00:02:25.250
Joanna Jackson (Host): Hello everyone, and welcome to today's webinar: Information Architecture Strategies and Large Drupal Sites. My name is Joanna Jackson. I'm a UX designer and researcher with Sandstorm, and I'm going to be your host today. My pronouns are she/her and I am a Caucasian woman with light brown hair wearing a black top and glasses. Today, we'll be talking about how to manage your information architecture strategy for your content heavy website, and showcasing some real-world examples as well. If you have any questions for the speakers, please submit them in the Q&A channel, and we'll address as many as we can. But for general comments, feel free to share those in the chat channel. If there's anything that you would like us to expand on or talk more about in our next sessions, please let us know. We're always looking for ways to help you and your team have the tools you need to succeed.
A couple more things here. Today's presentation will have audio captions and the session's being recorded. So next week, we'll provide a human-corrected transcript and a link to the recording, and immediately after the presentation today, we will share the deck. Thank you so much for attending. We are so excited you're here. Over to you, Anne,

00:03:34.260
Anne Lentino (Presenter): Great. Thanks so much everybody for being here. My name is Ann Lentino. I am a Senior Product Owner at Sandstorm Design. My pronouns are she/her, I am a Caucasian woman with curly reddish hair. I'm wearing a floral V-neck top, and I am really excited to talk about information architecture and Drupal today. A few quick fun facts about Sandstorm: We've been around for over 25 years. We work with purpose-driven brands. We have accessibility certifications and expertise across associations, global nonprofits, higher ed, and healthcare.

So just wanted to give a brief overview about us. And today, we'll cover in this talk how user research can help drive content strategy when we're thinking about reworking a site map, and how we are going to handle the architecture of a redesigned site, and how that information architecture is a foundation for really good user experience. And then we'll get into how we build that using Drupal strengths to manage a very complex and broad site, as Joanna mentioned, we are going to be showing some real world examples of a redesign that we've done somewhat recently, so that we have some some examples that we can point to for how we handle the strategy and the build.

00:05:09.800
So first, let's talk briefly about the research, or, as I like to say, "Tell us how you really feel". When we do the research, we really like to focus on learning about the needs of the business, the organization, and what those goals and KPIs are when we're tackling a big redesign, and also talking to content editors, because they're users of this site too. They are in the website, day in, day out, and their experience needs to be as good as our end users' experience. And we learn about the content calendar, how frequently they're publishing things, and what kinds of updates they're getting and requests that they're getting from other stakeholders.

In our example, we're going to be showing you some aspects of the Milken Institute redesign that we did recently.

We talked to project sponsors, the main core group who we were working with day in and day out on this project, different stakeholders across the organization, from different departments, again, those content editors and then, of course, end users, to glean what they would expect from how they want to find the content on that website and how they interacted with the website before we did the redesign, in the interest of our time today, we're just going to highlight a few of those findings that really helped drive the strategy that we used for the information architecture.

So one of the biggest findings we had is that everybody uses search, not just the end users, but folks at the Institute as well, and what we heard from is that it heard from all of them, really, is that it didn't perform the way that people really expected search to behave on a site like this. So the results were somewhat inconsistent, and functionality that folks have really gotten used to, like filters and facets, were hidden. You couldn't find them easily when you landed on that page

00:07:09.660
One of the other things that we heard, and this was also something that was a goal of the site, from the project sponsor perspective, is that related content was really important. So when we talked to end users, they tended to fall into one of two behavior categories: They either browsed the site because they were sort of generally looking for things in a specific topic, and they wanted to be able to find all sorts of published content in that topic, or they were super laser-focused. They came to a very specific piece of content because it showed up in their email inbox, and that's what they were there to read. But in both of those cases, there was a resounding agreement that being able to get to valuable content in that relevant topic that they were interested in was really important to them.

The other thing that was related to that is they relied on the credibility of the institute's leaders and contributors, so being able to access other work by those leaders was also very important to them.

So the good news is, there's the taxonomy to handle all of that. So we knew that to handle the related content needs, we needed very specific and very descriptive taxonomy terms and very clearly defined related content rules to be able to support the needs of the end users. We also knew that they wanted to be able to find the newest sort of research and reports and those very important nuggets of content on the site easily. But then from the business side, there were program leaders who had ownership over those assets, and they wanted to be able to feature that content on their specific pages on the website.

The previous website sort of buried that content on those pages, and we needed users to be able to get to that quickly, while also making sure that those different programs or focus areas had could showcase us as part of their overall strategy. So we needed to balance those, both of those needs at the same time. We also knew, again, search was going to be a big priority, so we would need those taxonomy to help drive those facets and filters so that we could expose them

and let people kind of whittle down a really large body of work, and then finally, no dead ends. We really want to make sure that we can close the loop on different content of a specific topic or a specific format, so that we allowed people to continue down a path that they were really finding information that was super useful to them.
So where did we start with all this? Well, I like to say we measured twice so that we could migrate once. The way that we approached this, the different needs for the navigation, the site map, the information architecture is what we call our yes and approach, if you're familiar with Chicago improv. So yes, we wanted to do what the business needed and what the different organization, organizational areas and stakeholders needed, and we needed to honor what we heard in the user research too, of how the users expected to find their content. So users told us our inspiration work, showed us our competitive analysis all confirmed that users wanted to be able to find content, either by topic or by format or type in as few clicks as possible.

So we needed to make topics a priority, and we needed a hub for published content where people could find specific types of content, podcasts, newsletters, thought leadership, those sorts of things. We also knew that we needed to make sure events, which was a very important part of the institute's model had a permanent home in the NAV but that was for the folks who maybe wanted to attend or wanted to know who was speaking. All of the content the sessions, a lot of which were recorded and lived on the website. Those needed to be surfaced as well, sort of in their own right, and as part of the sort of broader architecture of the site where users could find them, you know, when they were searching in a specific topic area. So we needed to be able to highlight those panels as part of the content hub in a way that they weren't before.

We also needed a way to communicate the institute's organizational focus areas or pillars for users who wanted to engage with the Institute without, you know, burying some of that highly valuable content only on those pages. So we needed to be able to create a home for each of the different programs and portfolios so that they could also show the value that they were bringing to the Institute and the institute's audiences.

00:12:13.170
So how did we accomplish all of that in one navigation? Well, we had to start with an audit. We inventoried absolutely everything.

So I equate this phase of a project to the big purge you do before you're moving from one home to another. So first you're wondering, how on earth you amassed all this stuff, and for the love of God, why do you have seven lemon juicers? And then it gets really messy while you figure out sort of what stays, what you could repurpose and move to a different room in the new house, and what goes and in the end, everything, after all of that work and all of that mess, everything ends up in nicely labeled boxes and gets packed into the moving truck.

But with all of that documentation of the current state site map, all the different top content types, all the taxonomies they had, media that they had on the site, we were able to confirm with the team what content was prioritized to remain on the new site, understand what was being used to display all of that content on the current site, and where the gaps were between the current site and our vision For what we wanted the redesigned site to be.

so then we had to give that prioritized content a new home. So we were able to organize the published content into two kind of similar buckets to meet the users expectations on how to find the content quickly, we elevated the experts and events into their own sections in the nav we also created a new hierarchy for the pillar structure, the way that the organization was had organized their portfolios and programs. And then we also wanted to make sure keep a nice tidy utility now for all those things that are very important, but maybe don't fit into the other categories that we had put in the main navigation, and then we wanted to give it a fresh coat of paint.

00:14:17.950
We needed to make the menu pleasing and easy to interact with. So when we're approaching this design of a nav, we often think, do we make a menu, or do we not so in order to determine whether or not that's a good fit, we tend to ask a couple of questions: “How many levels deep is this site map going to go?” “Are we talking two or three levels?” “Are we going down to level six?” “Could a mega menu potentially overwhelm a user with that many choices?” And then another question is, “Is the content evenly distributed? If I click on one top level navigation item, am I going to get seven options, and the other one I'm only going to get two?” So how is that going to perform?

As we're thinking about consistency in the user experience, a good nav keeps the user focused enough to find what they're looking for. So within the hierarchy of new site, in this case, we felt like a more streamlined approach is a better way to go than giving them sort of all of the options in one big drop down.

And on top of that, we also wanted to make sure that a user can backtrack if they wanted to get to something else within that one sort of focused area.

That brought us drum and roll please a very focused menu that also allows me to use the word bifurcated on a regular basis. So what this allowed us to do is really make use of the nesting of the different pages and terms so that users could get to the valuable content very easily in this high priority left side of the nav.

Here you can see how we created this hub where they could get to content by type, if they knew “Oh, I read a report the other day, but I don't remember exactly what the topic was. I want to be able to go find that so that they could get to that format really easily.” And then on the right hand side of the nav we have the sort of hierarchy of the pillars of the organization itself, so users can learn more about the departments, what sorts of content they're publishing, who's on those teams. But because there are programs within portfolio areas, we wanted to be able to show users how all of those things related to each other, and be able to backtrack if they want to go sort of back up the chain.

00:16:57.720
Once we came up with that strategy, we needed to figure out the way that we were going to build it.

As we align that strategy, we like to come up with and discuss these guiding principles as we're moving through the actual build phase for the site. Generally speaking, we like to keep what can be repurposed. We don't want to reinvent the wheel when we don't have to, it's a lot more efficient to use content types and different pieces of the architecture that are kind of in place, especially when we're doing a redesign, where the platform is already there. We're just kind of transforming what already exists.

The other thing is, be selective with new content types. We don't need to overwhelm our content editors with multiple new templates when one or two variations will do and that's why we want to be kind to our content admins, right? They're the ones who are going to be using the site day in and out. We want to make it very streamlined, easy for them to understand, and also make sure that we're giving them the opportunity to make different choices as other needs arise in the future. We want it to be scalable and really give them some flexibility as they're fielding new requests.

00:18:32.460
First, let's talk a little bit about the taxonomy strategy, because we knew we had a lot of related content needs, and we wanted to be able to showcase content in different areas, we really needed to do a deep dive on what terms were already there and what we knew we needed to account for in the future.

The thing is, when you're looking at a site that already has some tagging and some categories available, they can become very overwhelming and sometimes kind of duplicative. We see this a lot of times where you know this, the way that we tag content in multiple CMSes, you know, you want to have the flexibility to be able to add new vocabularies, add new terms as content needs change over time, but without A clear governance plan or strategy around how those are created. It's really easy to end up with a large list of very loosely related terms, one off cases where you have a category or a term that's only used on like one or two nodes, and it gets kind of messy.

So being able to do a clean house on the taxonomies can be really helpful and useful and unearth some some things that might help an organization or the content editors understand. Like “Well, why isn't the current way we're doing related content working the way we want it to?” Well, it's because you have three terms that essentially say the same thing. The CMS doesn't know that those things are related.

So we like to start with kind of the low hanging fruit. What of the existing vocabularies are are similar enough that we could kind of morph them together in a way that makes sense. Also look at terms that are duplicative, where we looking at terms that are basically the same, but one is capitalized, the other one isn't, or there are singular and plural versions of the same word, those are the really easy things that we can start to start to clean up.

And then, if there are terms that are kind of vague, where we think we understand what that term might mean, but the content that you find under there doesn't really make sense as to why those things are related together under that term. And as we're transforming the different vocabulary sets, we like to nest terms. So thinking of it as more of an outline, and how those ladder up to some of those larger categories or larger vocabulary vocabularies that we're putting together. This is an area that can be a little sticky in terms of what the content needs are. So, if you need an arbitrator, you can always look to AI to help find that middle ground. So we can start with a cleaned up list of those terms to say: “Okay, we think that this is kind of headed in the right direction.”

00:21:48.880
In the case of the Milken Institute site, we were also able to cross check that with some high priority SEO keywords that we knew that we really wanted to aim for, and craft a strategy around—and then used AI to say, “Okay, how would you look at these terms? Revise them into a taxonomy structure that's clear.” You know, that has some sort of hierarchy that makes sense, and is also optimized. And so, in this example, we condensed 17 different vocabularies, 11 of which were mostly unused, into six.

So really using that nesting structure where any term can be used at any level, nesting the terms within each other allows us to accommodate volume and specificity at the same time, so that the content ladders up and also can easily relate to the different programs and portfolio areas for Milken and in this example, so that they can relate more easily to each other and can get down to the very specific nitty gritty terms that makes sense for that particular program. Then, as we go through the taxonomy strategy, understanding what the rules are to make related content, we need to figure out how we're going to structure the different pages’ content types to help accommodate those needs.

Where we start is looking at the different content types. So when we're assessing the Drupal architecture based on what we've set out in our UX strategies and in our creative UI design, our goal is really to strike the balance between flexibility and consistency. That doesn't always mean to get us to as few content types as possible. The audit that I mentioned earlier, includes not just a list of current pages and how they're organized in the on the previous site site map, but also looking at what existing content types or templates are there, what fields are on those, where we have references or links to other pages or media that we know we're going to need to carry forward so that content editors feel comfortable that we're accommodating all of the content that they know they need to update or create on a regular basis. When we're thinking of that migration that we were talking about, we like using structured content as kind of the workhorse of the site. So when we're thinking of migrating or transforming hundreds or thousands of different nodes, we lean on the structured content to help editors get timely content published as quickly as possible. So when we're thinking about sort of using one content type for a large volume of similar nodes, we don't want each one of those variations to have its own content type, because that can get duplicative, and it gives the content editors kind of too many choices, they can migrate it into one and give them options, knobs and levers to be able to add the content that they need in there very easily

So we think of it sort of as adding the fields that act as the outline for the content editors so they can say, all right, I know I need a title, I need an image I know I need to upload this document to make sure that that's there so that users can access this report, or wherever it might be, and then giving them all of the different taxonomies they need to make sure that the related content works the way that they need it to.

00:25:53.250
So let's take a look at an example here.
So this is a report which is a very highly sought after piece of content on this Milken Institute site. So this is what the front end user sees, but the way that we build it is, this is a content type called an article. An article can be used for multiple different types of content on the site, but we're leveraging the different taxonomies and print fields to make this stand out as a unique piece of content. So we set up a format taxonomy so we know, okay, when a user lands on here, they can see, although it's kind of tiny on this screenshot, this is a report, the title, the date, the different topics that this particular report is tagged with. And then the content editor can upload their image, their the PDF that's associated with this report. And then some description, a little short description of what's going to be in this report, which is also really helpful for for search and for SEO. And then this view down here is just automatically displaying related content based on the topics that are already added here. So the content editor doesn't have to do anything with this block. It's just going to show up because they put in all of these topics on this particular node.

So to see what that kind of looks like, on the back end, we have the title, and then we give the content editor a choice of one format, because a piece of content is only going to ever have one format. So that's a drop down, but they can choose multiple topics. Now the way that we've set this up is they use the most specific topic, but we talked about how we nested those taxonomies, so the whole hierarchy of where that particular taxonomy term is nested shows up here, so that this report can display in related content on a page about economic landscape, on a page about finance, etc. And then this pillars taxonomy allows us to very easily put this report as related content on those more department specific pages, so that they don't have to manually choose which reports or which articles display on that page, because it's tagged with finance, it will just show up there. And then we also use this to build out our URL structure. So as we add these different terms and things, this gives us a sort of breadcrumb, not an on page breadcrumb, but in the URL, so that a user knows where they found or where they can get to this content and content like it

00:28:48.260

And then finally, talking about how we build for that flexibility. So the flip side of the coin of structured content is using a more layout based approach. So in Drupal, obviously we're talking about layout builder, and we love using that for the flexibility, but we can also set a default layout for consistency, to make sure that all of the content of sort of a similar hierarchy or similar sort of structure, like all programs are all portfolios kind of have the same layout so users know what to expect every time. The beauty of the default layout here is that it balances the dual needs of giving content editors the flexibility they need to kind of reprioritize or highlight a different topic or event at one time or another, but we have the structure of the UX principles that we established in our design phase carried forward over and over again. So when we put that all together, we can create something where we can combine fields and layout builder and the flexibility of a more drag and drop approach into something that is highly valuable, both for the end users and the content admin, so they can manually change out what report is going to be is either the most valuable at the time or that they want to highlight as part of their initiatives. We've also have a field here where they can link to the children of this particular portfolio. We have the most recent content.

Again, anything tagged with in this example, Faster Cures is going to show up in this most recent content and update regularly.

00:30:35.690
And then, not every portfolio has a newsletter, but Faster Cures does. So they can choose to put this newsletter, sign up block in here if they need it, and then a quick view of the Faster Cures team, where, if we add a person to the site, they can be tagged with Faster Cures in that pillars taxonomy, and that automatically adds them to this page and all of that work that we put in to the tax strategy and the way that we build out these nodes Power Search, so we're able to build out the different facets and filters that users need when they're searching for a particular topic, but they really want to hone down on maybe only research that's been done in this area, or they want to find something by a particular expert at the institute. They're able to do that very quickly here.

00:31:28.160
So some quick takeaways:

We always like to start with research. It reduces the subjectivity of content that an organization has been looking at has had their hands in for a very long time, and it really helps ground future decisions as we're thinking about these different information architecture and development strategies.

Investing that time in auditing and documentation, it can be a slog, but it is well worth it to help understand how migration of content is going to look in sort of the new world of things, and can help catch some of those tricky edge cases that we may not be aware of as we're just taking a look at the site from the front end, you got to think about every page as a side door to the website, so making sure that we have correct breadcrumbs tagging and a good URL structure to help orient a user if they don't automatically come to your home page.

And make sure that your strategy of updating content and the way that you organize it on the site is clear, and that the way that you build it out supports that vision.

We'll have some other links about sandstorms work with Drupal in our deck, but thank you so much to everyone for joining us today. I hope this was useful for you.

00:32:46.730
Joanna Jackson (Host): Thank you, and that was very helpful, especially about the content audit. I think we're big fans of auditing all that content and helping us find those edge cases quickly. Thank you again to everyone that has joined today and tuned in. We are so grateful that you joined. If at any point you have any questions, you've got Anne's information up here on the screen, feel free to reach out. She loves to talk about information architecture. So again, right after this presentation, we'll be sharing today's deck, and then next week, you will have the human correct transcript as well as a link to the recording, so feel free to share that so thank you all. Hope you have a fabulous rest of your day.