Clinical Decision Support Software - Final Guidance
October 18, 2022

Moderator: CDR Kim Piermatteo

CDR Kim Piermatteo: Hello and welcome to today's CDRH webinar. Thank you for joining us today. This is Commander Kim Piermatteo of the United States Public Health Service, and I serve as the Education Program Administrator in the Division of Industry and Consumer Education in CDRH's Office of Communication and Education. I'll be your moderator for today's program.

Our topic today is the final guidance titled Clinical Decision Support Software issued on September 28, 2022. This guidance clarifies the scope of the FDA's oversight of clinical decision support software intended for health care professionals as devices. We're holding this webinar to provide you with an opportunity to learn more and to answer any questions you may have about this final guidance.

It's my pleasure now to introduce you to our presenters for today's program. Dr. Sonja Fulmer, Acting Deputy Director of the CDRH Digital Health Center of Excellence in CDRH's Office of Strategic Partnerships and Technology Innovation; and Jessica Paulsen, a Division Director in the Office of Cardiovascular Devices in CDRH's Office of Product Evaluation and Quality, or OPEQ, who also serves as the Director of CDRH's Digital Health Focal Point Program.

We'll begin with a presentation from Sonja and Jessica, and then field your questions about this topic. Thank you all again for joining us today. I'd now like to turn it over to Sonja to start today's presentation. Sonja.

Sonja Fulmer: Thank you, Kim. And thank you, everyone, for joining us today to learn more about the Clinical Decision Support Software final guidance.

This guidance document is available online on FDA's website as well as on regulations.gov under this docket number.

Our learning objectives today are to describe the purpose and scope of the guidance. We'll be explaining FDA's current thinking on clinical decision support, or CDS, including which CDS software functions are considered devices.

We're going to identify how the guidance clarifies the criteria for non-device clinical decision support software functions. We'll provide examples of those CDS software functions, and we'll also explain how the guidance document complements other existing digital health guidance documents.

In December 2016, the 21st Century Cures Act amended the definition of device in the Federal Food, Drug, and Cosmetic Act to exclude certain software functions from the device definition. The software functions are defined in 520(o) of the FD&C Act and include five different categories of software functions. The Cures Act recognized that there are certain digital health products that are low risk, and many of these products were under enforcement discretion policies prior to the enactment of the Cures Act, which codified many of those existing policies by revising the definition of device to exclude those low-risk software functions.

These include software functions intended for administrative support of a health care facility, those intended for maintaining or encouraging a healthy lifestyle with no reference to any disease or condition, software functions intended to serve as electronic patient records, and those intended for transferring, storing, converting formats, or displaying device data and results.

These first four in section 520(o) 1A through D are described in FDA's guidance document, entitled Changes to Existing Medical Software Policies Resulting from Section 3060 of the 21st Century Cures Act. The fifth category are software functions intended for clinical decision support, and this is the topic of the guidance for today's webinar.

As I mentioned in December 2016, the 21st Century Cures Act amended the device definition to exclude certain clinical decision support software functions from the device definition. In December 2017, we issued our first draft guidance interpreting this provision. The guidance was entitled Clinical and Patient Decision Support Software.

Based on the comments received to that draft guidance, we issued a new draft guidance in September 2019, and revised the title to Clinical Decision Support Software. Finally, in September 2022, we issued the final guidance, entitled Clinical Decision Support Software, to explain the types of clinical decision support that do not meet the device definition.

Before we get into the content of the guidance a little more, I want to take a minute to note that the term Clinical Decision Support, or CDS, is used broadly and in different ways depending on the context. It's generally thought of as a tool that provides health care professionals and patients with knowledge and person-specific information intelligently filtered or presented at appropriate times to enhance health and health care.

However, this is not a definition that's used in the statute for defining non-device CDS software functions. In 2014, the FDASIA Health IT report included examples of CDS software functions such as computerized alerts and reminders for providers and patients, clinical guidelines, condition-specific order sets, and other contextually relevant reference information.

The 2019 revised draft guidance provided FDA's interpretation of the types of CDS software that are excluded from the device definition by the Cures Act. The draft guidance described three categories of software functions: CDS that are no longer considered devices; it also proposed an enforcement policy for certain CDS software functions that would remain devices; and finally, included a category of software functions towards the FDA would continue to focus its regulatory oversight.

The draft guidance used complementary information from the IMDRF Risk Categorization Framework to help make those determinations. We received 67 sets of comments on the revised draft. Some of those commenters asked for changes to the enforcement discretion policy, suggesting both a broadening and a narrowing of that policy.

Other commenters requested explanation of the difference between the term signal and medical information and other terms in the statute that are important for determining whether a software function is device or non-device CDS. Commenters also highlighted the need to better understand the application of the IMDRF Risk Categorization Framework to those determinations. And finally, commenters asked for further elaboration on how artificial intelligence or machine learning software functions can be explained and understood so that they can be considered non-device CDS.

In the final guidance, we provided additional clarity on each of the criteria from section 520(o)(1)(E) of the FD&C Act, including better explaining the distinction between many of the terms in the statutory criteria.

We also provided information and examples to illustrate how health care professionals can independently review the basis of software recommendations, including software developed by AI or ML technologies, so that those software can meet the non-device CDS criteria. You'll note that the IMDRF Risk Categorization Framework was not used in the interpretation of non-device criteria in this final guidance.

And the final guidance also does not include any enforcement discretion policies that were proposed in the draft guidance. The guidance does include additional explanation and examples to clarify these policies, including identifying the complementary existing digital health guidances.

So as I just mentioned, the final guidance's scope is focused just on the interpretation of the criteria for non-device CDS software functions. The guidance does not further explain any enforcement policies for the software functions that remain devices, and rather points the reader to complementary existing digital health policies. The guidance also does not address which FDA statutory or regulatory requirements apply to device software functions.

It does not address which regulatory requirements they apply to a device software function that is part of a combination product nor does it address the labeling requirements for certain decision support software disseminated by or on behalf of a drug or biological product sponsor.

Next, we'll spend some time talking through FDA's interpretation of the criteria in 520(o)(1)(E) of the FD&C Act. There are four criteria that describe non-device CDS, which is a term used to refer to decision support software functions that do not meet the definition of device in section 218 of the FD&C Act. It's important to note that the functions excluded from the device definition are independent of the platform on which they might run.

This slide includes the four criteria that must be met for a software function to be considered non-device clinical decision support. Stated simply, these criteria describe the types of CDS that are not regulated as devices. As described in criterion 1, non-device CDS software functions do not acquire, process, or analyze images, signals from an in vitro diagnostic device, or patterns or signals from a signal acquisition system.

As described in criterion 2, non-device CDS software functions display, analyze, or print medical information in order to provide recommendations about a patient's care to a health care provider as described in number 3. Taken together, criterion 1 and 2 describe the types data inputs used in devices-- under criterion 1-- and the types of data inputs used in non-device CDS under criteria 2.

Finally, the fourth criterion says that non-device CDS software functions provide sufficient information about the basis for the recommendations to the health provider user so that the user does not rely primarily on any of the recommendations to make a clinical decision about an individual patient. If all four criteria are met, then a software function can be considered non-device CDS. We'll be talking about each of the individual criteria next.

OK. We're going to break down the terms in each of the criteria, starting with criterion 1, which, again, says the software functions that meet the non-device criteria are not intended to acquire, process, or analyze a medical image or a signal from an IVD or a pattern or signal from a signal acquisition system. These software functions remain devices, and have been regulated as devices for many years.

So starting with medical image. We generally consider the term medical image to include images generated by the use of medical imaging systems, like X-rays or MRIs, to view any part of the body or images acquired for medical purpose, like for pathology or dermatology. It's also important to note that images that were not originally acquired for a medical purpose but are being processed or analyzed for a medical purpose are also considered medical images for determining whether a software function is a medical device or not.

So for example, if a software function is analyzing an image taken from a smartphone of a patient's skin to determine whether or not a mole is benign or malignant, that software function is analyzing an image, and therefore would fail the criteria for non-device CDS software.

Next up is the word signal. We generally consider the term signal to include the signals that typically require use of either an IVD, which can include an electrochemical or photometric response generated by an assay and an instrument that may be further processed by software to generate a clinical test result or some sort of signal acquisition system that measures a parameter from within, attached to, or external to the body for a medical purpose.

This can include, obviously not limited to, the use of sensors, the collection of samples or specimens, or the use of radiological imaging systems. Based on comments received to the draft guidance, we've also included an interpretation of the term pattern. FDA interprets the term pattern to refer to multiple sequential or repeated measurements of a signal or from a signal acquisition system.

So for example, for an ECG, an electrical signal that's acquired from the body is processed to create an ECG waveform and QRS complex, that's considered a pattern. For continuous glucose monitors, or CGMs, the signal that's generated by an assay in an instrument is processed to generate repeated glucose measurements over time. That measurement of glucose over time is considered a pattern.

Finally, FDA considers software functions that assess or interpret the clinical implications or relevance of a signal, pattern, or medical image to be software functions that do not meet criterion 1. So for example, a software function that's processing or analyzing an ECG waveform or QRS complex, such as measuring repeated complexes, measuring some variation from baseline or detecting heart rate arrhythmias, structural abnormalities, this sort of signal processing and analyzing does not meet criterion 1.

That is another example of software function that processes or analyze an electrochemical or photometric response generated by an assay and instrument to generate a clinical test result, such as for determining a potassium level. That software function also does not meet criterion 1 because it is processing or analyzing that response generated by the assay and instrument.

So now we've spent some time talking about what types of inputs non-device CDS software functions do not analyze, those signals, patterns, and images we just discussed. And now we'll move into criterion 2 that describes medical information that non-device CDS software functions can display, analyze, or print.

As a reminder, non-device CDS are intended to display, analyze, or print medical information about a patient or other medical information. The CDS final guidance describes medical information as demographic information, symptoms, test results, certain medical device outputs, such as a heart rate or blood pressure reading, patient discharge summaries. And then other medical information from the statute can mean clinical practice guidelines, peer-reviewed clinical studies, or textbooks, approved drug or medical device labeling, or recommendations from government agencies.

One way to think about the difference between medical information and signals or patterns is that medical information is the kind of information used by the intended user to make decisions about the prevention, diagnosis, or treatment of a disease or condition for an individual patient.

Medical information normally is and generally can be communicated between health care professionals or between a health care professional and their patients. FDA interprets medical information about a patient to be information whose relevance to a clinical decision is well understood and accepted in the practice of medicine.

An additional consideration for determining whether something is medical information is the sampling frequency, a single discrete test or measurement result that is clinically meaningful, like a blood glucose lab or test result can be considered medical information. And a software function that analyzes that information can be considered non-device CDS if it meets the other criteria.

However, a more continuous sampling of the same information-- for example, continuous glucose monitor readings that shows the glucose measurement over time-- that's considered a pattern or a signal.

Other examples of medical information include the report from a radiology study or summary information about the output of a legally marketed CAD software or an ECG report annotated by a health professional with a description of an abnormal heart rhythm, a blood pressure result from a legally marketed device, or a lab test result in an electronic health record. These can all be considered types of medical information that can be analyzed, displayed, or printed by a non-device CDS software function.

Next, I'll pass it to Jessica, who will describe the interpretation of criterion 3 and 4.

Jessica Paulsen: Thanks, Sonja. So now let's jump into criterion 3. Specifically, criterion 3 states that software functions must be intended for the purpose of supporting or providing recommendations to an HCP about prevention, diagnosis, or treatment of a disease or condition. So a key aspect to highlight for understanding this criterion is the intended user of the software.

Non-device CDS software functions are intended to support or provide recommendations to health care professionals about prevention, diagnosis, or treatment of a disease or condition. Software functions that support or provide such recommendations to patients or to caregivers and not to HCPs, therefore do remain in the definition of a device.

I do want to note that FDA understands the importance of CDS for both patients and caregivers, so FDA intends to be consistent with existing software policies in the regulation of CDS intended for non-HCPs. So for example, in our policy for device software functions and mobile medical applications, we describe

that FDA intends to exercise enforcement discretion for software functions that help patients self-manage their disease or condition without providing specific treatment or treatment suggestions.

So when thinking about what it means when a software function is supporting or providing recommendations, FDA interprets criterion 3 to refer to software that provides condition, disease, or patient-specific information and options to an HCP to enhance, inform, and/or influence a health care decision.

This software does not provide a specific preventive, diagnostic, or treatment output or directive. This software also is not intended to support time-critical decision making, and it's not intended to replace an HCP's judgment. I will add that an important concept described in the guidance is automation bias, which is the tendency for people to over rely on a suggestion from an automated system.

In the context of CDS, automation bias may be more likely to occur if software provides a user with a single specific output rather than a list of options or complete information for the user to consider. So when a single specific output is provided, the user is more likely to accept that single output as being correct without taking into account other available information to inform their decision making.

Similarly, when thinking about the time-critical nature of an HCP's decision making, this concept is important because in situations that require urgent action or urgent clinical intervention, automation bias increases because there's not sufficient time for the HCP to adequately consider other information when making their decision.

Here, you will see outputs that both meet and do not meet criterion 3. Some outputs that meet criterion 3 include software that provides a list of preventive, diagnostic, or treatment options to the HCP; software that provides a prioritized list of preventive, diagnostic, or treatment options to the HCP; or software that provides a list of follow-up or next-step options for the HCP to consider.

In contrast, outputs that do not meet criterion 3 include software that provides a specific preventive, diagnostic, or treatment course; software that provides a specific follow-up directive; software that provides time-critical alarms or alerts intended to trigger a potential clinical intervention to assure patient safety; or software that provides a treatment plan for a specific patient's disease or condition.

Some examples of non-device CDS software functions include software functions that provide evidence-based clinician order sets for an HCP to choose from, tailored for a particular condition, disease or clinician preference, or a software function that matches patient-specific medical information from records or reports to reference information such as clinical guidelines.

Another example is a software function that provides contextually relevant reference information about a disease or condition. And finally, I'll highlight a software function that provides reminders for preventive care or clinician orders. I only mentioned a few examples here. There are many more included in the guidance.

Now let's discuss criterion 4, which states that software functions must be intended for the purpose of enabling an HCP to independently review the basis for the recommendations that such software presents so that it's not the intent that the HCP rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient.

So this criterion really highlights the concept of independent review, and that non-device CDS software functions are intended to enable HCPs to independently review the basis for the recommendations so that they do not rely primarily on such recommendations but rather on their own judgment to make clinical decisions for individual patients.

When thinking about how to satisfy criterion 4 to enable independent review, we've included some software and labeling recommendations in the guidance. Sponsors may use alternative approaches as long as the approach enables an HCP to independently review the basis for the recommendations so that they don't rely primarily on such recommendations.

So let's go ahead and walk through the recommendations detailed in the guidance to satisfy criterion 4. First, FDA recommends that the software or labeling include the purpose or intended use of the product, including the intended HCP user and the intended patient population. FDA does not consider software functions intended for a time-critical decision to meet criterion 4 because an HCP is unlikely to have sufficient time to independently review the basis of the recommendations.

Next, we recommend that the software or labeling identify the required input medical information, including information on how the inputs are obtained, their relevance, and data quality requirements. We recommend that the software or labeling provide a plain language description of the underlying algorithm development and validation that forms the basis for the CDS implementation.

This includes a summary of the logic or methods, for example, any artificial intelligence or machine learning techniques used; a description of the data relied upon so that an HCP can assess whether that data is representative of their patient population, for example; and a description of results from clinical validation studies to allow an HCP to assess the performance and any potential limitations.

Finally, we recommend that the software output provide relevant patient-specific information and other knowns or unknowns for consideration by the HCP as they make a final decision for their patient.

Now that we've walked through all four criteria, we wanted to discuss some examples to help illustrate FDA's interpretation. On this slide, I'll highlight a few non-device CDS examples. I will note these examples focus on the first three criteria, and they're considered non-device CDS, provided that they also meet criterion 4.

So let's begin with an example of non-device CDS that includes a software function identifying drug disease interactions and contraindications such as notifying an HCP that a patient with asthma should not be prescribed a non-selective beta blocking drug.

Another example of a non-device CDS software function is one that aggregates possible post-operative care instructions, medication needs, and follow-up instructions to assist an HCP in assembling discharge papers for a patient. Next, we have a software function that provides alerts to HCPs regarding changes in formulary, and recommends alternatives. Another example of a non-device CDS software function is one that analyzes medical information on a patient's asthma diagnosis and demographics from the patient's medical record, and provides an HCP with a list of FDA-approved treatment options for asthma.

Finally, we have an example software function that analyzes blood glucose lab test results and a pre-diabetes diagnosis from a patient's medical record, and provides an HCP with a list of next-step options to consider, such as more frequent office visits or referral to a specialist.

Now let's go over a non-device CDS example with a focus on how it meets criterion 4. So here, we have a software function that recommends a prioritized list of FDA-approved chemotherapeutic agents approved for the patient's diagnosed cancer type to an HCP based on analysis of reported outcomes in a database of clinical studies using the patient's diagnosis and demographics from the medical record. To enable HCPs to independently review the basis for the recommendations presented by the software so they don't rely primarily on such recommendations but rather on their own judgment. The information described here in these four boxes is provided to the HCP.

First, the intended use, HCP user, and patient population are clearly identified. The use is not time-critical, and the intended HCP is expected to have sufficient time and training to assess the clinical studies that are the basis of the recommendations. Second, the cancer diagnosis and patient demographics are clearly identified as the inputs being used in the database search and analysis. Third, information about how the clinical studies included were selected, the full reports of the clinical studies being relied upon are clearly identified, with a brief summary of the strengths of each study-- such as number of patients, outcome metrics, randomization-- and the key elements of the diagnosis or demographics searched for in the medical record are noted.

Finally, the prioritized list of FDA-approved chemotherapeutic agents and the basis of the prioritization is provided to the HCP, and the studies that most closely match the patient specific diagnosis and demographics are identified. Other considerations, such as the warnings and contraindications from the current version of the FDA-approved drug labeling, are also provided to the HCP for consideration prior to making a final decision.

Now, let's walk through a few device software function examples focusing on the first three criteria. So first, we have a software function that analyzes sound waves-- when users cough or recite certain sentences-- to diagnose bronchitis or sinus infection. This software is a device function because it fails to meet criterion 1 in that it analyzes a signal or a pattern. It also fails to meet criterion 2 in that it's not intended to display, analyze, or print medical information. And finally, it fails to meet criterion 3 because it's providing a specific diagnostic output or directive.

The next example on this slide is a software function that uses a variant call format, VCF file, containing patient-specific genetic variants and mutations identified from a Next-Generation Sequencing, or NGS, analyzer, and provides recommendations for FDA-approved treatment options based on those findings. So this software is a device function because it fails to meet criterion 1 in that it's analyzing a pattern. It also fails to meet criterion 2 in that it's not intended to display, analyze, or print medical information. It does actually meet criterion 3 because it's providing support or recommendations to the HCP. However, all four criteria have to be met to be considered non-device CDS, so this example remains a device function.

Next, we have a software function that analyzes an ECG waveform, output from an FDA-cleared device to detect or diagnose arrhythmias such as atrial fibrillation. So this software is a device function because it fails to meet criterion 1 in that it analyzes a pattern. It also fails to meet criterion 2 because it's not intended to display, analyze, or print medical information. And it fails to meet criterion 3 because it provides a specific diagnostic output.

The next example is a software function that analyzes five sequential RR interval measurements from Holter monitor data found in a patient's medical record to identify a possible heart rhythm irregularity

and recommend follow up testing. This software is also a device function, and that's because it fails to meet criterion 1 in that it's analyzing a pattern due to the sampling frequency of the RR interval measurements. It also fails to meet criterion 2 in that it's not intended to display, analyze, or print medical information. And finally, it fails to meet criterion 3 because it's providing a specific diagnostic output or directive.

Next, we have a software function intended for HCP management of heart failure patients that analyzes patient-specific medical information, such as daily measurements of heart rates, SpO2, blood pressure, or another similar output from a wearable product, to predict heart failure hospitalization. So this software is a device function because while it meets both criterion 1 and criterion 2, it fails to meet criterion 3. That's because it's providing a specific diagnostic output or directive, and it's supporting time-critical decision making.

The next example is a software function that analyzes a radiologist's score or report of regional contrast discrepancies measured from a head CT of a suspected stroke patient to identify whether that HCP should initiate a specific drug therapy based on a scoring algorithm. So this software is a device function because while it's meeting both criterion 1 and criterion 2, it fails to meet criterion 3. It's providing a specific treatment output or directive, and it's supporting time-critical decision making.

The next example is a software function that identifies a specific FDA-approved chemotherapeutic agent to an HCP based on analysis of patient diagnosis and pathologist's confirmed biopsy results. So this software is a device function because while it does meet criterion 1 and criterion 2, it fails to meet criterion 3. That's because it's providing a specific treatment, output, or directive.

And then the final example we'll walk through today is a software function that helps a diabetic patient manage their blood sugar by calculating a bolus insulin dose based on carbohydrate intake, pre-meal blood glucose, and anticipated physical activity reported to adjust carbohydrate ratio and basal insulin. So this software is a device function because while it does meet criterion 1 and criterion 2, it fails to meet criterion 3. That's because it's not intended for an HCP, and it's providing a specific treatment, output, or directive while also supporting time-critical decision making.

So now I'm going to turn it back to Sonja to continue the presentation.

Sonja Fulmer: Thank you, Jessica. And thank you all for listening to the interpretation of the non-device CDS criteria. I know it was very dense, and a lot to digest. And so for that reason, we've developed this graphic that's available on FDA's website to provide a summary interpretation of the CDS criteria.

This graphic provides a more plain language understanding of the four criteria that must be met to be considered a non-device CDS, as well as providing some bullet points examples of each of the criteria. We hope that this might help you better understand the criteria for non-device CDS. However, we do encourage you to consult the guidance for the complete policies. And it's also important to note that the device examples that are identified in the graphic are illustrative only.

I'd also like to note that we've updated other digital health guidance documents to make corresponding changes to align with the final CDS software guidance, as well as with the April 2021 final rule on medical device classification regulations that were updated to conform to the medical software provisions in the 21st Century Cures Act.

We've updated the policy for device software functions and mobile medical applications guidance, as well as the medical device data systems, medical image storage devices, and medical image communications guidance. The changes to existing medical software policies resulting from section 3060 of the 21st Century Cures Act guidance has a new cover page that indicates that this guidance will continue to serve as a reference for the interpretation of the first four software functions that are excluded from the device definition. However, we do not intend to further update this guidance moving forward.

We've also issued an interactive tool to introduce digital health policies to our stakeholders. We understand that there's quite a few different guidances to navigate, and so we've issued this tool called Digital Health Policy Navigator.

The tool incrementally introduces policy considerations with plain language to establish a baseline understanding of policies and pathways and where to look if you need more information. The tool is interactive, with accessible questions that are simplified with yes/no answers, and each of the answers provides information about where to go for more information, which guidance documents-- sometimes even which section of those guidance documents to reference for more information.

Even when this tool doesn't help you find a concrete answer for a particular product, it will help you identify many of the right questions that you should be asking yourself when you're trying to determine the regulatory status of a software function. This tool is also available on FDA's website at the link here at the bottom of your screen.

So in summary, the CDS final guidance focuses on the statutory criteria describing non-device clinical decision support software functions. The guidance clarifies the scope of FDA's oversight of clinical decision support software that are intended for health care professionals as devices. The guidance provides examples of how FDA intends to consider different kinds of software functions, including non-device clinical decision support software and device functions.

If you're unsure of how to apply the guidance or the non-device CDS criteria, please reach out to us at DigitalHealth@fda.hhs.gov. You may also consider submitting a 513(g) for device determination or Q-Submission to come talk with us about how to apply the criteria.

CDR Kim Piermatteo: Thank you, Sonja and Jessica for that presentation. I'd now like to introduce an additional panelist joining Sonja and Jessica today.

Brendan O'Leary, Acting Director of the CDRH Digital Health Center of Excellence in CDRH's Office of Strategic Partnerships and Technology Innovation. Brendan will be providing a few additional comments regarding today's topic, and then assist with the interactive question and answer segment of today's program.

Brendan, I'd now like to turn it over to you for your comments.

Brendan O'Leary: Thank you Kim. And thank you, Sonja and Jessica, for that excellent overview. There are really just a few points that I'd like to emphasize that my colleagues have already called some attention to.

First and foremost, and what many people may not know, is that this guidance-- more clearly than ever before and with more detailed examples than ever before-- makes it clear that even some of the most complex machine learning technologies, including algorithms based on technologies like deep neural networks-- which are among the most powerful, the most promising, and the most complex algorithms in health care and other fields-- can meet the 21st Century Cures criteria for exclusion from the medical device definition. And through the policies and examples in the guidance, it provides a roadmap for developers who choose to go that route.

And then the second thing that I'd like to emphasize is that FDA is maintaining its long-standing final enforcement discretion policies that were already out there in existing guidance, including guidance documents like the device software function and mobile medical applications guidance.

This guidance is just about the statutory criteria. The CDS guidance is just about the statutory criteria for exclusion from the device definition. And it's not about an enforcement policy nor is it about any single disease or condition. Instead, it just helps you determine whether or not a software function you were developing is a device software function.

And with that, I think we can move on to questions.

CDR Kim Piermatteo: Thanks, Brendan. So before we take your live questions, I'd like to go over a few reminders and tips. Foremost, we ensure all hands are raised before the segment begins. So if you wish to ask a question, please click the Raise Hand button, which should appear on the bottom of your Zoom screen. I'll announce your name and give you permission to talk. When prompted, please select the blue button to unmute your line and then ask your question. After you ask your question, please lower your hand. If you have another question, please raise your hand again to get back into the queue, and I'll call on you again as time permits.

Please remember to limit yourself to one question only, and try to keep it as short as possible. And lastly, please refrain from asking about specific submissions. For these questions, we ask that you consider submitting a Q-Submission or consider emailing my division at dice@fda.hhs.gov.

Alright. Our first question is coming from Sophia. Sophia, I have unmuted your line. Please unmute yourself and ask your question.

Sophia Bessias: Thank you. I am dialing in from Duke Health in North Carolina, and my question is that the guidance mentions risk scores specifically as an example of device CDS. Can you clarify whether clinical consensus-based risk scores are always considered device CDS under criterion 3? And we'd be interested in answers for both time-critical examples like Wells score for pulmonary embolism or non-critical examples-- non-time-critical, that is, such as the ASCVD score.

CDR Kim Piermatteo: Thank you, Sophia, for that question. Jessica, would you like to provide a response?

Jessica Paulsen: Sure. Thank you so much for that question. So risk scores are a really good topic. It's actually a good example where the policy navigator may be helpful to you. As we've mentioned, there are several exclusions under 21st Century Cures. CDS is just one of them. We actually have a number of guidance documents that outline FDA's thinking regarding software functions and how they are or are not subject to FDA regulation.

So this policy navigator-- if you haven't taken a look yet, I'd recommend you do. It's on our website. It's a pretty useful tool. And it really just walks you through these policies for a given function to help a developer or a sponsor understand which policies may apply. So for risk scores, for example, or software functions that perform simple calculations that are routinely used in clinical practice, maybe it's a risk score from a clinical guidelines, that is something that we cite specifically that we intend to exercise enforcement discretion for in our mobile medical apps and device software functions guidance.

Sophia Bessias: Thank you.

CDR Kim Piermatteo: Thank you, Sophia, for that question, and thank you, Jessica, for that response. Our next question is coming from Haroon. Haroon, I've unmuted your line. Please unmute yourself and ask your question.

Haroon Hameed: Hi, everybody. Thank you so much for that presentation. I'm calling from Washington DC. I have a question. So if in an application, a patient has self-selected a typical, for example, pain pattern for an L5 radiculopathy, at that point, we suggest treatments across multiple domains-- like medication management, procedures, DME, like bracing, et cetera-- do the options given in multiple domains-- and so then the provider has the option to select whether they want to do physical therapy or they want to do medication management or they want to do a procedure. Does that suffice for the criteria that you need to have multiple options available or do they need to be multiple options within a single domain?

CDR Kim Piermatteo: Thank you Haroon, for that question. I'm going to turn it over to Brendan to address this question. Brendan.

Brendan O'Leary: Sure. Thank you. Yeah, and thank you for that question. And again, I think the policy navigator is an excellent tool. In this case, without getting into too much of the specifics about the specific clinical condition, what the inputs to this product might be, and really the specific nature of the outputs, I think that what you're describing is on the right track for meeting the criteria for exclusion from the device definition, especially the third criterion about providing recommendations or options, rather than specific diagnostic outputs or treatment directives.

And so I would just say that it does sound like you're on the right track here.

Haroon Hameed: Thank you.

CDR Kim Piermatteo: Thanks, Brendan. Thanks, Haroon. Alright. Our next question is coming from Mark. Mark, I have unmuted your line. Please unmute yourself and ask your question.

Mark Sendak: Hi. This. Is Mark Sendak from Duke, and I have a question about time-critical decision making. And so we can consider any outcome. And if we can assume that a health care provider needs x amount of time to read the relevant literature to make their own decision, when you're building an algorithm, you can increase the prediction window to be greater than that x amount of time.

Are there specific recommendations you would have for the way that you train algorithms to make sure that the decision making is beyond what time critical would be? And is there any detail you can provide

on what-- if there's a number of hours, number of days, something that you're thinking about as defining time critical?

CDR Kim Piermatteo: Thank you, Mark. Brendan, would you like to provide a response?

Brendan O'Leary: Sure. Thank you for the question. You know, the final policy does not identify one single cut off for what constitutes time critical and what doesn't. Because as you've highlighted in your question, it can really depend on the clinical condition, where in the health care system you are, and a number of other factors.

That said, the document does provide some really good discussion about the issue of automation bias, as my colleagues mentioned earlier, and other factors that can come into play when determining whether or not something is a time-critical decision.

Another thing that I think is important to keep in mind here is the difference between recommendations and alarms and alerts, and I think that it's important to-- if you're targeting a non-device CDS path, it's important to make sure that the outputs that you're providing are really within the realm of recommendations as opposed to being a specific diagnosis or being an alarm or an alert about a life-threatening condition, for example.

Mark Sendak: Thank you.

CDR Kim Piermatteo: Alright. Our next question is coming from Jason. Jason, I've unmuted your line. Please unmute yourself and ask your question.

Jason Brooke: Excellent. Thank you. This is Jason Brooke. I appreciate the opportunity to hear you guys explain the guidance, and ask questions. The portion of the guidance that I think I found the most confusing was the discussion around how medical device data can satisfy both criterion 1 and criterion 2. And in that discussion, the guidance speaks to a single discrete test or measurement result of a medical device data would be medical information rather than signal or pattern, which, to a certain extent, make sense.

What the guidance doesn't really articulate, though, is when does a series of single discrete tests or tests or measurement results become a signal or pattern? If we look at some of the examples, there's quite a diverse set of scenarios, whether it's five sequential measurements of RR interval, glucose measurements every 30 minutes, heart rate every hour.

So the question I have really is exactly that. When does a set of single discrete tests or measurement results become a signal or a pattern rather than medical information that would satisfy criteria 1 and 2?

CDR Kim Piermatteo: Thank you Jason, for that question. I'm going to turn it over to Brendan to provide your response.

Brendan O'Leary: Thank you. Yes. And thanks for that question. Similar to the last one, this is a situation where the final guidance doesn't identify one single cutoff. Again, because FDA recognizes that there can be a variety of situations that we're encountering and diseases and conditions where this factors in.

That said, we do have a nice rule of thumb, if you will, in the guidance document that can be useful to developers as you're considering this issue. And that gets into what constitutes the type of information that health care professionals might be able to exchange and might normally exchange in the context of a hallway conversation, for example.

It would be pretty regular or normal for somebody to be saying, hey, this patient's blood pressure is increasing. But it would be a little bit less normal in the context of that kind of conversation for hundreds or thousands or any large number of blood pressure measurements to be analyzed. And so I think that can be a useful starting point when having this-- when considering this criteria.

And really, again, would plug the digital health at DigitalHealth@fda.hhs.gov mailbox if you have any questions about a specific situation.

Jason Brooke: Thank you very much. Appreciate it.

CDR Kim Piermatteo: Thanks, Brendan. Alright, our next question is coming from Senan. Senan, I have unmuted your line. Please unmute yourself and ask your question.

Senan Ebrahim: Hi. This is Senan Ebrahim from Harvard. I wanted to ask about calculating a risk score that's specifically been defined by an individual clinician user, say, using my own clinical judgment, potentially including a machine learning model that I built for a non-time-sensitive, time-critical decision using low frequency data. Would that constitute a device? Thanks so much.

CDR Kim Piermatteo: Thank you for that question. I'm looking-- I think our panelists-- we're going to have Brendan take this one again.

Brendan O'Leary: Thank you for that question. There's a couple of considerations you've raised here. Probably the first one that I would look at is in the device software functions and mobile medical applications guidance, rather than the clinical decision support, guidance. And in that guidance, we talk a little bit about when an individual clinician offers something for use in their individual clinical practice or in their practice group.

And so if that's the sort of situation that you're facing, that may be a relevant policy to take a look at before you even look to clinical decision support. Beyond that, I think you are touching on a couple of challenging issues, and it might get a little bit specific. And so I would encourage you to reach out to DigitalHealth@fda.hhs.gov if you want to go over a specific example of a product that you're developing.

CDR Kim Piermatteo: Thanks, Brendan. Alright. Our next question is coming from Timothy. Timothy, I've unmuted your line. Please unmute yourself and ask your question. Timothy, are you there? Are you able to unmute your line?

OK. Hearing no response, we'll go ahead and move on to the next caller. The next stakeholder is Belinda. Belinda, I have unmuted your line. Please unmute yourself and ask your question.

Belinda Smith: Yes. I have an IRB-related question. I was wondering if FDA could confirm that-- excuse me, hang on one second-- that a study with a product that does meet the CDS exclusion criteria would be considered not FDA regulated under 50 and 56 it's not considered a medical device. Of course, the IRB would review under HHS regulations only, but I wanted to just check with that.

CDR Kim Piermatteo: Thank you, Belinda. Alright, we're looking to—Sonja, would you like to take that question? Or Brendan, I apologize.

Brendan O'Leary: Sure.

Sonja Fulmer: Thanks. No, go ahead, Brendan.

Brendan O'Leary: No, please.

CDR Kim Piermatteo: I think we're going to go—

Sonja Fulmer: Sorry, I'll—

CDR Kim Piermatteo: --with you Sonja. Go ahead.

Sonja Fulmer: --to cough. Excuse me. So that is a good question. I think that there's a few different resources available for you to understand what kinds of regulations might apply to the study of software functions that are devices versus those that are not devices. And this is another good opportunity to reach out to the digital health inbox. And we'd be happy to point you to those.

And I'd just check real quick see if I wanted to add anything too since he was also coming off mute.

Brendan O'Leary: I was just going to say that this guidance document is about what is and is not a medical device. And I think that those questions about investigational devices and other study questions are probably questions that are a little bit beyond the scope of what we can get into depth on today.

CDR Kim Piermatteo: Thank you, Sonja and Brendan, for that response. Alright. Our next question. I'm going to circle back to Timothy Johnson. I have unmuted your line. Please unmute yourself and ask your question. Timothy, we are still not able to hear you, so I'm going to give you one more opportunity to try to unmute your line.

Alright. We're going to move on to our next caller. Our next caller is Orest. Orest, I've unmuted your line. Please unmute yourself and ask your question.

Orest Boyko: Well, thank you. It's Orest Boyko. And I just wanted to get any insights where you mention that it was important for clinicians to be trained. Any thoughts or guidance in the future of where clinicians go to get this training?

CDR Kim Piermatteo: Thank you for that question. I'm going to refer again to Brendan to provide a response.

Brendan O'Leary: Yes. Thank you. And I think that, again, this guidance document is about what is and what is not a medical device. And really, the fourth criteria in the clinical decision support software portion of the statute speaks a bit to what it means for a software function to provide information about the basis of its output so that it's not the intent that the clinician rely primarily on that output.

And if you have a more specialized expert, that can be a little bit-- if your software function is intended for an expert with more specialized expertise, then I think that the interpretation of that criterion in that situation can be a little bit different than if the software function is intended for use by a broader group of health care professionals.

And I think that that's really what the guidance speaks to. FDA does not regulate the practice of medicine, and really is not involved in the initial training of clinicians. And so that's probably not something that we would want to speak to here.

Orest Boyko: Thank you.

CDR Kim Piermatteo: Thank you for that question and response. Our next question is coming from Elaine. Elaine, I have unmuted your line. Please unmute yourself and ask your question.

Elaine Tseng: Thank you. Appreciate your webinar. And my question relates to criterion 3. And are there considerations around the idea of a list of options? Are there considerations around the number of treatment options that should be included? And does an express recommendation to do nothing constitute potentially an option, a treatment option?

CDR Kim Piermatteo: Thank you for that question, Elaine. Jessica, I'm going to turn it over to you to provide a response.

Jessica Paulsen: Sure. Thanks, Kim. And thanks for the question. Yeah. So in criterion 3, it really focuses on providing recommendations to the HCP and that-- like we talked about and you highlighted-- can really look like a list of options. We don't specify a number of options. That might be a couple, a few. It could be a whole lengthy list. Maybe it's prioritized. That's really what the intent there is as we interpret it when it means recommendations. I think an output saying do nothing, if that's on its own, sounds to me like a single specific output or directive. But yeah, I hope that helps address your question.

Elaine Tseng: OK. Thank you.

CDR Kim Piermatteo: Thank you both. Alright, our next question is coming from Brandon. Brandon, I have unmuted your line. Please unmute yourself and ask your question. Brandon, if you're speaking, we are unable to hear you. Please make sure you unmute your line.

Brandon: Sorry, how about now?

CDR Kim Piermatteo: Yes, we can.

Brandon: Thank you. This is Brandon Blakely from Washington, DC. My question was about criterion 4. You have a lot of information you would like shared, especially regarding things like machine learning algorithms. I was wondering what kind of level of detail you would recommend be shared such that it still meets a criterion but it's still practical. Thank you.

CDR Kim Piermatteo: Thank you Brandon, for that question. I'm going to refer to Brendan to provide a response.

Brendan O'Leary: Thank you, Kim. And yes, thanks for the question. I think in this case, I would take a closer look at some of the examples that we provide in the guidance document that really get into what types of information could be useful for different types of machine learning techniques that are used and can still be accessible and useful for the end user.

To your point about overburden, though, it is very important that the information provided be targeted towards the intended user of the product and understandable by that user so that it's not the intent that they rely primarily on the output alone, but instead can understand the basis for the result. And in the case of a machine learning algorithm, that can include information about how that algorithm is trained, the context in which it was trained, the populations, and beyond.

And so, again, I think there are several examples in the guidance document that provide a roadmap that you can follow for machine learning algorithms.

Brandon: OK. Thank you.

CDR Kim Piermatteo: Thank you, Brendan. Thank you, Brandon. Our next question is coming from Sean. Sean, I have unmuted your line. Please unmute yourself and ask your question.

Sean Amato: Hi there. Sean Amato. I'm calling in from Southern California. Yeah. My question actually pertains to criteria 3 in terms of decision bias. So if you have an algorithm that basically matches a therapy, whether it be off-label or on-label use, based on the results of a companion diagnostic, would you consider that a medical device or not?

CDR Kim Piermatteo: Thank you, Sean. Brendan, would you like to take this one?

Brendan O'Leary: Yeah. I think we are getting a little product specific here, and it would probably be helpful to know a bit more about what the specific output you're talking about is, the companion diagnostic in question. This one raises some interesting challenges, and I think that it's a good example where we can work with you at DigitalHealth@fda.hhs.gov to really understand the specifics and help you navigate this policy.

Sean Amato: OK. Thank you.

CDR Kim Piermatteo: Thanks, Brendan. That is definitely a valuable resource, that email. Alright, our next question is coming from KMullen. I have unmuted your line. Please unmute yourself and ask your question.

KMullin (Karen): Thank you. Hi, my name is Karen, and my question is about the difference-- or how it would be handled, medical information versus a pattern. I understand that a pattern would be considered a device.

But if both medical information, one data point or a pattern, identifies a value that comes from a cleared medical device and is displayed for the user-- like I said, it's outside the normal range and, therefore, notifies an expert about that abnormal range, and then that expert reaches out to-- it triggers the expert to reach out to the user.

In both instances, because it triggers that expert to reach out to the user to discuss the value and their options-- would it be handled the same way or different way because of the fact that it's either medical information, i.e. one data point or a pattern of data.

CDR Kim Piermatteo: Thanks, Karen. I think we're going to turn to Brendan to get some clarification and provide you a response.

Brendan O'Leary: Sure. Yeah. Thank you for the question. I'm not sure I fully grasp the example that's being raised. However, I think that there are a couple of examples in the document that may be useful in this case. First, there's a section of the guidance that includes a number of product types that were in the FDASIA Health IT report back in 2014, and just really explicitly identifies FDA's approach to those products. And they were in the webinar slides as well. And so I would refer you to that section.

From there, identifying an out-of-range value is something that could be done in the context of a laboratory situation. It's something that could be done in the context of a bedside monitor for something going on with the patient's heart, for example.

And so it sounds like something that might be trending a little bit more towards alarm or alert, depending on the context of use. But again, if you have a specific product that you have under development that you're wondering where it falls within the myriad of policies that FDA's developed, one, you can take a look at that policy navigator, which provides a seven-step process for understanding which FDA policies are going to be most relevant to your product.

And two, you know as you're going through that, that will even provide links to that digital health mailbox where we can help you if you get stuck on any question. And so this is one where I'd really encourage you, one, to take a look at the FDASIA Health IT examples that are in the guidance document and then, two, to run through the policy navigator with and reach out to us if you have any questions.

CDR Kim Piermatteo: Thanks, Brendan. Alright, we're going to move on. Our next question is coming from Priyanka. Priyanka, I've unmuted your line. Please unmute yourself and ask your question.

Priyanka: Hi, everyone. So thank you for this opportunity. I was wondering if you can provide some clarification around time criticality element. So for example, would you be considering the software that is used in emergency department-- would that fall under a time critical element, and hence to be considered as a device?

CDR Kim Piermatteo: Thank you for that question. I am going to refer it again to Brendan to provide a response.

Brendan O'Leary: Sure. Yes. Thank you for the question. And I think that the guidance doesn't preclude the use of non-device clinical decision support software in any particular environment. But certainly an emergency environment or an ER might be an area where you expect to see more things that are going to trend towards that time-critical end of the spectrum where automation bias and other concerns are going to be an issue.

And so here, again, it's going to depend on the specific inputs to the algorithm, the specific outputs that it has, how directive it is, whether or not it's really intending to provide an alarm or an alert as opposed to a recommendation, for example. All of those are going to be at play here. But there is no categorical--

there's no categorical statement in the guidance that something can't be used in that environment and can't still meet the criteria. It just does have to meet the criteria.

CDR Kim Piermatteo: Thanks, Brendan. Thanks, Priyanka, for your question. Our next question is coming from Claudia. Claudia, I have unmuted your line. Please unmute yourself and ask your question.

Claudia: Thank you so much. I just wanted to get some clarification on the guidelines for CDS that are now already in the market but have been reclassified as a device but without any regulatory clearance. Is there a grace period or what's your take on those types of devices?

CDR Kim Piermatteo: Thank you, Claudia. I'm going to look towards Brendan. Again, would you like to provide a response?

Brendan O'Leary: Sure. Thanks. First, I'd like to point out that 21st Century Cures has been the law of the land since 2016, and this final guidance is consistent with FDA's implementation of Cures since then. And so there's no product that was previously not a device that becomes a device as a result of the finalization of this guidance.

On the contrary, this guidance outlines what can be excluded from the device definition under Cures, which really narrowed the device definition back in 2016. And so as far as grace periods and other concerns. If you have a product that you're concerned based on what you've learned in the context of this final guidance, maybe a device that requires FDA review or has other regulatory requirements related to it, the first step is really to open a dialogue and have a conversation with FDA about how we can move forward.

Some developers may choose to use the criteria in the guidance document-- or, excuse me-- the recommendations in the guidance document that help you understand the criteria and the law to really adjust their product and choose that non-device path. And other developers may choose the device path and wish to work with FDA on validating and getting that authorized, and we're happy to work with you as you consider those issues. So again, DigitalHealth@fda.hhs.gov.

CDR Kim Piermatteo: Thanks, Brendan. And thanks, Claudia, for your question. Our next question is coming from David. David, I have unmuted your line. Please unmute yourself and ask your question.

David Bragg: Hi. This is David Bragg calling from Dallas. Appreciate the webinar. It's helpful. My question is back to the alerts and alarms issue. It almost sounds like that anything that contains an alert or an alarm pretty much puts it in the device category. The spectrum of alerts and alarms is huge, though. I mean, an alert could be that the function of the software or monitoring device or whatever is no longer attached. So something as simple as an alert, hey, we're not receiving data, or whatever, all the way to complicated alerts.

So is there interpretive room within that category of alerts and alarms based upon the importance of it or can you comment on that, please?

CDR Kim Piermatteo: Thank you, David, for that question. I am going to turn it over to Brendan to provide you a response.

Brendan O'Leary: Sure. Thank you for the question. I think one thing to consider is—is this a reminder, for example, rather than an alarm or an alert. Certainly, hey, we don't see any evidence that there was an annual flu shot in this patient's records. Maybe consider whether or not this patient could be a good candidate for a flu shot or other preventative measures might be more along the lines of what would be appropriate and would meet the criteria for non-device CDS.

Whereas on the other end of the spectrum, we really do take a platform-independent approach to the regulation of device software functions, meaning that a software algorithm does not need to be built into a hardware medical device in order to meet the device definition and be regulated by FDA. It really has to do with its intended use and its functionality as that relates to its intended use. And so the issue of connection to the acquisition system isn't quite how this works in practice.

CDR Kim Piermatteo: Thank you, Brendan, for that response. Our next question is coming from Adrienne. Adrienne, I have unmuted your line. Please unmute yourself and ask your question.

Adrienne Hammond: Hey. Thank you so much for getting around to me. I had a question about the user. So I noticed that in the older guidance, there's a focus on the user and the HCP. And in this guidance, they really primarily focused on the HCP. So if you have a software function that uses assessments and guides the user through assessments and their user-driven assessments to provide a recommendation and then the outcome could be used by the user or directed towards an HCP-- I mean, where do you guys think that that would fall? And do you have any comments on that? Thank you.

CDR Kim Piermatteo: Thank you, Adrienne. I am going to look to Jessica. Would you like to provide a response to Adrienne?

Jessica Paulsen: Sure. Thanks. Yeah. Thanks, Adrienne. I think if I understand your question correctly, when you're saying user, you might mean patient, is that fair if you're still with me?

Adrienne Hammond: Can you still hear me?

Jessica Paulsen: Yeah.

Adrienne Hammond: What if the person's not necessarily a patient? What if it's just someone that actually is using an application?

Jessica Paulsen: Perfect.

Adrienne Hammond: And it's not really prescribed at all.

Jessica Paulsen: Perfect. Thank you for your question. Yeah. So criterion 3 specifies that the recommendations or support are provided to an HCP. So that's really what we focus on when providing clarity around how to meet criterion 3 is that the intended user of that software function must be an HCP.

And I think I mentioned this in the slides, but there are a number of other guidance documents that might be helpful in those kinds of situations. Again, not speaking to any specific examples, but you might consider walking through our policy navigator to see if there's another guidance document that might help you. But criterion 3 of the statute really focuses on HCP users.

Adrienne Hammond: OK.

CDR Kim Piermatteo: Thank you, Adrienne. Thank you, Jessica. Our next question is coming from Leo. Leo, I have unmuted your line. Please unmute yourself and ask your question.

Leo Espindle: Yeah. I have another follow-up question sort of around the risk calculator, but it's really more around an example of almost definitional calculators or simple calculators. So for instance, I'm thinking of like-- these are ubiquitous online, like a BMI calculator that says if you're a certain weight and a certain height, you're obese. Or if you have a certain blood pressure above 135, you have stage 1 hypertension. Very simple. Or anxiety-- GAD-7. There's a lot of these sort of things online. Would those fall under that simple calculator thing or are those more like diagnostic?

CDR Kim Piermatteo: Thank you, Leo, for that question. Sonja, would you like to provide a response?

Sonja Fulmer: Sure. So, yes, the enforcement discretion that we talked about earlier for simple calculators means that for software functions that provide that sort of simple calculation, including for BMI and other types of calculations that health care providers would be using in their normal practice, that would be a software function that FDA does not focus its oversight on. That's what we mean by enforcement discretion.

So we're not focusing or expecting folks to comply with any regulatory requirements that might apply for those types of software functions. And the MMA guidance-- or the device software functions and mobile medical applications guidance-- that one includes a lot of those examples, including BMI as an example of a simple calculator.

CDR Kim Piermatteo: Thank you, Sonja. And thank you, Leo, for that question. Alright, that wraps up our live Q&A segment. I want to thank you all for a very engaging question and answer segment. And at this time, I'm going to turn it back over to Sonja for her final thoughts for today.

Sonja Fulmer: Thanks, Kim. And thank you all for attending the webinar today and engaging with us during this Q&A session. As we've said before, this final guidance is the last piece of FDA's interpretation of the software provisions of the 21st Century Cures Act, and there are many types of software functions that are excluded from the device definition, as well as others that may meet the device definition but are not the focus of FDA oversight, those enforcement discretion categories that are described in other guidance documents now.

This final CDS guidance does provide clarity on the statutory criteria for what types of decision support software functions are not devices, but it doesn't change FDA's long-standing, final enforcement policies that are in existing guidance.

And because these policies are spread over several different guidance documents, we've mentioned a few times now that we've issued this Digital Health Policy Navigator to guide you through these different policies. And we hope that you use it. But if you find that after this webinar or after using that tool you have any additional questions, please don't hesitate to reach out to us at DigitalHealth@fda.hhs.gov to seek feedback on the regulatory status of a product we're developing or any other questions you might have about digital health policies. So thank you again for joining us today.

CDR Kim Piermatteo: Thanks, Sonja, for those final thoughts and for your presentation today on this final guidance. And I'd also like to, again, thank Jessica and Brendan for their participation today.

For our audience members, please remember, printable slides of today's presentation are currently available on CDRH Learn at the link provided on this slide under the section titled Specialty Technical Topics and the sub-section Digital Health. A recording of today's webinar and a transcript will be posted to CDRH learn under the same section and subsection in the next few weeks. A screenshot of where you can find the final presentation materials has been provided on this slide.

If you have additional questions about today's webinar specifically, please email us at dice@fda.hhs.gov.

We also encourage you to attend a future CDRH webinar. A listing of all of our upcoming webinars is available at www.fda.gov/cdrhwebinar. And with that, this concludes today's CDRH webinar. Again, thank you all for joining us today. And have a nice day.

<div align="center">

****

END

</div>