



CREATING AND MANAGING NEW KNOWLEDGE CONTENT

VIDEO TRANSCRIPT

Hi everyone and welcome to our second KM webinar of our Espresso series: Creating and managing new knowledge content. Just a few logistics before we go ahead and get started. Everyone has been muted to eliminate any background noise and if you have any questions during the presentation, feel free to type them into the Q&A box on your screen and we will cover those questions at the end. And now I like to introduce our speaker Paul Nelson. Paul was an early pioneer in the field of text retrieval and has worked on search engines for over 30 years. He was the architect and inventor of retrieval Ware now owned by Microsoft Corp. Paul served as the chief architect at Search Technologies until their acquisition by Accenture in July of 2017 and he's now the innovation lead with Accenture Applied Intelligence where he continues to provide architectural oversight, design technology, research and training. And now I'll pass to you Paul.

Thanks Susanne.

You know I think knowledge management and knowledge are such fuzzy subjects. You know we know that knowledge is so important for everything that we do. We know that we can't solve any problem at all with knowledge. And you know, the knowledge of the corporation is just critical in order for the corporation to do its job. But you know it's just so hard to manage. This is kind of an abstract thing where we talk about managing knowledge. It's like, I don't know, managing air or managing water or something like that. How do you manage something like that? So you know I just wanted to try and create a

presentation which talks a little bit about how we think about knowledge and how we go about understanding what knowledge is and how it can be managed and mapping knowledge flows. Then beyond that, how we can actually improve it and apply some artificial intelligence to really make it better. You know when we talk about what is knowledge management, we often use this quote from Gartner which formalizes the management of intellectual assets and really about the creation and capture of an organization's use of those intellectual assets to help people and organizations operate more efficiently. But even with this, it's like well, what really is knowledge management? So what I'd like to do is I'd like to get super practical about what is knowledge management. Let's talk specifics, personas and get just really really practical about that. And so you know only think about why is knowledge management so hard.

You know these are what I would say are the top 5 challenges of knowledge management when I talk to customers. One, knowledge is really everywhere. It's dispersed over the entire organization over many systems and silos. How do you determine what knowledge needs to be curated? What knowledge comes out about organically from the organization? How to gather it? How to organize it? How to curated? That's part of what knowledge management is really trying to solve. And whenever we talk about knowledge management, often I hear customers jumping immediately to the taxonomies, classifications, ontologies, you know, the idea of managing knowledge means classifying it, structuring it, putting into in the higher keys. And when is it appropriate, when is it not appropriate.



And then context and linkage. These are, I would think, the unsung heroes of knowledge management. How to make the knowledge contextual? How to link it to other objects in the business so that it becomes so much more powerful when it's connected directly to the business. And then finally how to justify, measure, and prioritize knowledge management initiatives? How do you know that your knowledge management is actually having a positive impact on the organization? How do you measure it, how do you prioritize, what functions to do? Now these are the top 5 challenges. But in this presentation, we're really just going to focus on number 2 which is talking about curated versus organic, and how to understand a knowledge flow, and when knowledge might need to be curated, and what processes. And I'd like to do that by talking about some very specific knowledge management scenarios. So these are scenarios that we see inside of companies all the time. I think I have sort of brought up 3 scenarios and they're all kind of very old. So you know I kind of doubt I'm going to be able to get through all 3 in just presentation, so you might need to contact us. I can go through the other in more detail but we'll see how it goes. Ok the 1st one and you know I think this is like the entry point for knowledge management for a lot of customers, is customer support. And even very small companies like my own when I started my search engine company, we had customer support and we had knowledge management for customer support. This is where I think knowledge of the organization is so acute because you need to communicate your products and how to use your products to your customers. So the knowledge of your product and your service, and how customers deal with that is where the knowledge management really comes together. So your customers have questions and so you feel like one of the first things you need to do is hire customer support representative, a c.s.r. And that c.s.r. is going to receive questions from customers through a variety of channels like chatting in email or phone calls and then provide some answers. Now we have this knowledge transfer between the what's in the brain or the customer support representative who presumably has gone through some training in this.

Talk to some people and then how that customer support representative is providing those answers to the customers. Ok so that all works great. Now the Customer Support Representatives, you know, they're judged based on how quickly they can answer questions and so what are they going to do? They create the ticket database, of course every transaction goes to a ticket database, so we know exactly all the questions that are coming in. We can judge the volume of questions, the kinds of questions we have, and we have some details about the questions that are in the answers that go back and forth to the questioners so we can sort of track that knowledge. That's very important because that's a great place where we can mine for knowledge. Anyway, so the customer support representatives, you know, they're judged based on how quickly they can answer questions right. So one of the first things they're going to do is create a Frequently Asked Questions list just for themselves. So they get the same question over and over, and when they get that question they find it in their f.a.q. they can just cut and paste it from the f.a.q. into their email and send it back to the customer. Problem solved and now they're just more efficient. So this is great. Now we get a bunch more c.s.r. and the other c.s.r. say "Hey you have a really nice f a q's database. I want it so now we have sort of knowledge sharing within the customer support community." And the c.s.r. is like "Ok fine, I will give you my f.a.q. But I want your questions as well." And now they start sharing on the f a q's list and now I have not just one person updating a document f a q's. I have a whole team of people working on f a q's. Ok somebody looks at this and they say "You know, what if we just asked our users to search for the f a q's or you know just scan through the F.A.Q.'s themselves." Well then you know they wouldn't have to call the support representative at all and we save all this money because we don't have to hire as many c.s.r. Now the customers are sort of getting the answers on their own.



Ok that's great but then someone else says "But that's the f a q's that the customer support representative originally wrote. You know they're just not for public consumption. They're full of acronyms; they're full of jargon. You know that they talk about things and tools which the customers don't have. So really we need to stick a tech writer who's going to take the f.a.q. and sanitize it for customer presentations and then probably a public content." Ok this is a really nice little knowledge management flow. Between the customers, the Customer Support Representative, other. C.s.r.s., tech writers, and the public. This is all working great until of course the volume increases. Now we get a lot more customers. We get a lot more tickets. We get a lot more f a q's. We start getting more complex F.A.Q.'s like Knowledge Base articles which actually give step by step instructions for how to do things. The more you start selling products, more varieties of products, you know the whole thing just starts to get more and more complicated. And now we've got like, 100 different f.a.q. documents, maybe a 1000 different Knowledge Base articles. Things start to get bad and now all the c.s.r. are like "You know this is really really getting complicated. I can't find anything." We're just you know getting overburdened by the volume of content. So one of the 1st things we can do is we can stick in a search engine. So search engine really helps break down barriers and makes it easy to find past tickets. So I can answer questions, make it easier to find f a q's, and then make it easier for both the public and the c.s.r.s. to search for their internal F.A.Q.'s as well as for the externally sanitized f.a.q. And this is good. And this is you know mostly what I see lots of customers and this is as far as they get right. But let's think about can we reinvent or refactor this whole process just to make it more efficient and cleaner? So let's first store all the f a q's into an actual document management system and not just some hard drive someplace so now it's a little bit better managed.

We have some traceability; we have some old versions. If somebody makes a change which turns out to be wrong, we can fix it. We can recover the old version people aren't accidentally deleting f a q's all those sorts of things and the document management system we can stick in the approval workflow. So the tech writers know what to do. They're not getting lost as to which f.a.q. needs work. The public content approvers can just get notified when they need to approve something. So that whole process works a lot more smoothly you know and so they're happy right now. Let's put a search engine on top of that so that the c.s.r. can actually search the document management system and the ticket database. So now here we're kind of breaking down the barriers instead of having to search for the ticket database and the 2nd one for the document management system. Now we have just a single search box and now the c.s.r.s. can search for everything from one search box. So it's a lot nicer and so now they're happy and so the next thing of course is we want to take those documents, those sanitized documents, directly from the dock and measurement system to the search interface for the customers. So immediately after the document is approved, it's put online in a search engine. Customers can search for it right away. Their data is a lot cleaner so our fresh old f.a.q. can be sunset or rewritten or notified for the tech writer to rewrite. So now the customers have better, cleaner, more accurate content. It's also probably better categorized as well. And so now they're happy. So this is all great and this is like a nice clean management system.

Let's see if we can make it even better. So one of the things we can do is we can perhaps remove one of these links. That is instead of c.s.r. having to go directly to the document management system to upload new content, perhaps they can contribute it through the same interface that they do with search. That's one less user interface they have to deal with and that lowers the barrier for them to provide new content. Right now they can do search. They search for stuff and they don't find it so then they say "Ok there really ought to be an f.a.q. for this" so they create the f.a.q. right inside the search interface and contribute.



Now they don't have to go anywhere else. It's like Wikipedia. You search for something and we if you don't find it, it says "Hey would you like to create an article about it?" You can create that article right there and then it gets added directly and then it goes through the whole regular approval process. And now all the other c.s.r.s. have access to it immediately. Then after it's gone through the tech writer in the public content, the customers will have access to it as well. So it just lowers the barrier, makes it easier to contribute. Let's do the same thing for customers. Let's gather comments from the customers so we can start creating their own f a q's or their own comments. Then start commenting on this so that we're crowdsourcing a little bit of our content. So this is cool now can we make this even better. What about a robot monitors the ticket database, looks for hot and new topics, things that customers are talking about.

For example, maybe there's a global pandemic and now how people use your products is changing very dramatically. What sorts of things are important to them are going to be changing very dramatically. If we can monitor those changes to identify hot topics or new topics or interesting topics in real time, then the robot can say to the tech writer "These are some trending topics. Maybe you should investigate those." And now the tech writer can look at those tickets and talk to the c.s.r.s. They can create F.A.Q.'s on their own without having to wait for the c.s.r. to create them and then contribute that to the document management and start answering those topics to the customers right away. We have one more input. We can take the same thing from customer comments and start looking at those 4 things that are hot topics. The f.a.q. is for those right away and now we have a nice system which is generating and monitoring content and looking for knowledge. And then you know, refining that knowledge for both public and internal consumption.

So this is nice but of course let's look at this in the context of the larger organization. We don't just have one support. Well we also have 2 support and 3 support right now with some documentation writers and we have new product discussions and tickets, and we have product marketing with the official product brochures and spec sheets and things like that. And we approach corporate marketing which talks about the company and sometimes you have questions about that as well. So now you know what's next. Well all of this is of interest to people. Lots of people want to search for product documentation. Supports may be of interest to one c.s.r. as you know, it may provide interesting new hot topics. And then Product Marketing may be of interest to customers as is corporate marketing. And now you know new topics may be of interest to documentation writers and product development. It starts to get quite complicated in terms of all the different kinds of places where people might be interested in topics and how knowledge might flow through this organization. aAnd we're just going to leave that as an exercise for the student. Now starting to interact with the whole rest of the organization. Ok so what does this tell us? Scale is critically important right. One f.a.q. is easier to manage than a 1000 f.a.q. is. At some point you will need to refactor to handle scale. Who needs to share what, what needs to be filtered, what needs to be written. Can they read original content responsibility? Who is responsible for creating what content? What is collaborative, what is not? Where is it stored? How is it managed? Can you learn from your content? Can you identify what content is needed? Listening to your content and listening to what people are saying is so helpful. Then adding automation to identify new topics breaks down barriers and then of course you know search engines. They're still important. Search engine is by far the best way to break down barriers, surface content, make content active, and enrich it, and you know, make it helpful for people. So I think that's really where we're going to skip this next scenario because we're running out of time.



It's about materials, test reports, and how we can manage them. Then skip directly to what I think is sort of the North Star or the kind of the gold destination for a lot knowledge management systems. It's how do we handle sales? How do you respond to requests for proposals? As you know, one of the most difficult parts of knowledge management and I think very poorly done in most cases. So you know we start with the customer. They send in an RFP to a salesperson. That salesperson creates a proposal but in order to do that they need input from engineering. They then put from business and pricing. They need input from marketing. They need input from subject matter experts. And then they bundle all that together, send the proposal back to the customer, and we win. That's great. Ok so now we get another customer and we have to do the whole thing again.

The teams get bigger and bigger you sell more products. You have more and more customers. And you know while we have 500000 people that are in this world, so this is a very large, very complex process. Write more and more proposals, more and more pieces, more and more content things just start to become overwhelming with so much content and so many pieces. Somebody says "Well let's organize. What if we stick all the proposals in a document management system? Maybe we'll have a proposal workflow." You know this work for other knowledge management systems, it should work for proposals right. And I never really liked this process because proposals as we all know, they come in on a Tuesday and they're due on a Friday. And you have to do it very quickly. If you stick in some heavyweight process like this, you know the heavyweight set of tools that have to manage the content in this way. It's just not agile enough and people start doing the proposals on their desktops or their laptops. And you're back where you started and I'm not sure if that works

Ok so what if we go to the other side and say "Ok we have a bunch of snippets like pricing templates, proposal templates, use case write ups, product features. Then we'll do approval workflow or that are not the proposals itself, but you know the components that go into the proposal. And even then you know what I've seen this done and in real life, these things get out of date super quick. You create a pricing template or proposal template. It works for a couple customers. It doesn't get updated. It becomes out of date very quickly. People stop using it and then you're back to where you were. It turns out to be quite expensive to have somebody assigned to keep these things all up to date all the time unless you're the kind of company that has a relatively large volume of customers with a very large volume of proposals that are very similar in that scenario. I could see this totally working but for most of us, we have a much more diverse customer base with a wide and diverse set of needs. And so we tend to have to create proposals that are a little more customized. So these templated approaches they just don't quite work well. Ok now we're stuck with this mess the first thing we're going to do is we're going to try and organize around a collaboration system. So this is like Microsoft Teams and so you know we're doing this at Accenture. We have a separate Team for every proposal and now everybody can coordinate in that one Team site and talk within that site and contribute to that site. And then collaboratively work on presentations and pricing spreadsheets and everything that all goes into one place. At least those materials for that opportunity are put together right. So then what we're do we go from here? Well maybe we should start asking questions. Of course, we could put a search engine on top of it.



That's helpful but I'm not sure if it solves the problem necessarily. Ok so let's start asking questions. What do people really need when they're writing proposals? What kinds of questions do you know? And by customer here, I mean the people writing the proposals. What they really need to know? Well they need to know if it's very practical things. I think like what have we done for this client previously. Can I get intelligence about the client? What sorts of proposals worked and which ones didn't? Who worked on the proposals but for this client before? Who are trying proposals for this subject before? Can I find previously written paragraphs or slides? Maybe I don't want the whole proposal but I'd like a slide just about cloud regulation or maybe I want to slide on outsourcing or a slide on credentials. You know just finding a whole doc. I want to filter through it to find the nuggets that are important. That's just too much effort. Can you just tell me the paragraphs I need? Can I find what paragraphs often go together? Maybe with cloud migration topic, I also have cloud migration tools or technologies or processes and those often go together. So that will help identify gaps and things you need and shouldn't be including in your proposal. And then finally what worked and what can we analyze this content to get better at writing proposals.

Ok let's take a look at trying to do that, so what we're going to do first is we're going to start tagging these collaboration areas with opportunity I.D.'s and client I.D.'s and we're going to link that up to our CRM system. So now we can start asking those questions like who worked on this proposal? Who's the customer? What's the industry? What are the contacts? Can you really leverage that CRM system and tie it directly to the knowledge and now start giving context right context to this knowledge? What was the opportunity? And so on? So that we can use that to answer these questions. Now let's look at the next few questions. Can you know about looking deeper into the content and this is where I think the new neural networks can be very helpful.

First, we start by slicing and dicing all of the proposal data into smaller pieces. That would be one Power Point slide, one section, maybe even one paragraph. And now we apply neural networks to understand that those neural networks will come up with these vectors and then these factors will go to a semantic search system. Now the inputs can be queries, paragraphs, or sections, and the output can be paragraphs, Power Point slides, and proposal pages. So it's like you're mining right. You're digging up your yard and you're looking sifting out. All that dirt and finding those gems in the yard so that's much more what a proposal process is like. It's like digging up a big area and finding the pot that was buried in your yard. Somehow using neural networks to find those gems that are appropriate to provide a proposal output. And then of course we can actually start selecting those pieces and starting to assemble an initial draft just by finding pieces that are appropriate, that are successful, that look right. And now we have a little start of a proposal that we can really accelerate the process. And each of those will be tagged with the people who worked on it, what client it was for, and what industry it was for, and also if it was successful. And so we can start now cross correlating those paragraphs to find paragraphs that often go together so we can start recommending new paragraphs and new slides. And then finally what worked and what didn't. We can now start cross correlating that with our database. So what does this tell us? It's all not always about document management, you know collaborative systems, and other ways of managing content. And knowledge is absolutely there and needs to be included. Find ways to tag content to business data, find the way to leverage new neural networks which can actually read and understand things, and then finally correlating the content will really help find relevant high-quality materials.



And just some closing remarks. Make sure you understand the business value. It really needs to solve a business problem. Common content management systems are always helpful or often helpful but not always. You still need a great search engine. Search engines really smooth the process out and then what's next? Taxonomies and classifications. We'll talk about that in our webinar. We just have a few minutes for some questions. And I hope this was interesting and helpful.