

SOFIA Thanks for joining this webinar entitled "Where Do We Go From Here? Web Accessibility Across Three National Systems." I'm Sofia Enamorado from 3Play Media, and I'll be moderating today.

I'm joined today by Cyndi Rowland, executive director, and George Joeckel, web accessibility specialist, of WebAIM. And with that, I'll hand it off to Cyndi and George, who have a wonderful presentation prepared for you all.

CYNDI ROWLAND: Thank you so much. And George and I are just excited to talk about this topic. And this was some research that we did actually last fall. But we love to give it more press. We know that it's critical that we are continuing in the field to look at data in terms of web accessibility and where we are.

George, if you'll advance that next slide. And then you probably can just push them from here, George.

All right, so we know a couple things. As we are looking at national samples of web accessibility, we know that we have had issues in the past, and we know that we can, in fact, do better. We've got a historic range of inaccessible content, really, across all kinds of sectors.

And George, you can keep pushing those buttons there.

And our thinking was that even from folks in the disability field, we would want to be looking at promising practices that are happening. Of course, those in the disability field have even a greater vested interest in making sure that accessibility is available, really, for everyone. But is that true?

So looking at a national snapshot, it really seems like now was a very good time to redo it, to look at it again. And we started, really, with the question-- where should we look? Should we go ahead and look at what we think would be a place where web accessibility is kind of handled and look at areas where we might be approaching this from strength? So that's really what we did. We decided-- go ahead, George.

I'm sorry. It's weird for me not to have control of these slides. Thank you.

WebAIM is housed at one of the University Centers of Excellence in Developmental

Disabilities. They're called UCEDDs. And every single state has at least one UCEDD. Some of the larger states-- California and Texas-- they have more. But WebAIM is housed at one of these UCEDDs. And since every single one of these is housed at a university, and they have this focus on disability, it seemed like this would be a good place to take that national snapshot.

Now, the mission of the UCEDD, as you can see here on the slide is that-- and I should tell you that this network was started back in the early '60s. Actually, two weeks before John F. Kennedy was assassinated, he signed this legislation that started what became this particular system, this nationwide system.

But the mission of the UCEDDs is to provide leadership and advice to federal, state, and community policymakers to promote self-determination, independence, productivity, and full integration of individuals with developmental disabilities. So it seemed like this would be a perfect hotbed for the kind of peek into accessibility that we might like to see.

George?

[LAUGHS]

Thanks.

All right. So we've got a shared context between the UCEDDs and the universities, because every university-- or every UCEDD is housed at a university. And of course, every one of these pairs of UCEDDs and universities are within an individual state. So we've got this shared context. And we thought, you know, it would be an interesting thing to look at those three systems, to go ahead and look at accessibility of the UCEDDs, because they have a disability focus and they are a national network, to go ahead and look at the accessibility of the universities where these UCEDDs are housed or hosted, and then also look at each of the state governments and how their accessibility compares.

So George, if you want to move that slide again. And you can even move it along again.

And we've got this perfect combination for a snapshot, looking across the three systems-- again, disability, university accessibility, as well as state government accessibility.

Go ahead, George.

All right. So we were asking three research questions. And as we get through each of these, of course, we're going to give you this preview. We're going to then talk a little bit about our methods, and then come back to what it was that we found in our research.

Our first question was really, to what extent are the UCEDDs, their host institutions, and the state governments using web content that conforms to current accessibility guidelines? To do that-- well, George will talk about this in a minute. But we looked at just a small section of some success criteria to look at how folks are doing. Now, we weren't looking at accessibility per se. We were looking more precisely at failures of a subset of these success criteria. You'll see how we approached that in a minute.

The second question that we had-- George, thank you-- is we wanted to answer the question as to whether there were any relationships between the web accessibility on these three systems. So again, the UCEDDs that are focused on disability, their host institutions, as well as the state governments. Very specifically, we were thinking to ourselves, gosh, if we have certain UCEDDs that are doing well, is there some relationship to, let's say, being in a state that has a law or an institution that has policy? Those were some of the things that we wanted to get to. But first, we needed to establish whether or not there were relationships between these particular entities.

And then the third question that we wanted to dig around a little bit at-- we wanted to look at whether or not we were seeing some patterns of accessibility errors across our entire sample. And of course, if the answer to that is yes, then that was going to help. That really would help the entire field more surgically assist folks.

If we know that here are the areas where folks appear to be struggling more than other areas, those are places we could go in and provide more immediate help. If we were to find, however, that the patterns of errors are fairly random, then we just need to proceed maybe an entity at a time, or look at maybe within each of those groups if there are patterns. But at this level, we were just looking across the entire sample to see if we had patterns.

So with that, at least you've got a little bit of an intro into what our questions were. And now, I'm going to pass it off to George to talk about how we gathered our sample and what our methodology was.

GEORGE
JOECKEL:

OK, thanks, Cyndi. So as Cyndi mentioned, we're looking at these three different types of organizations-- UCEDDs, hosts, and governments. And we had 62 UCEDDs. And this isn't the

full sample. Some had to be eliminated because they didn't have enough web pages for our sample. We looked at 62 hosts and then 52 governments, which, again, would be states or territories in the United States. So if you put all those together, that gave us 176 different websites that we could look at.

So what we did is we created a spider for each of these websites. And the spider, its job was to go out and find, or harvest, 250 URLs.

Now, I do want to make a shout out. A big thing that I want to get across to people is everything that we did, all these resources that we used in our evaluation were either open source or no-cost tools. So I hope one of the messages that you take from this webinar is that this is a very achievable methodology that anyone can replicate. And so I used this webcrawler. That's no-cost. And that's one of the reasons why we chose 250 URLs. That's as many as it would gather.

But what we did was I created a filter for that spider to say, ignore these different types of things. We're not interested in Word documents or PDFs or movies. So all of that stuff got filtered out.

And then after we created that spider, then we created 175 more of them. As you do with things when you're using open source or no-cost, or no-fixed-cost resources, there's a little bit of legwork or elbow grease involved in building those.

And then once we had these individual pools of 250 pages, what we did is we took out 13 random pages. And that number actually comes from the literature. And I won't get into-- that gives you a very high degree of confidence with a random sample.

Now, one of the other things that we did, and one of the ways that, again, without spending a lot of money, you could really increase the validity of each sample, is every now and then-- so that would leave us with 237 pages after we took out the 13 random ones. But we would have pages that would come back as invalid for whatever reason.

But what we were able to do-- there were a limited number of random or invalid pages. So we could just poll again from that pool and replace those individual pages. So in that regard, we were able to make sure that we had 13 random pages for each website. And then we also added in the home page to that sample.

Now, the tool that we used to evaluate these pages is called WAVE. And it's a tool that we provide free to the community. And there's the address, wave.webaim.org. And there's actually a plug-in for Chrome so that you can use it right in your browser. And again, this is something that's available at no cost.

And what we used the WAVE evaluation to do was analyze about 2,500 web pages, record 50 types of machine-discoverable errors. And just to break that out a little bit, this means that in terms of these first evaluations, there was no human intervention. So these are all errors that a machine can discover through algorithms that it'll run on the page. And then from those pages, we had reporting on 12,351 errors.

Now, what we wanted to do with our WAVE errors-- and there were 17 of them that we used. And again, I think it's important to understand that these are machine-discoverable errors, so this is something that our software program can find, which really translates into two things. One is these errors are actually fairly easy to remediate because a machine can discover it. So there are some things where you get into a gray area. But because these are such black-and-white errors, they're among some of the easiest to fix. The downside, or the relevance, I guess, of these errors is that most of them have a pretty significant impact on accessibility.

And so what we did is we took these 17 WAVE errors. And there are seven WCAG 2.0 success criteria that we mapped them to. And so 1.1.1 is Non-text Content. 1.3.1 is Info and Relationships.

And I'm sorry. I probably should have defined what WCAG 2.0 is. That's the Web Content Accessibility Guidelines. They're put together by W3C, the people that run the internet. And it's in the title. They're guidelines. They're not really meant as evaluation criteria, but that's how they're being used right now. So that's what we mapped to. That's actually what WAVE maps to as well.

OK, so moving on, Pause, Stop, and Hide. And when I show you some of the results, hopefully it will become a little more obvious of what these are. 2.4.2 Page Titled, 2.4.4 Link Purpose, 2.4.6 Headings and Labels, and then 3.1.1 Language of the Page.

And so what we did is we took these seven success criteria and mapped them to our 17 WAVE errors. You might not be able to notice right off the bat here, but there is a little bit of overlap between some of these categories.

Now, I'm going to hand it back over to Cyndi to talk about our analysis.

CYNDI ROWLAND: You bet. And George, I thought I'd let you know that I just got kicked off, so I'm just going to talk here from what I know that we're doing.

So with the analysis that we conducted, of course, an awful lot of it is just going to be straight-up descriptive statistics. So we were looking at averages and confidence intervals, things like that. We did some Z-tests looking at whether there were differences across groups. And then we also wanted to look at, were there are some relationships between the groups? So we did some correlation coefficients between the UCEDDs, the host universities, and the government, the state government web pages.

Now, something that's really, really important that we want to share before we get into any of our analysis of the-- I mean, as George says, nearly 2,500-- but I'm going to go ahead and say 2,464 pages across our national sample-- is that remember that we were only looking at pages that had these machine-discoverable errors. So if we have a page that doesn't have a single error, that does not mean it's an accessible web page.

But our thinking was that if somebody has web accessibility in mind, there's going to be a good chance that they won't have these kinds of problems. But again, it does not mean that we have an accessible page. And also, because there are an awful lot of accessibility errors that need manual checks, that's another reason we can't say that just because in our sample we don't have an error that somehow it's acceptable. So make sure you're not getting that piece confounded at all.

So for us, if on a page, if there was a single error from those machine-discoverable WAVE heuristics, that was an error of the success criteria. And so that would be considered a failure of that page. And of course, if there isn't an error, it's the same thing as my first point. It doesn't mean that it is a pass of that particular success criteria.

So I am going to start with the first question. But then I'm going to have you guys use a little polling here. So think about this as I ask.

So now, what we want to do, I want to give you the results from our first research question-- to what extent are UCEDDs, their host institutions, and state governments using web content that conforms to current web accessibility standards?

So one of the poll questions that you should have available to you right now-- hang on one

quick second. Let me get to-- OK, so you can see the poll up there. What do you think that we found across our sample for the proportion of failing pages? We just threw it into quartiles. Do you think that 0 to 1/4 of all of our pages failed? Do you think the second quartile, up to 50% failed? Do you think it was over half going 51% to 75% failed? Or an awful lot, 76% to 100%?

Now, I've got to say that I am asking that as our total sample. And of course, we are going to be sharing each individual system. But if you are thinking about the total samples, what would you put there? So we're going to come back to that in a second.

And then the second poll that I want you guys to go ahead and weigh in on, all right, is, do you guys think that there was any difference in our organization types? Now, remember, there could be reasons-- and I guess we should have had, in fairness, none of the above, or I don't know, or something. So apologies that we didn't put that on our poll.

But it could be that UCEDDs, in part, because they are a disability organization, did a little bit better. It could be that universities that do have more resources and seem to have this on the bubble a little bit more, lots of policies and practices happening, that that's going to be an advantage. Or it could be that the state government web pages were advantaged, in part because an awful lot of states are having some emerging law and policy work on this.

So we're really interested in your ideas. And of course, this will give you an opportunity as well to see was your thinking in line with what actually happened here.

[LAUGHS] All right, so now, let me go to-- hang on. Let me go now to our results. And I've got the funniest thing happening on my computer here. Here we go. All righty.

So let's look at the results of the first poll question. And George, you're going to have to chime in here. What did our listeners think about the proportion of failing pages in our sample?

GEORGE OK, so we have the first quartile, 4%. The second quartile was 25%. And then the majority of
JOECKEL: people, 50%, said 51% to 75%. And then for the last quartile, 76% to 100%, that was 21%.

CYNDI ROWLAND: OK, all right. And how about for the second poll? Did they [AUDIO OUT] one of these was going to do better?

GEORGE So like you said, as UCEDDs, their whole entire mission is to serve people with disabilities.

JOECKEL: That expectation is borne out. 58% said UCEDDs. 20% said university hosts. And then 22% said state governments.

CYNDI ROWLAND: Yeah, and sadly, I think if we had had a none of the above, we might have had a few [INAUDIBLE]. We really forced you guys to make a choice here.

OK, so let's go to our bar graph here of the results. And interestingly enough-- so let's all think about the first polling question. The majority of you considered the proportion of failing pages would fall somewhere between 51% and 75%. The average across all three samples was 81%.

Now, to us, this was actually a little alarming, that out of 2,464 pages-- and now, remember, this is only the tip of the iceberg in terms of accessibility. What we were looking at were just these 17 machine-discoverable errors. The fact that 81% of the pages across this sample had at least one of these errors was an unexpected result for us.

Also unexpected was really the level of errors across the three systems. We really did think, as most of you guys did, that the UCEDDs themselves, because they are a disability group, would be a little bit ahead of the curve from some of these others. And in fact, that wasn't true.

And you can see also on here that the confidence intervals are really quite tight. So we didn't have a ton of variance on even just the descriptive statistics that we had run.

But the UCEDDs overall had 80% of the pages in their sample failed. Or I should say, 80% had failures in at least one of those 17 machine-discoverable errors. The host universities, 85% of their pages had some of these problems. And for government-- I mean, whoever thought that governments were doing slightly better, first, you're right that it was governments, and second, it really would be slightly better, right, with only 1% fewer failing pages than UCEDDs and only 6% fewer failing than the host universities.

Just because we are part of a UCEDD ourself and we thought that it was maybe important to get the message across to the UCEDDs how they're doing, we went ahead and did a little bit of a conversion of our own scoring. We mapped it onto each of the UCEDD's performance. And we thought to ourselves, you know, they're all at a university. All these UCEDDs are at a university, so we'll just go ahead and give them a grade.

So here is what we did. We didn't grade them on a curve. We graded them with standard grading. If you look at the As, that means that those five UCEDDs, they could have one page with errors, but they couldn't have had more than one page with errors. But we had five that

fell into that category, and they had an A.

You can see the the Bs, the Cs, and the Ds. But now let's move on. Well, actually, let's move on to the Ds and the Fs.

So when you're thinking about the fact that there are 62 UCEDDs and the fact that of that, 53 of them had so many errors that they would have received a D or an F, we've got a crisis in terms of what our awareness, as well as the application of knowledge and skills of web accessibility is inside of just the disability organizations themselves. So this is, like I said earlier, a little unexpected.

So as we look at our first question-- to what extent are UCEDDs, their host institutions, state governments conforming to current standards-- well, they're not. So we're going to stamp that one a big fat fail at this point in time. Now, we certainly hope that in the future, these data will change dramatically.

I also wanted to mention that as we looked at differences for groups, we just did some simple Z-tests, difference of proportion tests. And again, none of our groups did that great anyway. But we still were trying to see if there were some group differences.

And interestingly enough-- and probably it's the rule of large numbers here-- we did have some statistically significant differences between some groups. And that was at the 0.05 level. What I can tell you is that on Z-tests, the UCEDDs actually did slightly better than their host institutions.

And governments did slightly better than the host institutions themselves. But on the final test, there was no significant difference at all between the UCEDDs and the government sites. But again, keep in mind that even though we're detecting some differences, none of them really did very well at all.

We also wanted to look at the relationships across the three different systems. And as you can see on this correlation matrix, we just did simple Pearson r's here. We did not have any that were statistically significant. And all of them are fairly weak associations, weak correlations that we found.

So just to sum up this piece of it, this second part, the differences between the groups were fairly small. The relationships between the groups were quite weak. And the groups in the sample really all performed poorly in equal measure here, which is not quite what we were

hoping.

But I've got to say, just as a little minute of an editorial, that that's, in a way, kind of the joy of research. You go in with your hypothesis, with your expectations. But the data are the data. And in this instance, the data took us in directions that were unexpected for us. But now, this is information that's going to be very useful as we move forward in the future.

But let me go back to conclude this second piece, well, the second research question-- are there relationships between the accessibility of UCEDDs, their host institutions, and state governments. And actually, the answer is a no, which confounded us a little bit. But that's where we are.

Our third research question, I'm going to toss back to George, because he's going to talk a little bit about patterns of accessibility errors. Go ahead, George.

**GEORGE
JOECKEL:**

Thanks, Cyndi. OK, so again, just repeating. Research Question 3 is, do UCEDD, host, and government failures in a subset of success criteria-- because again, we're not looking at all of them-- suggest any patterns of accessibility errors?

And again, we're going to throw it back to you guys. And we give you a not sure option this time. So do you think we uncovered patterns of accessibility errors? Yes, no, not sure. You probably don't need too much time to just say yes or no, so we'll go ahead and move on from that.

CYNDI ROWLAND: Although to me, it's palace intrigue. I'm always going, OK.

**GEORGE
JOECKEL:**

OK, so what we get is 82% said yes, 7% said no, and 11% said not sure.

All right, so what we're going to do is we're going to take a look at the success criteria page failures across the groups. And so what this bar is showing right here is four of the different success criteria that we looked at with the results by UCEDDs, hosts, governments. And then there's a blue bar going across, which is the average. And we're going to look at each one of these components individually. We just wanted to show you the overview.

OK, so success criteria 1.1.1, and this WCAG 2.0. It's a specification. This language is technical and somewhat jargon-laden. But we wanted to make sure that you had the actual language, that we'll try to make that a little more understandable. So non-text content-- all

non-text content that is presented to the user has a text alternative that serves the equivalent purpose, except for the situations listed below. And we won't get into all the exceptions. This is only an hour-long webinar.

OK, so we'll put that to the side. And what we're talking about is really images here. Right? And what WAVE is saying is basically, does this image have an alt attribute? And so to expand that out a little bit, we weren't looking for something that was equivalent, let's say, or the right alt attribute. Some content management systems will just throw something in there. This means it was just not present, which means, to me, it's just completely off the radar. And obviously, it's a very easy thing to fix to make sure that all the images have those.

And this, what we're looking for are alt. If it doesn't-- let's say if it's decorative, or doesn't have content or function, doesn't serve a purpose, it would just look like that, alt equals quote, quote.

And again, going back to these different groupings, the UCEDDs and governments failed pretty hard on these. The hosts-- so this is the state-- not as much. And I don't know if that's necessarily due to content management systems or just more awareness, but it suggests that hosts are much more aware of at least needing to have that attribute in there.

OK, success criteria 2.2.2 Pause, Stop, Hide. For moving, blinking, scrolling, or auto-updating information, there's a mechanism for the user to pause, stop, or hide it. And this touches on a couple of different types of disabilities. At an extreme level, blinking content at a certain rate can trigger an epileptic seizure. As well, if you have content that just auto-loads and doesn't give the user the ability to pause it, a lot of times, that content is just delivered too quickly.

And the biggest criminal that we're seeing right now is carousels, the slide decks that we all love to put up with the nice, big images. And they're very attractive. But often, those widgets, let's call them, don't give users a way to pause them and resume them. And also, there's many users that don't use a mouse. They have to use the keyboard only. And many times, those controls are not keyboard accessible. Autostart is a big problem.

And then, if you're wondering if you should use a carousel on your website, you can go to this website created by our colleague, Jared Smith, shouldiuseacarousel.com. And it has lots of information for you.

So if we look at this, almost nonexistent for UCEDDs and governments. But hosts really got

hammered on this. And I would speculate-- and this is just speculation-- it's because universities are very into carousels right now, variants of content that's loading up.

All right, success criteria 2.4.4 Link Purpose (In Context). So the purpose of each link can be determined from the linked text alone or from the linked text together with its programmatically determined link context, except where the purpose of the link would be ambiguous to users in general. Now, again, just because we didn't record any errors-- what we were looking for is images that were linked and didn't have any text in them, or different types of things-- doesn't mean a pass. So again, linked images.

Social media bars, many of these social media bars use a font icon to put up these images. And so if I hear F, I don't know what that means as a person that uses a screen reader that's speaking out this content to me. So those social media bars need to have proper labels so that users can know what their content is or their function.

So hosts, again, did pretty well on this. UCEDDs, unfortunately, were the worst. And then governments were kind of in-betweenish there.

Success criteria 2.4.6 is Headings and Labels. Headings and labels subscribe topic or purpose. And again, we weren't checking to see were these headings effective, or were they structured in the right way. If there's an empty heading, an empty label is what we're looking for.

And one of the things that I want you to just imagine, headings are a way that people that can't see web content can navigate through that content-- the first-level heading, second-level heading. It's like a table of contents. So this is what Google News would look like without structure, just a big blob of information run together. And so again, very critical to user accessibility.

So UCEDDs and governments did very well on this. The hosts did not. They did very poorly. And again, we don't really necessarily have time to speculate on all those different things. It's not quite as cut and dry. Hosts, I think, for the most part, probably have more use of structure, and so they might tend to have empty headings. But again, this is somebody that's using these elements, but not putting anything in there.

So for our third question-- is there some pattern-- the answer to that is yes. And I think, again, as Cyndi said, that's very encouraging because that suggests that there will be some pattern

of remediation that we can follow as well.

OK, so I'm going to throw it back to over to Cyndi.

CYNDI ROWLAND: Yeah. So this is just the last couple minutes of our presentation. And then I know George and I are both very interested in what it is that you guys have to say, either vis-a-vis questions or comments.

But what does all this mean, all this fuss and bother that we did this past fall? Well, most sample pages-- well, most of the pages in our sample-- and again, I do want to just remind folks these were national samples. These were randomly selected pages with some pretty solid methodology. And yet, we were only looking at 17 machine-detectable errors.

And everyone in the accessibility field understands that this does not mean it does not equate to accessibility. But most pages in our samples really had basic errors that affect accessibility writ large. And we were a little amazed with the level in 20-- well, at the time we did this, it was 2016-- that the rates of pages with these kinds of errors were as high as they were.

Also, the UCEDDs that have a core focus on disabilities were not more accessible than either their host institutions or the state governments. And again, as you've heard us say before, that was a hypothesis that we were going in with. We were genuinely thinking that we would have a place of strength here from which we could then launch some more work.

OK, George, why don't you share other things?

GEORGE

OK, yeah. So what else did we learn? I think we demonstrated that an enterprise-level

JOECKEL:

evaluation can be conducted with free resources. So just shooting straight from the hip, even if you have thousands of web pages in your website, you can do this.

Now, is this the end? No, this is more the beginning. But as Cyndi said, this basic level of accessibility errors for anybody that has accessibility on their radar at all, these are things that just shouldn't be on your pages. But the good news is that you can do it, and it does not need to cost a lot of money.

The second thing is that automated tools can be augmented with manual input. So going back to this case of where we gathered the sample, and then WAVE would return some errors-- oh, well, this is invalid. Well, the number of those pages wasn't so significant, especially since we had tried to filter out a bunch of different things, that I couldn't go back in there and replace

those pages with a randomly sampled URL that was valid. So that was just one example of augmenting with manual input.

So I'm going to hand it back over to Cyndi for what is next.

CYNDI ROWLAND:Right. So what's next is we're certainly interested in comparing what we found with other national samples. We are aware that there are other researchers doing some things out there.

Actually not long ago, we were on a hunt. We hunted high and low for really large data sets that might set forth what we would consider to be a national-level benchmark of web accessibility, or even researchers that did really large samples. And we're talking large samples. I mean, we don't consider our 2,464 pages a large sample, although we do understand that for many, that would be considered large. But those data just don't exist.

So in as much as we're interested in looking at how do our results compare with others-- and we know others are doing it-- we're very hopeful to see in the near future that we can get really big samples. In Europe, there are folks that are doing tens of millions of pages to look at the kind of stuff we're looking at. And we think that's really, really valuable. And as a nation, I think we're going to need to get to a spot where we recognize the importance of this, and we can make movement towards that.

We also want to delve into variables of influence. So it's really time. And George, I don't know that I've actually talked to you about this yet, but you'll hear it now. This fall, I'd like to redo this and add in some other variables so we can actually look at, for example, which of the state governments have a state law, which of the institutions have a policy in place.

We were initially going to look at it in a backwards compatible way if we saw some differences. But I think we need to be proactive and collect those data on the front end. But I think that we may find-- of course, that's a hypothesis for later-- that there are some mediating or moderating variables that are going to account for some of the accessibility.

We also think it's important right now to study those A organizations in the UCEDD sample. So we know we've got five that were more successful. It would be a fabulous opportunity at a qualitative level to find out how it is that they're pulling this off, because these are still large organizations-- I mean, not as large as a university.

But for example, at our University Center for Excellence, we've probably got-- I don't know-- over 220 professional staff. We're running about 80 different projects. Webbing is just one of

them. But each of those has a website. We have quite a large business component, PR component, all of those things. And most of the UCEDDs are that equivalent size.

So why is it that those five did so much better? This would be a good thing to look into.

And then the other thing that is next is, of course, a no-brainer. Since most of the sample show that there are problems-- actually, for us, problems that were unexpected, or at least to the degree that we saw the problems. This is opportunities for those of us that are coming out of accessibility to reach out to these folks, make sure that they're getting the training, the support that's needed, that they're being linked to resources so that we really can turn this around.

Folks need support. They need technical assistance. They need a lot of stuff. I think before we're going to see a marked change in the data that we found-- of course, assuming that our sample is representative of the population, which is always something that you wonder about.

So here we are at our last slide. And George put that in. And I love this little picture. It's a picture of a little boy looking quizzically at a woman. And the title of the slide is Questions or Comments, which is, I'm sure, what's running through his mind.

And, Sofia, I think that you are going to broker those for us. And George and I'll just respond to what it is that people have to either ask or say.

SOFIA

Yes, thank you so much, Cyndi and George. We're getting ready to begin the Q&A. So I want to encourage everyone to continue to ask questions by typing them directly into the Questions tab.

ENAMORADO:

So someone is asking, what was the most common failure? And did it vary by organization?

CYNDI ROWLAND: So, George, you've got that. But I think you might want to pull up your spreadsheet.

GEORGE

Yeah, let's see.

JOECKEL:

CYNDI ROWLAND: And of course, we have all of that. If you want, Sofia, what we could do is go to the next question while he looks that up. Of course, we do have all of these data, and we're very happy to make sure that that gets out to the group, even if we don't have time to present that here.

GEORGE

Yeah, I do have it here, actually. I just wanted to make sure for the sake of surety, which is it's

JOECKEL: 1.1.1, which goes back to making sure that non-text content, which are, for the most part, images, have alternative text, which doesn't always have to be an alternative text tag per se. Sometimes, that alternative text is provided nearby on a web page.

But if that's the case, it still requires an attribute. And in that case, it would be a null tag, alt equals quote, quote. What this does is that if it is a decorative image, or if the alternative text is provided elsewhere, then a screen reader can skip that image instead of it just saying graphic with no content or no function.

SOFIA
ENAMORADO: Thank you, George. Someone is asking, are these failures a lack of education or understanding of web accessibility or a refusal to comply for some reason? How can we improve this?

CYNDI ROWLAND: Of course, we would never know the answer to that. I mean, at the heart of it, the question you're asking-- and it's a very important question-- is, is the issue that they can't do or they don't do?

[LAUGHS]

And I think for those of us that are in the field, we've got to guard both of those. We have to do everything we can to make sure that folks have both the knowledge and skills to do it and also work within the system to ensure that there are sufficient motivations. And of course, you guys know what those are. Those are both carrots and sticks, but that folks have a sense that this is important work that needs to be completed.

I know a very long time ago-- and so I'm going to be sharing a tale from, oh, probably about 10 years ago-- we had done a large survey with web developers. It wasn't just education. We actually did it across different organizational types. But anyway, and when we got to the question of whether or not folks that had the knowledge and skills to do accessibility were doing so, we were actually quite shocked that a good number of them said, you know, I can do it, but I don't because my boss doesn't see that as a priority, they want me to work on other things. Of course, back then, folks were looking at this as an add-on rather than it being baked in there.

But I think the takeaway here is that in some entities-- and I'm going to hazard a guess that it's at play in the three systems we're talking about-- that if you don't have administration that understands the criticality of accessibility and supports it, some of the developers may be put

in a bit of a bind as they're trying to get their work completed.

Great question. What's at issue here? Is it knowledge and skills? Well, actually, those are two separate things-- knowledge, comma, skills, comma, a motivation to actually get the job done. And boy, we would welcome anyone who is going to kick themselves into the research fray on that one. That would be fabulous.

GEORGE JOECKEL: I just want to make a quick comment on the end of that, which is, certainly, we don't know if this is a predominance of ignorance or lack of motivation. But I think it's very fair to say that the level of these issues and their discoverability, nobody who's trying to implement these things correctly isn't doing so because they just don't know how. Right? Because again, going back to those images, we weren't trying to say, oh, is your alt attribute correct, or is it equivalent. We're just saying, oh, you've got an image with nothing.

CYNDI ROWLAND:[INAUDIBLE] [LAUGHS]

GEORGE JOECKEL: That's one of the things, I think, is really important to remember is none of these errors were from issues that were so esoteric that somebody just doesn't know how to do it. More that whatever they are doing, like using headings, or using a label, they're not doing correctly because they're empty. There's no value in there.

SOFIA ENAMORADO: Thank you. Someone is asking, it appears that the UCEDDs hosted by universities are having no impact on their host universities' awareness of accessibility issues. Is that a fair statement? To what do you attribute the UCEDD hosts' lack of awareness?

CYNDI ROWLAND: Yeah, whoever typed that in, that is such a great question. One of the things that UCEDDs really are supposed to be doing is working at their own institutions on many different levels of disability. And I would agree with you. I would say that we may not have an influence of UCEDDs to hosts. But interestingly enough, since the UCEDDs did so poorly, I don't know that we'd want to have that influence. In fact, the hosts did a little bit better. So maybe the direction of effect should maybe be going in the other way.

Again, if I ask myself, what explains this outcome, what explains this effect-- and although I know the question is in the context of the UCEDD and post-institutional relationship, I'm going to expand it out to all three. I think that as a field, we've focused on knowledge and skills, which is critical. It is necessary. But I also think it's insufficient.

I think we've got to do a better job as a field looking at the entire system engagement, because that's where the real change is going to happen. And I think that we're still at a spot where some people know about a policy, or a commitment, and others don't. Some have been given the tools that they need for their jobs, and others haven't.

So in the big scheme of things, we have so many giant potholes down this highway, we're surely going to fall into some of them at some point. So in answer to your question, I think we just need to continue our focus on knowledge and skills and figure out how it is that we're going to better impact systems.

**GEORGE
JOECKEL:**

Yeah, I guess I would add to that. I feel like we were able to, at least in this specific domain, move the conversation forward because actually, when we went to present this to the national group-- we did a poster session-- we printed out all the results for all of the UCEDDs and handed those out individually.

And so if there was a lack of awareness before or an ignorance, there isn't now. And as Cyndi said, we also are going to be replicating this study. So we'll be able to see if people have been able to make progress over the last year.

I do think the encouraging thing is that-- one of the concerns I might have is a lot of times, the UCEDDs are pretty locked into whatever the host is doing, the university, in terms of a course management system. And it would be my fear that maybe they wouldn't be able to perform better, or there was one root cause, or if the state was really poor, there would be a correlation.

But because there was no correlation, I think each of these organizations are responsible for themselves. And by identifying which of those three organizations is performing the best, hopefully, they could help bring the other two along.

**SOFIA
ENAMORADO:**

Thank you, Cyndi and George. We have time for one last question. Someone is asking, do you have a centralized list of tools you used in the audit somewhere that we can visit or direct others in our organization to review?

CYNDI ROWLAND: George, I'm going to toss that back to you. It would be relatively easy for us to list the things that we used in the audit. I mean, there's really just a couple. It was the spidering system. And you might be able to just reference that. And then it was WAVE. But I'll let you take this.

GEORGE

Yeah, yeah. So like Cyndi just said, we did use WAVE.

JOECKEL:

Now, in all fairness, we have access to a way to wave these large amounts of pages. We didn't have to do them all individually. But again, when you're considering putting together a sample of 14 pages, even if you have multiple organizations in your organization, I don't think it's unreasonable to do each one of those by hand, which you can do either using our plug-in for Chrome or our server-based one, which is at wave.webaim.org.

And then the tool for the crawling was-- I'll get the URL, and I will put that into the PowerPoint, and then just Excel, which there is a cost for. But I imagine most of the things that we did on Excel could be done on an open source program such as LibreOffice.

CYNDI ROWLAND: Yeah, any spreadsheet application would really do the same, yeah.

SOFIA Thanks, everyone, for joining. And thank you to Cyndi and George for such a great
ENAMORADO: presentation.