexity of the Supply Chain - interactivity between millions of different nodes - and built a digital-twin that models these interactions and allows us to see how changes affect the system as a whole.

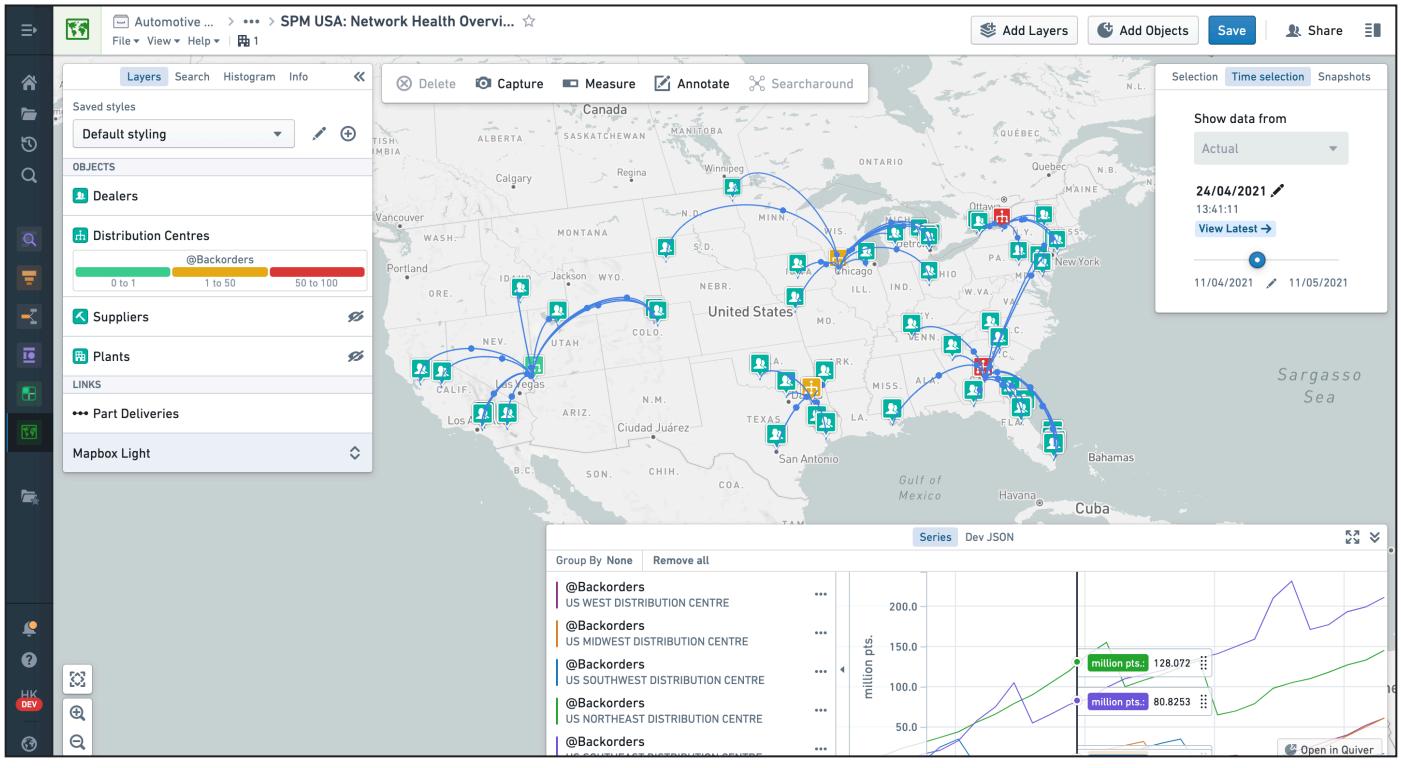

Industrials Cont.

Turning that impossible complexity into a manageable asset is the real differentiated aspect of Foundry, and I can now use it to drive intelligent decisions on an integrated and up-to-date understanding of supply, demand and logistics across my global network.

For example, as a Logistics Manager of Spare Parts I'm using the Supply Chain Archetype to forecast network health and respond to unexpected shocks to the system.

Our network is healthy today. But as time goes by some nodes you can see some of these nodes start to turn red, highlighting upcoming issues. I'm monitoring potential backorders of our spare parts but we can easily update the properties I want to monitor in just a few clicks: tracing costs, quality, on-time performance and CO2 emissions in one place.

Once I've defined my risks in this holistic view, a personalized Inbox gives me early-warning of potential issues I'll need to resolve, from responding to everyday disruptions such as low inventory or late deliveries, to actually flagging potential bottlenecks in the system.

Selecting the highest priority alert, I can see I have 11 days before the South-East Distribution Centre has a backorder of High Voltage Connectors - this is a critical component related to the battery of my Electric Vehicle line.

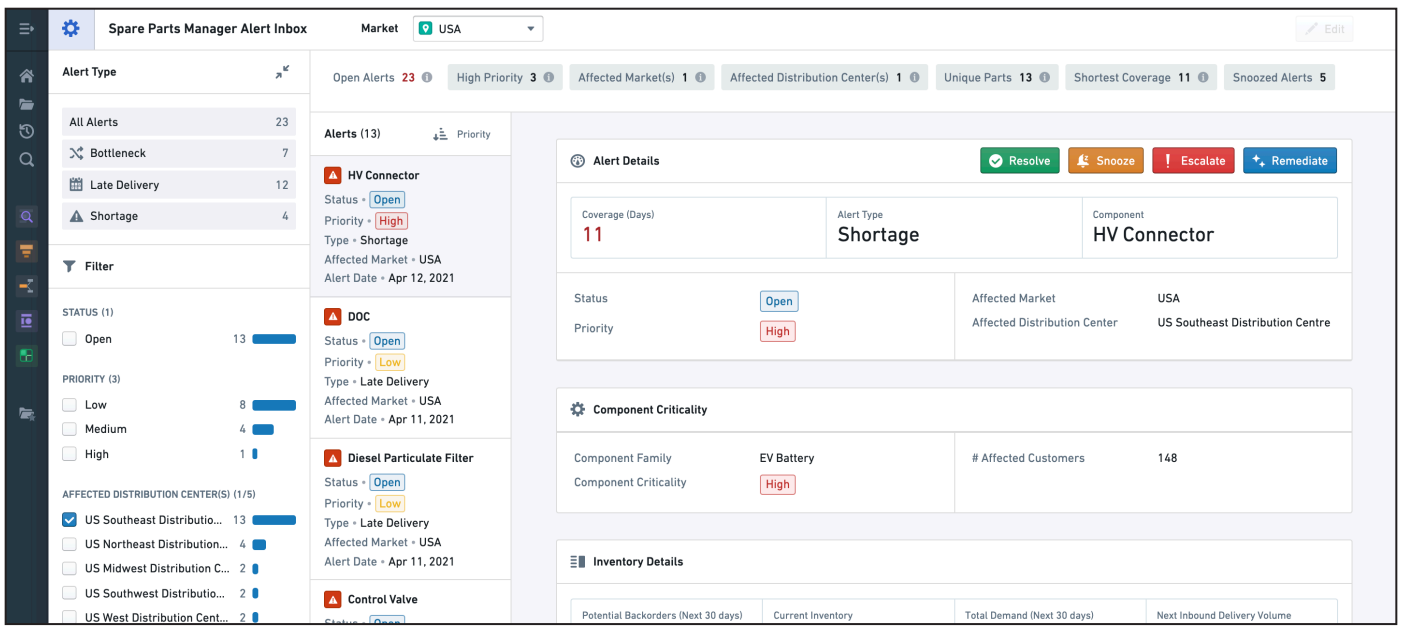

Industrials Cont.

If I don't react now, my backorder situation could quickly escalate, leaving sales orders unfulfilled and customers without working batteries.

Let's take a closer look at what's going on.

Integration with Advanced Shipping Notices show me the incoming Supplier Order has been delayed.

And I can see there are a few places causing an increase in upcoming demand, firstly real-time feedback from vehicle sensors analyzed alongside recent sales records and historical trends show me there's an increase in the expected regular customer demand, plus there are two open technical campaigns for this component raised by the Quality department.

What's critical here is that it's the combination of these different pieces in the network that's causing this disruption. Whereas previously, I would have to manually gather these data from across the business, now Foundry proactively connects the dots for me to help me identify issues sooner.

Now I have a holistic understanding of the situation, let's remediate the problem.

In one place I have visibility into all the actions I can take to resolve this backorder and at the same time a clear understanding of the impact this will have on my business. I can create new orders, modify existing orders or even scan through potentially hundreds of plants and warehouses to see if I can re-allocate my stock from within my network.

Industrials Cont.

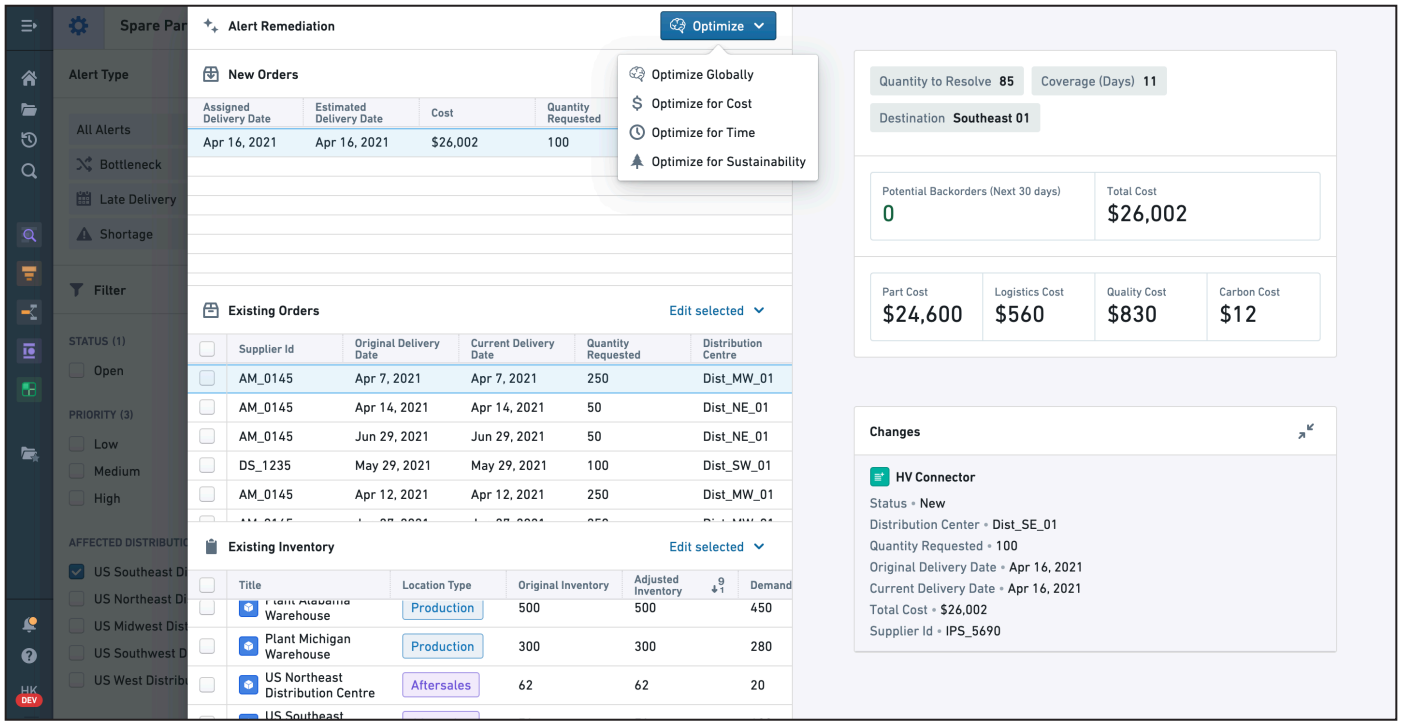

With every potential decision, I can instantly compare the estimated shipping cost, time to delivery, and expected quality and CO2 implications.

Previously it would have been impossible to respond with this level of understanding. But what's even more exciting is that instead of manually comparing a few scenarios, I can actually choose the exact set of decisions that will optimize for the factors I care about most - which will change, depending on the current circumstances.

While this might look simple, behind the scenes Foundry is comparing potentially 1000s of possible scenarios to help this Logistics Manager make a final decision on the best outcome - keeping business running smoothly and customers on the move.

Once I'm able to fully understand and react to today's challenges, the chained models underpinning the supply chain critically allow simulation of future events to rebalance trade-offs effectively - starting to shift from reaction to anticipation in the global supply chain.

Industrials Cont.

What happens if certain Suppliers go offline? What happens if I double my production of a given vehicle in the next 18 months? Or what happens if I were to change these inventory levels?

Let's simulate what will happen if suppliers in my network unexpectedly go offline.

I can review the potential impact this would have on my supply chain health, and also be given suggestions on how to optimize my network to take on future disruptions.

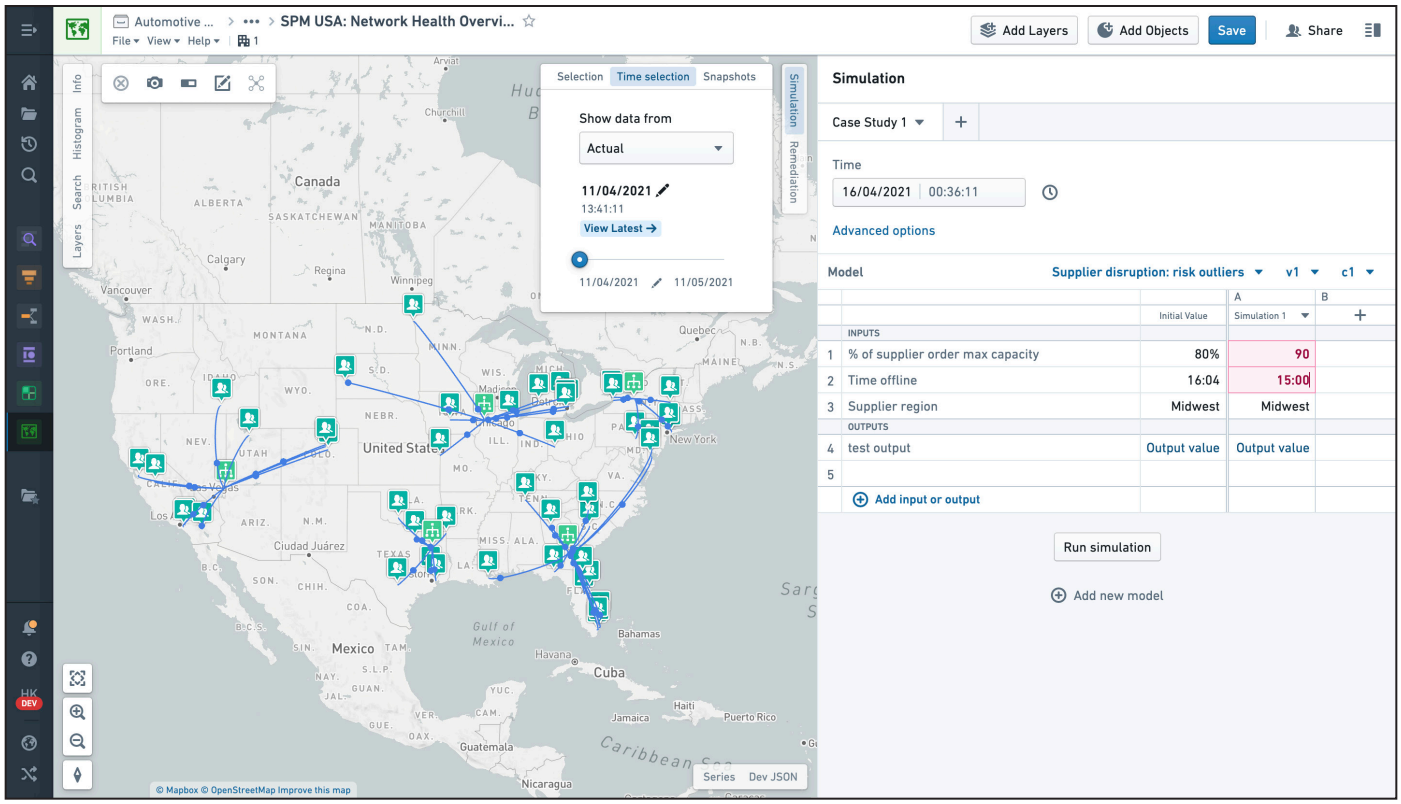

Here for example, I can see I have several critical dependencies that rely on a single supplier, so I should look at broadening my supply base to reduce my exposure. Supply Chain or Logistics Managers were previously forced to make these incredibly complex trade-offs with only a narrow view of the available context. But now with the Supply Chain Archetype, we can not only gain an accurate and granular view of today's risk, but can also simulate future scenarios and make adjustments to make sure our network is ready to take on the challenges of tomorrow.

Industrials Cont.

To give you an idea of the speed and impact of the Archetype, one of our Industrials customers had an incredibly complex supply chain data landscape with dozens of data sources - including 27 separate ERPs. With Foundry, in hours they had their first integrated view of the supply chain, within 2 days they were pro-actively alerting on potential bottle necks and in just 2 weeks had identified around $50 million in working capital, whilst simultaneously improving the robustness of the supply chain to react to future shocks.

Put simply, these archetypes can generate enormous businesses outcomes in just a matter of weeks.

Conclusion

While these archetypes can be deployed as individual building blocks, Foundry has been designed to be the connective tissue between these pieces - both within and outside of Foundry. Both the inputs and outputs of these different archetypes can seamlessly flow between each other - and any current investments outside of Foundry.

What we showed you today is just a peek into what we've been building, and we're excited to show more in the future. I've got to get back to de-bugging pipelines, so with that, I'll pass it back to Shyam.

SECTION

Conclusion

SPEAKER

Shyam Sankar, Chief Operating Officer

We are excited about the hard problems our customers are throwing at Foundry problems — problems they were unable to solve until Foundry. At N3C, the NIH is using Palantir to host the largest patient-level data asset of COVID-19 clinical trail data in the world. An effort like this usually takes years to put together. N3C put it together in weeks.

What you saw today is only a small slice of our Archetypes. We are so excited to show you more — a lot more — in the near future. Thank you so much for joining us.

Disclaimer

This presentation and the accompanying oral commentary include discussion of Palantir products, features and capabilities, including recent updates to our products, as well as potential product direction. They are intended for information purposes only and shall not be deemed to be incorporated into any contract or agreement and do not constitute a guarantee or warranty of any kind. They are not a commitment to deliver any material, code, or functionality, and should not be relied upon in making procurement, purchasing or investment decisions. The development, release, and timing of any features, capability, or functionality mentioned herein remains at our sole discretion.

This presentation and the accompanying oral commentary contain "forward-looking" statements within the meaning of the federal securities laws, and these statements involve substantial risks and uncertainties. All statements other than statements of historical fact could be deemed forward-looking, including, but not limited to, expectations of future operating results or financial performance, market size and growth opportunities, plans for future operations, competitive position, product development, technological capabilities, and strategic relationships, as well as assumptions relating to the foregoing. Forward-looking statements are inherently subject to risks and uncertainties, some of which cannot be predicted or quantified. In some cases, you can identify forward-looking statements by terminology such as "guidance," "expect," "anticipate," "should," "believe," "hope," "target," "project," "plan," "goals," "estimate," "potential," "predict," "may," "will," "might," "could," "intend," "shall," and variations of these terms or the negative of these terms and similar expressions. You should not put undue reliance on any forward-looking statements. Forward-looking statements should not be read as a guarantee of future performance or results and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all.

Disclaimer

Forward-looking statements are subject to a number of risks and uncertainties, many of which involve factors or circumstances that are beyond our control. Our actual results could differ materially from those stated or implied in forward-looking statements due to a number of factors, including but not limited to risks detailed in our filings with the Securities and Exchange Commission (the "SEC"), including in our annual report on Form 10-K for the fiscal year ended December 31, 2020 and other filings and reports that we may file from time to time with the SEC. You can locate these reports on our investor relations website (investors.palantir.com) or on the SEC's website (www.sec.gov). If the risks or uncertainties ever materialize or the assumptions prove incorrect, our results may differ materially from those expressed or implied by such forward-looking statements. Except as required by law, we assume no obligation and do not intend to update these forward-looking statements or to conform these statements to actual results or to changes in our expectations.

This presentation contains statistical data, estimates and forecasts that are based on independent industry publications or other publicly available information or our internal sources. This information involves many assumptions and limitations, and you are cautioned not to give undue weight to these estimates. We have not independently verified the accuracy or completeness of the data contained in these industry publications and other publicly available information. Accordingly, we make no representations as to the accuracy or completeness of that data nor do we undertake to update such data after the date of this presentation. All data shown in product demonstrations is notional or publicly available and any resemblance to actual persons, entities or events is purely coincidental and should not be inferred. Certain visualizations and capabilities shown in product demonstrations may rely on or reflect third party data sources or products that are not included as part of Palantir's standard product offering and may require separate third party licenses.

Disclaimer

This presentation may contain links to publicly-available websites, data, or other information. We have not independently verified the accuracy or completeness of such websites, data, or informa-tion and accordingly we make no representations as to their accuracy or completeness nor do we undertake to update such data or information after the date of this presentation. The inclusion of external links does not constitute endorsement by Palantir of the linked websites or the data or information contained therein.

By attending or receiving this presentation you acknowledge that you will be solely responsible for your own assessment of the market and our market position and that you will conduct your own analysis and be solely responsible for forming your own view of the potential future performance of our business.

Unless otherwise noted, all product, feature, or service names, logos, and trademarks, including without limitation Palantir and the Palantir logo are the intellectual property of Palantir and/or its affiliates in the United States and/or other jurisdictions. All third party product and company names are the property of their respective owners. Use of such names is for identification purposes only and does not imply any affiliation with or endorsement by such third-parties.