

Transcript for “The Health Catalyst Data Operating System (DOS™): Lessons Learned and Plans for the Next Evolution,” Webinar, September 30, 2020

Dale Sanders:

Hi, and thanks Brooke and hi friends out there in the audience. Thanks for joining us today. We'll do our best to use your time appropriately. Thanks so much for joining. Bryan, great to join with you on this. This is a fun topic, yeah, we were sharing in the preparation for this webinar that we originally had this webinar scheduled for February and then with COVID and everything else going on we kept delaying it and now it's, I think particularly coincidental that we scheduled this on the last full day of my employment and my relationship at Health Catalyst as CTO. So lots of symbolism here today. Thanks everybody.

All right, let's go. So the agenda is pretty simple today. I'll spend some time kind of looking back and reviewing the why behind DOS, the concept and the vision, and then Bryan's going to look forward. Talk about how far he and his team have carried the vision and implemented it and what's left and whether it's ever going to end or not. Hopefully it goes on for many, many years.

I wanted to talk about why we're giving this webinar today. In part it's the change of the guard. It's old to new, it's from me to Bryan and establishing his role at Health Catalyst in the industry. It's also just the recognition that DOS was a very ambitious vision and endeavor that we publicly launched less than four years ago. The first time we launched it and mentioned it publicly, it was January, 2017. So we just wanted to give a status update on that, and we made kind of a big deal of it in January, 2017. And so I think it's appropriate that we reflect and be as public about whether we've executed on that vision or not.

And of course, DOS is about Health Catalyst but it's also about, I think new concepts in data management, not just in healthcare, but in general. And in fact, I think it's a little bit of a different approach to computer science. And so I think this is a checkpoint to see whether that concept has been validated or not to some degree. I think there's still quite a bit of future ahead to determine whether the concept is valid, but let's see where we are right now. Just a disclaimer, today might sound a little bit salesy. Typically, in my webinars, I try to avoid any sales pitches. So apologies are offered in forgiveness would be appreciated. We will do our humanly best to give you the truth and then you can decide whether or not it's a sales pitch or not.

And for those of you who are interested in trivia you might Google that lithograph there on the right. It's kind of a fun reflection back on my life in music. Okay. So where does DOS fit in the tech stack? And my assertion is that we've done a really good job. I think sort of the renaissance of computer science is occurring right now. If you look at what the cloud has done for us, the infrastructure and the operating system layer, you look at cool operating systems like Android and iOS. You look at the incredible capability of fourth and fifth generation software languages. What's happening with GPU and things like that in the chip space and how smart we've become with user interfaces. That's all been just kind of amazing, right?

And we're truly, I think entering the renaissance of computer science finally, after many, many years of slogging it out in the mud. But it's still very difficult for and especially on complex data environments for software developers and data analysts to work with data. And it's especially the case for the incredibly complex world of healthcare and life science data. So the data operating system conceptually fits between the traditional operating system and application software with the intent being that it makes life easier for data first application development, analytics, and interoperability. So that's the fundamental concept.

And I might add too that even though this is sort of a Health Catalyst branded concept, the truth is it originated when I was at Northwestern as the CIO. So I was the CIO there with my good dear friend [Timsoff 00:04:58]. I also served as sort of the chief analytics officer, chief data officer, but I had to wear the CIO's hat for affordability. And so I frankly got tired of paying for separate infrastructures to handle all sorts of what amounts to the same management or different management and the same data.

So Northwestern back in 2005 is where I really started thinking about this notion of the data operating system even though I didn't call it that. As this slide indicates, applications will come and go, but data is here forever. And the notion is free your data and free your mind. There's a big opportunity to locally innovate at the application layer by leveraging the data that you already paid handsomely to collect. And if you look at all industries now, I would say at the exception of healthcare, software and data are becoming the differentiating brand in all other industries. So companies are leveraging software and data to innovate and become different than their competitors. But if you look at what we're doing in healthcare, it's all pretty much the same. Right? There isn't great variability in the software or the data from one healthcare system to the next.

So there's very little differentiation at the software and data layers still in healthcare. There's a few exceptions, but not many. In the old days when Intermountain, places Vanderbilt and Mass General had their own electronic health records there was great differentiation because they leveraged homegrown data warehouses and electronic health records, but that's not the case anymore. So I would encourage all of you that are in healthcare systems today, that by laying down an infrastructure like the data operating system, you have the ability to build applications and analytics on top of that, that can help you innovate locally and become a competitor that's head and shoulders above the rest of the market.

As the second bullet indicates here, I had frustrations with separate technology infrastructure for managing the same data. It was just a violation of common sense. It's expensive, and it leads to all sorts of inconsistencies in data management. And I would argue that there are essentially three missions for data. And those three missions for data are analytics and AI, data first

application development and interoperability. Data sharing on a transaction basis between partners in the ecosystem.

So that's essentially what we have been endeavoring to create is this Health Catalyst data operating system, as a single cloud-based, API based architecture with a common, consistent layer of data to support those three missions. So start reducing the cost and the expense of those three separate infrastructures that we've had in the past, peel the data out of those existing systems, repurpose that data for your local innovation. You can still keep those source systems like electronic health records, ERP, et cetera, but peel that data out and use it locally to differentiate yourself from your competitors is what I would suggest. Okay? So those were the concepts around the three visions of data.

The other thing that we believe is that the human health ecosystem is much broader than just healthcare encounter data and claims data, which is traditionally where we're stuck right now in healthcare. So Brian and I have been designing and looking at the data operating system with the notion of a 25 to 30 year life span, at least. Knowing that if we really want to understand the human at the center of this cartoon, we have to be capable of adjusting, analyzing and repurposing all of the different data sets represented in those clouds around the patient.

Our digital understanding of the patient requires that entire data ecosystem. So for the most part, if you looked at our clients today, we're still back down there in the lower left corner, but we're starting to make progress a little bit at a time in the rest of that ecosystem of data. And I would encourage all of you that are out there in healthcare systems that you peel those bubbles off of the cartoon, lay those down on a timeline, and that should be kind of your strategic data acquisition roadmap for the best care and the best treatment possible for the patient at the center.

Another thing that motivates us in the design and the development of the data operating system is the healthcare analytics adoption model. I published the first version of this sometime around probably 2002 or three, something like that, and then was more public about it, more formally using it in the market beginning of 2013, I think. We collaborate very closely with HIMSS, essentially gave them the rights to use this. And this now is used by HIMSS to help organizations benchmark themselves and measure their progress against this model. I'm in the process of publishing a new version of this, that addresses the direct patient analytics and artificial intelligence.

But anyway, if you look at what Brian and team and Eric, just in the act are doing, they're building out capabilities in each of these layers so that we can progressively move clients and the industry up these levels with our products, just as fast as we possibly can and as cheaply as we possibly can. This is the

marketecture slide. I won't go into much detail here just to show you kind of in general, how we portray this to the market. At the foundational level of what we offer at Health Catalyst is the data operating system with all of these components below and then on top of that are essentially three, maybe even four categories of applications.

We have what we call foundational software applications that cross many different use cases. We have domain specific software applications that address very specific workflows, very specific decision support. And then we have what we call tailored analytics accelerators, and those are kind of traditional BI dashboards and reports and things like that, that you would see and Tableau and click to that kind of thing. A fourth category that is not yet well represented is the third-party operating systems.

We've acquired a couple of companies over the last couple of years, we're plugging them in so that they don't have to move data around, but we can leverage their logic, their application layer on top of the platform. And we want to encourage clients and more and more vendors to do that because frankly it's quite expensive to lay down the data operating system, every single healthcare system. So economically we think it makes more interests of the industry, just a single layer of data to support multiple uses, including third-party application development. So that's the market texture slide.

I want to talk just a little bit about a paper that was published back in 2015, a lot of authors from Google and to a really interesting paper called the Hidden Technical Debt in Machine Learning Systems. And this statement on the right summarizes it and I'll read it. It is dangerous to think of these quick wins in machine learning and AI as coming for free. Using the software engineering framework of technical debt, we find it as common to incur massive ongoing maintenance costs in real world machine learning systems. I think this is one reason that AI and machine learning has kind of hit a flat spot because it's really costly and expensive to build out the infrastructure, to leverage the full potential of AI and machine learning.

And so, what we've tried to do, and we are endeavoring to do.... this is a diagram from that paper, which indicates the infrastructure around the machine learning code in the yellow circle. So relatively speaking, the machine learning and AI code is very small. This is the surrounding infrastructure required to extract the full potential of AI and machine learning. Okay. So this is our version of that infrastructure and here's our little box here around AI and data science. And this is the surrounding infrastructure required to leverage that. This is not easy to do. This is not easy to build on your own. This is not even easy for other vendors to replicate. It's taken us a lot of time and money to build this out.

And I can say as a CIO who had to build these systems before they were commercially available the cloud makes it feel it's cheaper and easier to build these kinds of platforms and it's true. It has made it a lot easier, but it's super expensive and hard to build these things and maintain them on your own. And I do see, I think a little bit of naivete in the market now because of the cloud. I see a lot of organizations going down to build their own platform. And I think they'll find out that it is just not economically sustainable.

Likewise, in the vendor space, there aren't very many vendors that have pulled this off as evidenced by the competitors we see in the market, which you could count on one hand, this is tough stuff to do. Okay. So in that January, 2017 lecture, we published the Seven Attributes of the Healthcare Data Operating System. So what we wanted to do is kind of hold this up against our own progress and grade ourselves, but also publish it to the industry so that you can use it to grade yourself or other vendors.

So those seven attributes are as follows, reusable clinical and business logic. So things like registries and value sets and other logic that lies on top of raw data that can be reused and updated through open APIs and enables third party application development. It has to support streaming data. So, near a real-time data streaming from source all the way to expression of the data on the right-hand side of the diagram, super important.

Data latency, right? Squeezing data latency out as much as we possibly can. It has to integrate structured and unstructured data. So texts and categorical industry data in the same environment, it has to be able to support images too. It has to support what we call closed loop capability, and that is the ability to ingest data, analyze data and then fold and close that analytic knowledge back on the production of that data in the workflow, especially in a clinical workflow. We'd like it to support a microservices architecture so that upgrades come in tiny little pieces that you don't even feel or know about rather than every fall 1000 hours and two months of painful upgrades to these.

Again, as a CIO, I had to live with that and we want to adopt microservices architectures support and I give Brian and his team great credit for making a lot of progress on this. It has to support machine learning. It has to natively run AI and machine learning models. That can't be an afterthought and it has to be embedded. So DevOps applied to machine learning and AI has to be a part of this, and it has to be somewhat agnostic as a data lake, meaning that the data lake has to stay fairly generic so that you can ingest and leverage just about any kind of application on top of it.

And it'll be interesting, Bryan and I can talk about how good we feel about grading ourselves with each of these, but given that it hasn't even been four years on a very ambitious endeavor, I feel pretty good. I'd give us an overall kind

of a B+ maybe on what we've achieved. And the good news is there's still room to achieve an A+ and that's where Bryan is headed in the future.

And to that point, one of my observations in life is that the greatest leaders, companies and organizations in history, exhibit chronic constructive dissatisfaction. And Bryan Hinton personifies that trait. And I'm just honored and thrilled and grateful that we've been partners for these many years and that Bryan now is going to take over the reins of this and make it even better than what we've got so far. So with that, Bryan, I will turn it over to you, my friend. Thank you so much. Brooke back to you to shift the presentation to Bryan.

Brooke MacCourtney: All right Bryan. You can present. You should be ready to go.

Bryan Hinton: Yep. Confirm that you can see that.

Brooke MacCourtney: Yep, we can see it.

Bryan Hinton: Okay. Well, Dale, thank you. It is interesting that it comes at this important and kind of watershed moment for us as a company. And I just want to thank you. It has been a personal and a professional privilege to learn from you to be a part of the team and have great love and appreciation for you and look forward to that personal relationship continuing even as our professional relationship changes. So thank you, Dale. I appreciate that.

Dale Sanders: Thank you friend.

Bryan Hinton: So as we pivot perhaps just a little bit and talk about kind of what we've learned over the last four years and then, and where we're headed, I think it's appropriate for a lot of reasons. To start with that healthcare analytics adoption model. So, we had the adoption model and Dale published that, and it provided that framework to communicate and measure the various levels of healthcare analytics. At the time, Health Catalyst had what we would call our late binding data platform, and we were using that in conjunction with clients very successfully and enabling success at level five.

Primarily, there were perhaps other levels, but level five was an area that we were finding great success with our customers. But when we launched DOS in early 2017, we did so as an inflection point, we knew and Dale described this very well, that if we truly wanted to be a catalyst to change how data is used in healthcare and to provide that massive measurable and data informed improvement, that is part of our mission that we needed to expand and rethink what we were trying to be as a data platform to more naturally address all levels of the analytics adoption model, and just really change in many ways the relationship with data that many people have.

So as we thought about that, and as we've approached that it truly involves creating a more hybrid data platform environment that could early bind some data elements into common data models that would standardize terminology which is needed to address the lower levels in the analytics adoption model. In addition to addressing those, by doing that, that work would provide an interoperability foundation that could be used for many other things, including fueling things like FHIR as the next wave of interoperability in the industry.

Now that same platform, however, to continue to address the higher levels in the adoption model needed to continue to be flexible and open and have a focus on, as Dale mentioned, advanced analytics, data science, big data and Dale hit on this concept of a data first application development platform. And that was a part of where we started. And you can see some of our forays in there as you think about level six around population health management, and some of the data first applications that we have endeavored to build have been in that space.

And we wanted to change the relationship applications have with data. Historically, if I wanted to be an application that had rich analytics, my options were really to try to extract much of the data into my own application repository and run things, and the cost and infrastructure needed to do that meant that applications cost more to build and were more difficult to build and maintain and sustain. And if we want the kind of innovative applications that we see fueling innovations in other industries, we have to lower the barrier to entry here. And so we envisioned DOS being that.

Now that's what we've been working on. So that kind of previews some of the work that we've done over the last few years, and along the way of doing that, we've continued to use DOS in practical settings, across healthcare with our customers, and it's yielded great improvements for our customers. We have a thousand documented improvements over the last three years with great economic benefit. And that's just the very obvious documented tip of the iceberg and our customers and clients are using DOS to drive many more day-to-day improvements in their businesses.

That on the ground experience has also allowed us to assess what components of our DOS vision are working well. What improvements did we need to make, what tweaks and features needed to be developed. And in some cases, we learned that some parts of that vision and strategy needed significant work, and we needed to pivot and learn from it.

And we showed that transparently and I think that's part of the process of trying to build something that is innovative and new and we have made those pivots and shifts and learned from them. But also true to our cultural attribute of continuous learning, we've worked in many different ways to continue to learn,

not just from our own practical experience. We appreciate the great work that industry analysts do to provide thought leadership here. And we've learned and evolved our vision as we worked with them.

On the screen, I have just some screenshots of a few of the reference architectures that exist out there from these organizations and as well as from vendors and we've studied them and discussed them, we've assessed them. And also tried to reach out and analyze where other organizations have concretely tried to implement these architectures. What's worked? What were the strengths? What were the weaknesses?

All the time with an eye for evaluating. What role? How did the data operating system enable these and what role did it play? So we could continue to evolve that vision. And we found some interesting things. So to provide the data infrastructure needed for that high fidelity picture of a patient, of population health, of an organization's financial analytics that we require a variety of different data technologies. And that's perhaps a no-brainer there, hopefully.

But that often includes and is validated as you look at these reference architectures, that includes systems for streaming, systems for machine learning, systems for traditional ETL, data curation and memory analytics. Those come with them, their own operational tools and techniques. And they come with their own deployment models, which causes an interesting challenge for data and analytics departments in healthcare and Dale alluded to this, is that people end up having to build the integrations.

And that comes with significant costs and challenge, frankly, to scale it. And as such organizations end up spending on sustaining and integrating rather than value creation for their business. And perhaps even more concerning is that's just to deployed the technical capability. That technical capability has limited true healthcare knowledge built in. It's just the foundational technical building blocks. And so as a result, the cost often comes with large sustain custom technology support costs or maintenance or development costs.

And when you talk about success in that context, because of the large cost, even if there is value derived in comparison to the cost, it often is a commiserate and we see ... Well, it's fairly well documented out there, talking about the analytic reboot that happens as people try to find their path through this. And so as we assess that and evaluate it an interesting dynamic happens. These tools are great. The tools that you would use to do this are great because they're open.

You can do many, many, many things with them, but we mentioned how they have little in the way of domain specific content or optimizations for the business domain, which is particularly important as you consider healthcare's messy data landscape. And so the burden falls on those organizations to how do

I deal with this? And so, as we stared at what DOS was trying to be, we said, "All right, with our documented history of results with improvements, can we both be open and integrate for people to accelerate their efforts here, but can we also introduce the optimized elements that they'll need to deliver from a healthcare perspective?"

Now, if you look at that space we often think about them, or we talk about them as closed or niche data systems. They often portray themselves as platforms when reality, what these closed systems have enabled is they're optimized for a specific set of use cases. They may provide the ability to ingest data from a subset of sources. Sometimes they force you to do that and provide it in some kind of a standard format. They might have a limited set of AI and machine learning capabilities and that fixed model enables certain types of data and analytics to occur, but it will never satisfy the breadth or flexibility required by the analytics adoption model. They're simply not open enough to be a general purpose platform.

And so as we go to build DOS, the challenge becomes to be open. So we address the spectrum of the analytics adoption model and to be optimized so that we can help people scale those investments. A third area that we've learned from is other industries. It's interesting as you look at companies that have transformed other businesses and industries and industry leaders is that one of their significant competitive advantages and Dale alluded to this as well, is that they know how to make the most out of their data.

And interestingly, not only do they know how to curate, mine and use that data to inform decisions, but they've structured their businesses to maximize the utility and impact of those insights. It's a fundamental part of their DNA in how they run their businesses. And that's an important part to call out because DOS is a technology solution to really empower data, to change your business, means that your business processes and employees have to learn how to change in response to that.

Now, fortunately, as we think about DOS and Health Catalyst customers change management has kind of been a core foundational principle for catalysts from the very beginning. So we know if DOS succeeds in providing a healthcare optimized and open data platform that we will be well suited to help our customers change and transform their people and their processes in accordance.

There's a quote I ran across a while ago that I love, that I think demonstrates the shift that organizations and individuals need to make. And it's the bottom part of this, where it says, "He uses statistics as a drunken man uses lamp posts for support rather than illumination." Brian Eliason, a colleague of mine here at Catalyst said this a little differently that resonated as well. He says, "Using data

to confirm or support is necessary, but not sufficient." And that's the challenge that we find in healthcare today. And frankly in many other modern businesses that we're using data to support rather than to really illuminate our decisions or drive our decisions.

And Jason Jones, our chief data scientist, he has some great insights. If you ever have the chance to hear him talk on this, I would strongly encourage it. It's just enlightening as he talks about our relationship with data and how we can use it to more effectively illuminate and point us at decisions rather than just confirm the decision we've already made.

The other element that we would acknowledge looking at this is they've learned, these industry leaders have learned how to use data, to transform their businesses quickly and to leverage data quickly and in a world where across the globe, we're facing the challenges relative to COVID-19 and specifically in healthcare quick, has been perhaps underscored in a way that we underappreciated in the past and how quickly we have to learn and adapt and integrate.

Now, it's been fascinating to dive into the ways that these various companies approach data success like, "Okay, how do we do it? And how do they do it?" And while they run at a scale and velocity that healthcare doesn't frankly have quite yet, in most cases, they also, while they have that complexity, they don't have the complexity that is represented in the competing standards, the inconsistent implementations that is the everyday reality of healthcare.

But even without the complexity of healthcare data, those other complexities with data has forced Netflix to invest significant, significant money and investment in both development and operations to develop a data operating system that works to run their business. That's simply a level of development and ongoing operational investment that is impossible for most organizations and certainly for most healthcare systems and providers and the like to maintain and sustain.

Another important insight from our learning and evolution, as we've used DOS and put it out there is to see the data landscape evolve. And simply put, data exists today, even beginning in the volumes that we deal with at Catalyst, at a volume or at a velocity and the variety of types of data that Dale shows in his wonderful cartoon of all the data, that it is impossible for a single data store or a single execution platform to meet all those demands.

Now, clearly bringing data together to integrate it and analyze it in a data warehouse, an enterprise data warehouse or a data lake, will have a vital long-term role in any part of a data platform. But increasingly, equally important is this data orchestration layer or fabric, or the data operating system layer that

connects these various data stores and compute platforms and allows the data to flow to where it's most needed.

As it flows the data it's cataloging it, it's mapping the dependencies and the relationships and the industry is starting to name this. And the name that often is found is just talking about it as a data mesh where you have that connective fabric together, that operating system, connecting the data stores in a very rich and meaningful way. Now, two as we looked at that landscape, and as we look at the variation and the complexity that begins to be required and think about it in the context of computer science and the broader software community, clearly things need to change in the data engineering community and world. And Dale once again, touched on this. Software engineering has been vibrant and evolving very richly over the last, certainly over the history of computing, but in the last 20 and 25 years, we continue to see great evolutions. Data engineering and data tools and platforms are not evolving at that speed.

We see pockets where you get a Hadoop rise and then a little bit of a Hadoop fall, but the level, the sustained level is not quite there yet. And many of the tools being used in data engineering look, and frankly act way too much like the tools that we've been using for the last 15 to 20 years. And that simply will not bring about the change with data that is needed. And so we think about this as applying software engineering principles to data engineering, and you can kind of see an interesting mapping where we've taken some of the concepts and showed them together.

And I might just make one comment on this that I think is interesting. So a few years ago, the software community was having a rich and robust discussion and conflict around the value of dynamic languages versus static type languages and performance improvements. And what's interesting to see how quickly they came to the fact that look static languages, static type languages, started adopting dynamic concepts. And dynamic languages started adopting static concepts and because they saw that it isn't an either or it's and you get that hybrid.

We went through the data space, went through its own and frankly, it's not quite cross the chasm on this argument, I don't think yet. Personally, between the arguments of schema on read and schema on write where schema on read was the solution to so many problems. And it actually is a very useful technique for more dynamic schema for data and just, it's vital actually. But as often is the case, an overcorrection can lead to the lead to disadvantages. And when you look on the heavy schema on read vendors what they started to realize is oh, for many scenarios that schema on write is advantageous, and we're seeing the schema on read vendors build out things to enable schema on write for a variety of other scenarios.

And it's a very similar corollary to that dynamic versus static language conversation that happened in software engineering many years ago. So there's things that we can pool and learn from and apply the data engineering. And that's what we're doing as we build DOS is how do we apply software engineering principles to data engineering?

Dale Sanders: Hey, Bryan I got interrupt you to compliment you on that slide. That is a super powerful slide. There's a whole white paper, there's a whole undergrad or graduate level course on that paper and how to move everything that we've learned about software engineering into data engineering. That was an awesome slide.

Bryan Hinton: Thanks, Dale. That concept, which frankly was my exposure to that I think first came very much from you is that, that's just super an important concept. So you will continue to see us innovate there on the DOS platform relative to that, for sure. So summarizing this kind of learning piece, I would just maybe reflect that there's a lot of things that we gathered in, and we found that we needed to organize them in a construct that was helpful. And the analytics adoption model, wasn't that perfect framework to do that in terms of what it was trying to accomplish in some of these learnings.

And so we've started to distill them into taking a cue from our healthcare success framework that Catalyst has in summarizing the learnings and what, we're right now just temporarily calling the data success framework, that defines foundational pillars that need to be present in your data and analytics infrastructure and strategy to generate insights, to govern those insights, to deliver those insights at scale. And then we categorize the various capabilities that we've learned and sussed out from all of these different disciplines and individuals and the like into a set of categories that will enable and empower successful enterprise data and analytics strategies and form the underpinnings and principles of what we want DOS to be.

And here you have it kind of more fleshed out with a variety of different pieces and parts, and an important part to capture here is that a framework like this hearkening back to what we talked about earlier, it lists a lot of technical capabilities, but those aren't sufficient. It has to include data and, or sorry, people and process components of the framework that work together to deliver success. Because one of the greatest mistakes made in our industry is applying technology as the only way to solve it without bringing along the people and processes.

We've also started applying this framework to the analytics adoption model and marrying those two approaches. So that as you think about a given analytics adoption model, what technical and people and process capabilities from the data success framework should I be investing in and do I need to focus on to be

successful at a given level of the adoption model? We've also used use it to score us. If you want to think about it from a DOS current capabilities perspective, both to help us see where we have gaps to guide our future investments so that we can continue to deliver that data success that we want our customers to have.

So from that comes the continued evolution of what DOS is and what we want DOS to become. And so running through that, that data orchestration layer or data fabric or what truly is kind of the heart, if you will, of the data operating system is our DOS data engine. And it's that connectivity tissue, the common data infrastructure needed for operating and deploying data pipelines, regardless of whether you're running, R or SQL or something else that it provides a common way to deploy and an operational console to manage so that you can integrate new technologies at scale, without significant increases to your operational costs.

In doing this, it's important to understand that DOS data engine is not replicating the functionality and features of the underlying systems that it's orchestrating. Those do their job very well. The DAS data engine's job is to manage the coordination between the various steps and optimize those steps. So a data form might look something like a SQL data curation step that's piped into our machine learning analysis, followed by a data extract to file that's sent out to some consumer. The data engine manages concurrency in that process.

It manages the dependency between those various steps and provide standard processing patterns to help accelerate the deployment of new pipelines and data flows and decrease the long-term maintenance of them. So let's take a look real quick at how that looks like. So as we break down the anatomy and the details in a data pipeline, what we find is there's a substantial amount of code that really is just infrastructure code represented there in those green boxes. It's things like incremental versus full loading strategies, logging, data persistence code, making sure that you're loading the data in such a way that minimizes business impact. Schema creation, index management goes on and on. That infrastructure, especially the fact that in data tools is not overly reusable often it quickly overwhelms the meaningful business logic that is really the value creation for the data engineer or analytic engineer is trying to accomplish.

Additionally, the cost of that infrastructure in each step in a pipeline becomes a disincentive to data engineers to write modular and maintainable and debuggable, if that is a word, code. And modular code has long been a recognized best practice in software engineering and software development. And so one of the ways that we have built and are building DOS is to remove the disincentives and writing more modular data pipelines so that your code becomes more maintainable over time. And the way that we do that is using a metadata driven approach where you have your business logic that you focus on

with a couple of other metadata elements that wraps in a logic package that we call a binding.

That enables DOS to do almost all of that infrastructure code for you, allowing the pipeline builders to focus on their business logic and sequence it together and build it out in steps that is maintainable and modular long-term. And as a result, they get the benefits of maintainability over time with significant code reduction. They don't have to think about that infrastructure code. It's been optimized and will continue to be optimized in code. And a name that's very popular or common in the software engineering industry around this is just declarative programming. It's programming at a higher level rather than imperative programming where you're being very explicit about all the pieces that you want to do.

Another element that that DOS does and is investing in is the aspect of dependency management. So here shows a mildly complex data flow that runs on top of DOS. And there are many more complex data flows than that. In traditional ETL tools and even some of the more modern database processing engines that exist out there, the data pipeline developer has to make these dependency connections. And that comes with a cost in changing the system, as it evolves over time as well. And DOS manages that for you. It automatically detects those dependencies, understand how they work together in a graph, in a directed acyclic graph and will execute them.

And so, as a result, it just takes that complexity off of the data engineer and allows them to focus on their specific business problem rather than the infrastructure of connecting them together. Additionally DOS, coming back to the diagram, in that desire to continue to be open, it supports a variety of execution environments and data stores, and we have plans to add more.

It enables you without to build your own data models and your own data flows. That's the foundational principle of DOS to allow you to interact with the data in a very rich way. While you're doing that DOS is cataloging the new data assets, tracking the lineage and doing those things for you but it's allowing you in a very open way to integrate with the data.

We also provide, and many of you probably know this or have heard this, but we provide many, many pre-mapped data sources that you can connect to, to get you started then to get data into DOS. And we're adding to that all the time. So we've covered some of the open parts of this. I want to pivot now as we wind down and get ready to open it for questions, to talk about the healthcare optimized portions of DOS, because that's equally important. Is DOS endeavors to enable both how do we provide healthcare optimized capabilities, and be open.

And in one of the other foundational principles is how do we allow that openness to also take advantage of that healthcare optimization? The first place we start is with our DOS Marts, which is our healthcare optimized data models. They come prebuilt with content to help you wire them up to standard data sources in the industry, but that is a huge lift and an important list in that standardization process that we talked about earlier. We also provide patient matching, which is a foundational part for data integration.

We have standard patient matching rule configurations that have been vetted out across many, many clients that you can apply and use quickly. However, keeping with that open philosophy, you can come and you can work with those and you can change those if you if you so desire and wish.

Now it's vitally important as we think about healthcare data to move past just standardization and talk about normalization. The key to do that is a robust, very rich terminology infrastructure that provides both the standards to normalize your data too, as well as the mapping infrastructure to ease, automate, and speed up that process. We certainly use these capabilities built into DOS to populate our own DOS Marts, but we also make them available for those custom data flows that open aspect of things that you can take advantage of that mapping terminology and knowledge that's been created to enhance your own data flows as well.

So you see the optimized aspects of the system interacting with the open aspects of the system. The platform as a whole allows us to then innovate in new and interesting ways at the healthcare knowledge level. So our Population Builder tool in our population platform allows us to engage a new set of users with data, people that are in the clinical domain or in the business domain can interact with Pop Builder and create and stratify populations, which are in turn used in products like our care management tool.

But that same healthcare content is readily available for custom data flows to be used. So you don't have to duplicate the definitions, pull it out and the dependency between all of those things are tracked. We're enabling the same capabilities at the measures level right now. That's one of the things that we're investing in heavily and building it out. Also, we're using those standard data models that come that's DOS Marts to build out the FHIR support. So that's something under active development right now.

And like all of the pieces that we've talked about, that flow into FHIR won't just be for those optimized data models that Catalyst ships, but will be a platform to allow you to integrate extensions, to that FHIR data model that you can add various data elements in using FHIR's extensibility mechanisms. So you get both of those things playing in there. And our DOS Marts model while heavily inspired and mapped to FHIR, are also being informed from other key and

important healthcare data models and standards, so that they can also be used to feed DOS using the same foundational capabilities that we're using to populate and serve up FHIR.

Now, we're super excited, in general, how DOS is coming together. It's a lot of work. We have stumbled, we've corrected and the momentum is building, and we're excited for DOS as an open and optimized data platform. There are many places that we could have covered today that we didn't. We could have talked about data quality. We could have talked about API level enhancements to make the app development experience better, but those we'll probably have to save for another day in another webinar.

But it's been great. Hopefully this has been helpful taking you through what we've learned, where we're going. We're excited about the opportunities that DOS will provide to healthcare in general. Dale, personally, it's been wonderful to present with you and wonderful to be inspired and led by you and thank you for all that you've done. And Brooke, I think probably I'm not sure how much time we got left, but a few minutes for questions, at least.

Brooke MacCourtney: Yeah, absolutely. If anyone has questions, please submit them now into the control panel and we'll do our best to get as many answered as we can. We do have one poll question I'm going to launch right now, as we jumped into the Q&A and while today's webinar is focused on the evolution of the Data Operating System, some of you might want to learn more about Health Catalyst products and our services. If you'd like to learn more, please answer this poll question.

And I'll go ahead and leave that open for a few minutes while we jump into the Q&A. It looks like we're getting lots of great questions. So we'll get through as many as we can. First question, and Bryan and Dale you can kind of both answer, whoever wants to jump in and answer feel free. First question says, what are some specific examples of what the future of DOS will look like?

Bryan Hinton: Well, sheesh. I hope we try to provide some of those, I guess. I may opt for a couple Dale, and you can chime in, but certainly I probably maybe provide three. One, as we build out that healthcare optimized content and enable FHIR and interoperability, DOS becomes a powerful interoperability platform for people to integrate FHIR. FHIR as a hope for standard would allow you to push this kind of information into any FHIR compliant workflow experience, more readily and easily. So I think there's a variety of different opportunities that light up there because of the flexibility that would come there.

The other scenario, and we talked about this a little bit, but I would maybe push in just a little bit is as that API layer builds out ... Right now, we have a little bit of a, I think Dale alluded this too. Stagnation in the application space because of

how hard data is to get. And so as we unlock that through the API's layers and allow applications to be developed more richly and robustly, I would hope to see the innovation with data start and the innovation and workflow applications be unleashed and unlocked in ways that we've yet to accomplish in healthcare.

The other element I was just going to mention is there is tons and tons of data to still go get. As good as we are today with the work that's going on to digitize with EMR, we're just at the tip of the iceberg. And so DOS by making it easy to consume and integrate data or easier to consume and integrate data, I would hope that that broader data ecosystem just becomes easier and easier for us to plug in and make sense out of and use. What would you add, Dale?

Dale Sanders: All of that friend, and also kind of reflecting on, gee, darn it, I wish we would have done more of this in the last few years and that's text and NLP processing. Pattern recognition and clustering sort of working at the app layer around clustering, but also getting our text data a little bit better processed, I'd say, and be more proactive with that. Few more curated data models, that more early binding curation of data. Eventually, inquiry by image content has to be a part of the platform.

Let me think here. I made a couple of other notes. Obviously, the streaming sensor space and making sense of that time series data, which is a very different kind of analytic environment than the kind of data that we typically engage in, in healthcare. So borrowing ideas from aerospace and automotive around streaming data and making sense of that streaming time series data to benefit patients. But there's lots of head for this, right friend? Yeah, plenty of roadmap-

Bryan Hinton: Totally.

Dale Sanders: Yeah. By the way, I can stay over as long as we've got reasonable attendance, Brooke. I can stay past the hour if Bryan can.

Brooke MacCourtney: Yeah, we've still have quite a few people on the line and so if you're able to stay on, we'll try and get to your question and if you're not able to stay on, we are recording this, you'll be able to listen to later on. So we'll just keep going through as many questions as we can. Next question says, "What challenges have DOS customers and potential customers experienced during COVID-19 that prevented them from purchasing more applications or the platform itself?"

Bryan Hinton: That one is huge in itself. I would say maybe perhaps trying to be brief, but address it would be ... COVID was challenging because of the underlying data foundation, the value sets what COVID actually was and how do we solve it? Raise just some foundational problems and so what was interesting about that was creating vendor created applications that could robustly survive in that dynamic data environment where things were changing so quickly. I think it was

difficult for all vendors. Everybody had kind of a COVID solution out there. How many of those actually were used is questionable I think. But what we did see with the platform was because it was open and flexible and we weren't constrained, people got in there and they rolled up their sleeves, which is pretty typical for our DOS customers.

And they developed solutions to enable COVID dashboards on their own. They used our idea app to do contact tracing because there wasn't appropriate tools for that. People were getting in and working on the analytic data to assess the financial impact when elective procedures were gone and what would that mean for the businesses as they came on the other side of COVID? So that openness was a huge advantage and in some cases it's tough for packaged applications to keep up with that, with just the dynamicism that existed in the COVID realm. So that's probably what I would say there.

Dale Sanders: Yeah. And I'm trying to read the underlying question here that Kyle is asking is maybe the economics of what prevented them from purchasing more apps or the platform itself. As Brian mentioned, the cool thing is the flexibility of the platform proved itself to the clients who have the apps and have DOS, we were able to pivot and support their reaction to COVID pretty darn easily really. And that's the good news. The bad news is, of course, the financial impact that it had on our clients and the enormous distraction that it had on them from just sort of standard operations and procurement, and strategy put a dent, I think in the ability to make sales during COVID but as evidenced by ... We've signed some pretty significant clients during this timeframe who stuck with their strategic plan around data. And Northwell is a good example of that.

Northwell, there's not a healthcare system that was more in the front lines in the COVID battle, the Northwell. But to their cultural credit and leadership credit, they stayed with their commitment to become a very data driven organization. And they recently signed a great agreement with us. So we appreciate that very much. But we anticipate everyone's going to realize just how important a flexible data platform is going forward. So we expect the interest in the platform apps to increase, not decrease going forward. I see the next question, Brooke. I got to ask this one. Can I ask this one to Bryan?

Brooke MacCourtney: Yeah, go for it.

Dale Sanders: Can you provide any information on the plans for Snowflake? That's from Robert Daniel.

Bryan Hinton: Yeah. And we actually have another one further down by B Murphy that asks the same thing. So we certainly are evaluating Snowflake. It is certainly the most interesting relational technology to come along in a long, long time, very transformative. So, we've been working with it and working with DOS on top of

it and experimenting with it. And so I don't think I represented my slide that we're yet fully committed to that. We're partnering with Snowflake in the evaluation, but it looks really compelling. I think one of the challenges that we have to figure out, and this becomes our important commitment to our customers as well, is how do we deploy it economically for you and make sure it's justified and Snowflake is very powerful. It is very expensive.

It especially depends on where you're at in your data journey. As your more sophisticated and more complex, the advantages that Snowflake bears out comes. We think will be more evident as people are beginning those journeys, how much is too much? That's what we're trying to work through in all transparency. But it would mean to the last question about Snowflake, it would mean that today many DOS implementations have SQL server as a core data store. It would mean that snowflake is maybe that core EDW data store as a replacement there.

Remembering that one of the foundational concepts is in a data mesh, the EDW is super important, but it is a node in the mesh. And so it serves a purpose. But as the universal data store for everything, I think it's unlikely as you consider IOT data and other types of notes data, there are other data stores that we'll likely want to use as part of that as well. So ...

Dale Sanders: If you go back to our criteria, right? For a data operating system, the seven criteria, the seven attributes I mentioned earlier, the last one was staying essentially agnostic to the underlying data lake or data content. [crosstalk 01:02:13] And so to Bryan and his team's credit, they've done a great job with that. So you could imagine Snowflake just kind of folding in underneath DOS with all the apps and logic sitting on top of it.

Bryan Hinton: Yeah, absolutely. And that is kind of the abstraction layer, if you will, Dale, that we're going for is how flexible can you make that for people to use. So ...

Dale Sanders: Back to you moderator for the questions.

Brooke MacCourtney: Okay. Next question is, do you envision the integration with low-code/no-code platforms like Appian and Mendix that would allow healthcare citizen developers to fully leverage data and drive a broader array of data driven applications?

Bryan Hinton: Yeah, certainly low-code/no-code platforms are super-hot right now. Everybody has one and wants one and/or just trying to use one. And I think there's a lot to like there. Something that we're paying close attention to. My answer would be, and this is perhaps a little opinion. So take it for that. But this isn't our first go around in the industry. I would argue with low-code and no-code that the technology stacks are better today and with the cloud, the scenarios are better.

So there's a lot of things that make this go around perhaps land and stay more foundationally.

So I think we're excited about it. We're evaluating it. I think they could serve a purpose understanding how they connect with healthcare data and how people would use them, I think is a, and how we think about consent and security, if you blow out the broader range, I think is an important element of how we would interact with low-code/no-code platforms that requires more work to be done before we empower the healthcare citizen developers.

Dale Sanders: I mean, I love the concept and I love the question Mark. So good for you thinking along these lines. The truth is our platforms are pretty specific to clients and healthcare systems, so it would kind of be up to them whether they allow citizen development or not, fundamentally. But we do have a product called Touchstone, we haven't talked much about that's a repository of curated data from across all of our participating clients.

And depending on how you count the denominator, there's at the low end, probably 50 million patients in that repository now, upper end, maybe 150 million. And I hope over time, and again, now that I'm leaving Health Catalyst, this'll be somebody else's decision but my vision was that we would someday allow citizen development around the Touchstone platform. Around de-identified data. Especially for patients who have rare diseases, right? So that you can, as a citizen access this national repository of de-identified data, and you can learn more about patients like you and if that meant building an application around a community of patients who all had a hemangiopericytoma, wouldn't that cool? That would be awesome. So I hope that that's what Health Catalyst has in the future.

Bryan Hinton: I would add maybe one other thing to that as we go back to the health systems side. I think there is a real opportunity inside the health system to expand people's interaction with data. We talked about that in terms of our tools like Pop Builder, where we want to engage more people in kind of their domain specific areas. I think as we talk about business process automation and some of the work that low-code/no-code platforms can help there, there's probably an opportunity for inside the enterprise, them to use those tools and interact with DOS, not as healthcare, citizen developers, but as driving efficiencies within their business enterprises. That's probably a little bit easier of a lift for most organizations to start with.

Dale Sanders: Yeah.

Brooke MacCourtney: All right. Our next question, how many provider practices are utilizing DOS?

Transcript for “The Health Catalyst Data Operating System (DOS™): Lessons Learned and Plans for the Next Evolution,” Webinar, September 30, 2020

Bryan Hinton: Yeah, I don't have a specific number on that, unfortunately. I think one of the areas of investment I'll just mention is how do you go get as many of the practices that are out there as possible? It's easier clearly when people are part of an affiliated group or standardized on an EMR, but there still is a lot of work to be done to bring in the spectrum of data that would represent a broad swath of provider practices that aren't totally affiliated together. And so we're actually spending a fair amount of time on that, figuring out how that might work for, in particular becomes interesting.

At least for us, for some of our client needs around measure submission and measurement to having that kind of full spectrum of data. And so the more that we can get too, and the more quickly we can integrate it, the better, but a specific number I don't have today.

Brooke MacCourtney: Okay.

Dale Sanders: I'm guessing this is somewhere around maybe 10% of our client base is provider practice sort of kind of independent of hospitals. Stan Procop asked that question. Having been a CIO for one of the largest, most forward thinking physician groups in the country, I can tell you firsthand that most physician groups aren't that interested in investing in the kind of capability that DOS provides. They've always kind of considered IT investments as a bit of an overhead as opposed to a differentiator which I think is unfortunate because platforms products like DOS can actually help those provider practices be a lot better financially, can help physicians feel more autonomous and empowered with data.

So it'll be interesting. We do have a lot of physician practices that subscribe to the our health information exchange and unable health, our quality measures platform. That's numbered in probably the hundreds, but DOS is a pretty small number. So push that agenda Stan.

Brooke MacCourtney: So, we're about 10 minutes past the hour. We still got 110 people on, do you two want to keep going? What are ...

Bryan Hinton: I'm good to keep going and answer the questions. So ...

Brooke MacCourtney: Okay. [crosstalk 01:09:12] All right. Let's do another one. All right. How do you expect payers, employers, and providers to all cooperate with the competing end products of such powerful tools?

Dale Sanders: Jeez.

Bryan Hinton: [crosstalk 01:09:28]

Transcript for “The Health Catalyst Data Operating System (DOS™): Lessons Learned and Plans for the Next Evolution,” Webinar, September 30, 2020

Dale Sanders: Yeah, it's hard, Robert. I mean, this is a fundamental challenge in the healthcare culture, right? Economically there's almost no incentive right now for those payers, employers and providers to cooperate much. We see it and we do a little bit of it but unless there's vertical integration among those three, it's all pretty much on goodwill and goodwill doesn't last very long and economically.

Bryan Hinton: And I would just add to that. I think we try to help whoever they are. That's our customer in that mix to be the best that they can be. And I think that's where we feel empowered as a vendor and as a service provider. We also, I think it would, and I'm not the expert on this, but we engage in the legislative space to help make sure that the incentives are as positive as possible for interoperability and support that way. I think it's tough like you said Dale. But trying to engage in that and encourage and nudge the market to head towards a better place is some of the things we're trying to do.

Brooke MacCourtney: All right.

Dale Sanders: Yeah.

Brooke MacCourtney: Okay. Eric asks, "Do you already have some European hospitals that you work with? What is your timeline for addressing this large opportunity?"

Dale Sanders: We have a couple in the UK. Guys in St. Thomas being one of those. It's not been an easy market to get into, frankly. I think generally speaking Europe doesn't want to spend their money on US business anymore, right? They want to keep their money at home at the macroeconomic level. But we're certainly open to it and we've proven the ability to work in that space both technically as well as culturally. Frankly, we've been so busy in the American market. Also, it's you always have to question whether you want to be distracted by time zone differences, cultural differences, that kind of thing.

Bryan Hinton: We are ... Yeah, go ahead Dale, sorry.

Dale Sanders: No, go ahead friend.

Bryan Hinton: So we do have international work going on in Asia and in the Middle East. That's where the focus has been from an international investment perspective. Ironically, that's probably a little more challenging from a time zone perspective in Europe to be perfectly honest, but that's where Health Catalyst has been investing their time Dale, I think answer to the European space. Certainly, there will come a time, hopefully that that's a place that we go and we help and can help address the large opportunity there for sure.

Dale Sanders: Yeah. And Eric Schumacher asked that question. He asked a subsequent question down there below, Germany, all in caps with a question mark and I

have huge respect for the German healthcare system. And what we largely advocate in healthcare is great care at a lower cost, but economically the great care to lower costs doesn't mean a lot sometimes to a fee for service environment, the US. But a healthcare system like Germany has that economic incentive is there.

Higher quality, lower cost is fundamentally an economic motive. So not lots of line-up there.

Bryan Hinton: Yeah.

Brooke MacCourtney: All right. I'm going to jump down to this question from Kevin that says, "What does DOS mean for the future of the Medicity HIE product?"

Bryan Hinton: I don't see that one. Where's that one? Oh. So- [crosstalk 01:13:48] Kevin, I don't recognize your name, but I'm still learning. The Medicity HIE products all report into my organization and we're on an integration path where we're DOS empowering them. That HIE business is an important part of the fabric of healthcare. And so we're taking a careful approach for how we integrate it together, but we would anticipate that that interoperability capabilities that are in Medicity are integrated with DOS and they become one platform. That's how we talk about it, is integrated together as one platform.

And so that's what we're marching to, and we have some milestones coming up at the end of the year that will ship kind of meaningful pieces to some of our, well, to the first kind of set of those HIE clients and it will just continue to roll forward that way and accelerate. So we're excited for those opportunities.

Dale Sanders: Yeah. In general, where the HIE space needs to move and is moving is from a message-oriented architecture to a services-oriented architecture, to API. And so Bryan's team is all over that, it's way, way over doing healthcare, right? I came into healthcare, what? In 1997 And we were already building services-oriented architectures and APIs in other industries in the mid '90s. And now the ability to do that is just orders of magnitude easier than it was in the '90s.

So we're way overdue in healthcare to move away from message-oriented architectures to APIs and that's fundamentally what you'll see with our HIE strategies, moving those concepts into the DOS platform. So you don't have a separate HIE infrastructure over here to maintain. You're running it all through the data layer of DOS through APIs. Listen to me talk like I'm still going to be around.

Bryan Hinton: What next?

Transcript for “The Health Catalyst Data Operating System (DOS™): Lessons Learned and Plans for the Next Evolution,” Webinar, September 30, 2020

Brooke MacCourtney: Let's see. Still a handful of questions in here. Has FHIR helped in the standardization of data models or has it been purely to enable interoperability for third party apps?

Bryan Hinton: Clearly, it's both. Well, it certainly has helped inspire data models, I think. And so from that perspective, I guess FHIR's helpful. Although, and my perspective would be this as, data models are totally necessary, but we've had data models in healthcare for a long time and they're vital and they have their uses, but they don't solve all problems. And so FHIR, I think will help drive greater standardization to common models. And so if you looked across vendor models, historically, those could look very different. I think increasingly you'll see them look more similar.

Which I guess is good. Where the rubber meets the road there and we've seen this as we've experimented with different FHIR implementations, by various vendors and understanding what other customers are trying to do is that unless you're committed to the normalization part of FHIR, you still miss on the overall interoperability value.

As you think about, if I'm getting a, just a simple thing, if I'm getting a lab result and it's not a lab result in a code set that works for me from interoperability perspective. And while the standard might discourage that, or even say that that's not an option, we all know that the implementation of the standard is very different from the standard itself. And so we still are working through that from a FHIR perspective, at least from what we have seen. And so I like that it's helped the standardization of data models, and it certainly is helping to have a more modern data exchange format for third party apps.

There's still work to be done at the core inside the data that FHIR ships around and the commitment that we as an industry need to have if we truly want to be interoperable and for FHIR to realize the promise that it holds.

Dale Sanders: There's a related question about why does it take so long? And you kind of touched on it. At the highest level, the EMR really motivated FHIR. Were not developed with the notion of open APIs, technically or culturally, right? Technically, those systems were never designed with the notion of open API for third party application development, technically.

So to their credit, they're trying to pivot that older software architecture towards something that can support the notion of APIs and FHIR. So it's technically difficult. Then you have every single EHR is deployed a little differently. So it doesn't mean that FHIR makes that easier to simplify truthfully.

Bryan Hinton: Yeah. And I think, oh, sorry, go ahead Dale.

Transcript for “The Health Catalyst Data Operating System (DOS™): Lessons Learned and Plans for the Next Evolution,” Webinar, September 30, 2020

Dale Sanders: Well, just the last thing is that just culturally and contractually, those EHR vendors grew up with the notion of a closed system. They were economically de-motivated for many years to open application development. So there's a technical, and there's a cultural shift that has to be overcome in order to realize the potential of FHIR.

Bryan Hinton: Yeah. And maybe just adding onto that with Eric was also, I think there was an element of SMART on FHIR there that I would just mention that at least from the analytic perspective and data perspective, that Catalyst takes, as we think about the SMART on FHIR aspects. SMART on FHIR is great from the authentication and security protocol and how do you embed? And then how do I get back? And I ask FHIR from a transactional perspective about people, but when you talk about analytic applications, leveraging SMART on FHIR the bigger challenge comes I usually don't need a single person's data. I need lots of people's data to provide additional insights and stuff.

And FHIR in general until bulk FHIR comes along and maybe bulk FHIR is that maybe it's not, but bulk FHIR to actually get data out of the EMR into an open compute platform where you can actually do something with it is where greater insight really comes and to use FHIR as a transactional data extract mechanism is very difficult, especially at high volume ...

Dale Sanders: Yeah.

Bryan Hinton: ... which is another problem. So ...

Dale Sanders: There's couple of questions about DOS implementation. Can we address those from the second to the bottom there? Mark Templeton asks, "Are there fundamental innovations that are needed to make DOS adoption faster and easier e.g. modular versus full-line implementation?" Then there's another one that Eric Schumacher asks about the process a hospital or a hospital group goes through to adopt it. Bryan, you ought to take that friend, take those.

Bryan Hinton: Yeah, I would say that's an area of ... Some of it is packaging innovation for clients to buy, I suppose. And how do they buy? DOS at its smallest can be very small. And so it's very easy to implement. And then what do you layer on top of it? So that modular piece is there and in terms of the ship cycle for the different components and splitting it apart, that part there's been great innovation there. Some of it is like, "Okay, how do we sell you an open platform that you can do anything with?"

But you want to start small or whatever. How do we balance that out, is something we've had to work through from a packaging, a licensing perspective and a balanced perspective. And so I think we've made some real progress in that. And actually would anticipate that we will see some meaningful shifts in

how people adopt DOS and they can do it very lightweight for very specific areas.

And then as they're ready to grow, as their organization's ready to grow, it grows with them, which has been something that we haven't been historically the way that Catalyst has done things.

Dale Sanders:

Let me throw another thought in there too. Sometimes there's a missing persona in our client base to lead the full exploitation and utilization of DOS. In other words the executive team might buy into the notion of digital health and analytically driven healthcare and all the things that you can do with DOS and the executive team is excited about that.

And they buy into it and they get some value out of it. But the truth is organizations need to have someone this kind of a mix of a domain expert in healthcare with a technical expert, with a data expert, that all needs to come together in a persona that's going to drive the full utilization and value proposition for DOS. And of course, going back to my chronic constructive dissatisfaction, I think there's a lot of value, a lot of fruit left on the vine that's not being harvested because a lot of our clients are missing that persona to drive and innovate around DOS' capabilities.

So if there are any clients out there or client prospects out there, I would just ... It's a call to arms really. Someone needs to take on that persona.

Bryan Hinton:

Maybe one other thing I would add there to Eric's question that just talks about the process. There's a contracting process about what problems do you want to solve. Catalyst, fundamentally, the way that we think about how we sell is we want you to focus on outcomes, not technology implementation. And so the process we're doing that. Once you've identified what the scope is of what outcomes you want to achieve implementing DOS is simply, it's cloud-based, we implement it, and we're partnered deeply with Microsoft. So we use Azure, we put it in the specific Azure region that makes the most sense. In the US it's mostly deal ... How do we get it closest to you? So the network ops are close. If it's international then it's a question of what's the data software question to play, but we can place it wherever and we can turn that on and turn it around very quickly in terms of the infrastructure stood up.

Generally, there's a VPN that's established because of the amount of data. And I'm not sure if Eric was technical or business process, but that's the technical process. We'd wire up a VPN, and then we're up and running, pulling data in and beginning to work on the platform. And then the steps after that have to do with, "Okay, whatever outcome or outcomes you're attempting to target, what is the need?" In some cases, those require deep data integration into the Dos Marts. Well, that's the next thing that we go tackle when they don't ... If it

Transcript for “The Health Catalyst Data Operating System (DOS™): Lessons Learned and Plans for the Next Evolution,” Webinar, September 30, 2020

doesn't require that, it can be more of a very specific value base. We're getting the source data end that you need, and we're working on the specific business problem that you're trying to solve.

Brooke MacCourtney: All right.

Dale Sanders: I want to put a plug in while I'm thinking of it for something I'm very passionate about. Speaking of adoption. We have a patient safety surveillance module that I think could, well I know, I don't think, I know could dramatically improve the safety of care that we provide to patients depending on which study you believe medical errors could be as high as the third or fourth leading cause of death in the US but it's largely under reported and it's largely under appreciated. And that's one thing that I think every one of our clients should be adopting as a fundamental part of the DOS utilization value proposition.

And I'm passionate about that for a lot of reasons, but also it just, in an enormously valuable, capable product that and it's a topic that I think healthcare executives don't spend enough time worrying about patient safety, fundamentally do no harm first.