UNITED STATES OF AMERICA

NUCLEAR REGULATORY COMMISSION

\+ + + + +

34TH REGULATORY INFORMATION CONFERENCE (RIC)

\+ + + + +

TECHNICAL SESSION - W12

FUTURE-FOCUSED RESEARCH -
WHAT WILL WE THINK OF NEXT?

\+ + + + +

WEDNESDAY,

MARCH 9, 2022

\+ + + + +

The Technical Session met via Video-Teleconference, at 1:00 p.m. EST, Raymond Furstenau, Director, Office of Nuclear Regulatory Research, presiding.

PRESENT:

RAYMOND FURSTENAU, Director, Office of Nuclear Regulatory Research, NRC

ANGIE BUFORD, Member, Digital Twin Project Team; Branch Chief, Reactor Vessels and Internals Branch, NNR/NRC

STEPHANIE BUSH-GODDARD, Senior Health Physicist, Radiation Protection Branch, Division of

Systems Analysis, RES/NRC

JESSE CARLSON, Digital Twin Project Team; NRR/NRC

CANDACE DE MESSIERES, Member, Digital Twin Project Team; Chief, Advanced Reactor Technical Branch 2, NRR/NRC

DOUG ESKINS, Member, Digital Twin Project Team; Reactor Engineer, NRC

AMY HULL, Senior Materials Engineer, Materials Engineering Branch, Division of Engineering, RES/NRC

MATTHEW HUMBERSTONE, Senior Reliability Engineer, Performance and Reliability Branch, Division of Risk Analysis, RES/NRC

RAJ IYENGAR, Digital Twin Project Team; Chief, Reactor Engineering Branch, NRR/NRC

JEANNE JOHNSTON, Digital Twin Project Team; Chief, Electrical and I&C Long Term Operations and Modernization Branch, NRR/NRC

DANIEL JU, Digital Twin Project Team; NRC

ANYA KIM, Computer Scientist, Instrumentation Controls & Electrical Engineering Branch, Division

KIM LARSON-JENKINS, IT Specialist (Cyber), Cyber Security Branch, Division of Physical and

Cyber Security Policy, NSIR/NRC

SHAH MALIK, Senior Materials Engineer, Materials Engineering Branch, Division of Engineering, RES/NRC

JAMES STECKEL, Program Manager, Regulatory Guide and Programs Management Branch, Division of Engineering, RES/NRC

VAIBHAV YADAV, Digital Twin Project Team; Research and Development Scientist, Idaho National Laboratory

CONTENTS

Page

Opening remarks....................................6

Presentations:

I. Digital Twins, the Regulatory Viability (Video)

Jesse Carlson................................9

Angie Buford...............................10

Daniel JU..................................12

Candace De Messieres.......................13

Vaibhav Yadav..............................14

Jeanne Johnston............................16

Doug Eskins................................18

Raj Iyengar................................20

Questions..................................21

II. Licensing Modernization Project Methodology on an Operating Reactor

Matt Humberstone...........................23

Questions..................................35

III. Drones and Virtual Reality Tools to Analyze Radiological Surveys in Decommissioning

Stephanie Bush-Goddard.....................38

Questions..................................50

IV. Review of Advance Manufacturing Technologies

Amy Hull...................................53

CONTENTS (CONT.)

Page

Questions....................................63

V. Zero Trust Paradigm for Cybersecurity in New Reactors

Anya Kim.....................................66

Questions....................................79

P-R-O-C-E-E-D-I-N-G-S

1:00 p.m.

MR. FURSTENAU: Welcome to our afternoon technical sessions. The session we're in today is on Future-Focused Research - What Will We Think of Next? I am Ray Furstenau, the Director of the Office of Nuclear Regulatory Research, and I'm really excited to share this session.

I wanted to give you a brief history of future-focused research, because it is a new program here at the NRC. It was supported by the Commission in mid-FY20. And we got started on some projects in FY20, and were funded again in FY21, as well as FY22.

And today you'll hear presentations and be able to ask questions of folks in the NRC that have been working on some of these projects throughout the year. So they'll be at different stages. Some were started in '20 and are finishing up. Some in '21 and some in '22.

And what I'm excited about with this program, this gives us a chance to look ahead, look beyond what we normally see in the Office of Research where we're responding to user need requests from the business lines. Which is very important. That's our bread and butter. That's what we should be doing.

But this gives us a chance to get good ideas from not just staff and research, but throughout the agency on what should we be looking at with regard to new advanced technologies that have applications in nuclear. And what should we be doing as a regulatory to get ready for those new technologies so we're not the holdup when a licensee comes in and says, hey, I want to do this and we say, well, we haven't even seen that before. We're trying to stay at least even with the game, or even ahead.

And so what we're -- the way that we're going to do the session is we'll have a presentation from each of the projects. We have five projects today. And I'll introduce the speaker and topic area. And then at the end of each presentation we'll have a little bit of time for questions and answers. Then we'll go to the next speaker, question and answers after that and so on until we're finished with the session.

So with that, oh, I forgot. I also want to thank Jim Steckel. Jim, you did a great job organizing and coordinating everything. And of course to all the speakers and those of you who have been working on our future-focused research projects.

So with that, the first project I wanted to introduce is called Digital Twins, the Regulatory Viability. And I'm introducing, not a speaker, but a team here.

The digital project team has been led by our Reactor Engineering Branch in the Office of Research. And includes staff from across the agency, as well as some technical experts from the Idaho and Oak Ridge National Labs.

And our reactor engineering branch plans and executes programs to develop technical bases, methods and tools for advance technology applications that influence the safety of operating, and advance reactors in other facilities we regulate.

The digital twin project was initiated in 2020 to study the potential nuclear domain applications of digital twin technology and the infrastructure needed to regulate such applications. The project wraps up this year with a publication of two final technical reports.

And with that I'll ask the technician to start the video slide show. And then we'll conclude with some remarks from Dr. Iyengar.

(Video playback started.)

MR. CARLSON: Hello and welcome to our digital twin future-focused research project. My name is Jesse and I'll be kicking things off today. So let's begin with a question. What exactly is a digital twin?

Well, that's a great question. I'm glad you asked. For our project we describe a digital twin as a part of a nuclear digital twin system with four basic parts.

The first part is the real-world or physical plant. The physical plant includes all the plant structures, systems, and components as well as associated physical phenomena, procedures, and even the human actions needed to safely produce nuclear power.

The digital twin is the virtual representation of the physical plant. Depending on needs, a variety of plant systems can be twinned, from single components to an entire plant, and two categories of technologies are needed.

Data and information management gives a digital twin the ability to gather, process and disseminate data. And modeling and simulation makes that data useful through advance technologies such as data analytics, artificial intelligence, and data informed models.

Of course, for the digital twin to maintain accuracy, otherwise known as state concurrence, it must constantly communicate with the physical plant. Real-time plant data and performance is transmitted to the digital twin to maintain this concurrence.

The information is processed and returned to the physical plant as actions and recommendations, such as diagnostics, operational guidance and even control signals.

Together these four parts form a nuclear digital twin system enabling novel, real-time plant state awareness that can improve both plant efficiency and safety. More on the potential applications of a nuclear digital twin in the next segment.

MS. BUFORD: Hi, my name is Angie Buford and I'm the branch chief of the reactor vessels and internals branch in the NRCs Office of Nuclear Reactor Regulation.

And I'll just briefly explore why there is so much interest in digital twins among nuclear stakeholders. From the stakeholders point of view digital twins can offer significant opportunities in design, testing, operations and maintenance.

Some of these opportunities include faster design, testing and validation of reactor components and systems, reduction of cost using condition based and predicted maintenance, improvement of operational efficiency while maintaining safety with real-time plant issue identification, diagnosis, and operational recommendations. And reduced staffing due to improved system control, increased awareness of plant state and automation.

Of joint interests to both licensees and the NRC is the potential for digital twins to improve the efficiency of regulatory processes while maintaining or even improving safety margins. A digital twin can provide a single, trusted source of always up-to-date plant information, which could facilitate such things as a common platform for plant validation, verification and uncertainty quantification, a shared framework for acceptance and approval of licensing actions, improved communication between licensees and the NRC, increased regulatory responsiveness by identifying safety significant issues in real-time, and improved inspection efficiency by enhancing NRC activities to prepare and gather information.

Given these potential opportunities, interests in nuclear digital twin applications is understandable, and this interest has inspired the NRC's digital twin future-focused research project discussed next.

MR. JU: NRC's digital twin future-focused research project was developed to identify and evaluate the regulatory impacts of a nuclear digital twin system, such as technical issues and gaps and needed regulatory infrastructure.

To that end, we have published several technical reports and information letters to help prepared the NRC for digital twin. We also held two widely-attended workshops on digital twin applications. Here are some of the workshop participants.

As you can see, there is a significant interest in digital twins among nuclear stakeholders.

And the insights and activities discussed in the workshops have been captured in two NRC information letters.

DR. DE MESSIERES: Good afternoon. My name is Candace and I hope you're having a great experience here at the RIC.

The team also researched various technical aspects of digital twins and published technical reports assessing the current state of digital twin technology, as well as the technical challenges and gaps for these enabling technologies.

One of the things we have learned is that digital twin technology is rapidly evolving, and nuclear stakeholders are already beginning to implement it.

To give you an idea, here is some applications discussed in the state of technology report. As you can see, digital twin technology is already being applied to improve design, testing, validation, and operation and maintenance efficiency including several projects that are sponsored by the Department of Energy's ARPA-E GEMINA program.

For example, X-energy is using an immersive digital twin environment during design to optimize plant operations and maintenance, and identify where cost reductions can be safely achieved.

Kairos is using digital twins for iterative development to directly reduce uncertainty and cost overruns.

We have covered a lot of information in a short time, but if you would like to know more, both the technical reports and information letters are publicly available on the NRC website. This project has been a collaboration between the NRC and Idaho and Oak Ridge National Labs.

We also coordinated closely with the Department of Energy's ARPA-E and the Electric Power Research Institute, or EPRI, to share information and co-sponsor workshops.

While the future-focused phase of this project will be completed soon, digital twin research continues. We will discuss continuing research later, but let's discuss what we're learned so far.

DR. YADAV: Hi, I am Vaibhav. One of the things our team identified early was the importance of certain technologies to the implementation of a nuclear digital twin.

These enabling technologies are the key to the feasibility of digital twin and must be successfully developed, matured, and regulated. We identified the following enabling technologies.

Advance sensors. Now, advance sensors will be required to measure and communicate the large volumes and new types of plant data required by the digital twin.

These large volumes of heterogeneous data flowing from the plant will also require new analytic and information management. Analytic and information management is an enabling technology that's needed to capture, correct, categorize, integrate, protect and process this data into useful forms.

Machine learning and artificial intelligence will be leveraged to characterize and model complex plant systems. As well as produce predictions, recommendations and even plant control signals.

Advanced and multi-physics models. Now, multi-physics models and advanced models will be needed to represent complex plant systems and processes, as well as their interactions in real-time and with high fidelity.

Finally, the key to maintaining concurrence with the physical plant is the use of data informed models. Such models will use real-time plant data to update and correct their performance while running.

Of course, challenges exist to the implementation of these enabling technologies. And they will be discussed next.

MS. JOHNSTON: So far we have presented an overview of digital twins, the opportunities for their use in nuclear, and some of the key enabling technologies. Now we will move on to discuss challenges and gaps that exist for digital twin implementation.

We have broadly categorized the challenges into three groups. Challenges with advance sensors and instrumentation. Challenges associated with modeling and simulation. And challenges associated with data and information management.

The path to a nuclear power plant digital twin is one that will see a more fully instrumented plant. And also introduces challenges with advance sensors and instrumentation.

There will be a greater number of sensors, including more varied sensors with new technology, and even multi-modal sensors, to feed the digital twin the data that it needs. These sensors will supply real-time data to models and simulations.

Other considerations include optimization of sensor selection and placement, infrastructure for real-time sensor data collection and sensor maintenance.

Modeling and simulation challenges arise from the complex and interdependent systems that form a nuclear power plant digital twin. The digital twin will integrate real-time, heterogeneous data into similarly real-time heterogeneous physics-based and data-informed models.

This further includes selections of machine learning models, scaling data analytics, computational efficiency and quantification and propagation of uncertainty.

The challenges related to data and information management encompass both physical and procedural requirements for digital twin implementation. The specific nuclear power plant digital twin physical infrastructure requirements for data storage, data sharing and computation are not yet well defined.

Additionally, nuclear specific standards and guidance are likely to develop regarding technologies like cybersecurity, redundancy and encryption. The implementation of a digital twin will also see the transition from a document-centric approach to one that is data-centric.

While digital twins offer great opportunities, additional effort is needed to meet the challenges of implementation and regulation.

DR. ESKINS: Hi, I'm Doug, and it's my pleasure to ramp up this presentation. Let me start by summarizing some key project take-aways.

As we've mentioned previously, there is substantial stakeholder interests in digital twins, and the technology is rapidly developing. Already we have identified several likely nuclear use cases, such as maintenance optimization, real-time radiation mapping, and even the use of digital twins to detect and diagnosis plant issues.

Also, and this is my favorite one, the digital twin offers great promise as a novel source of real-time, trusted information. Information that can inform activities like plant design, testing, operations, and even regulation.

Additionally, there is a growing need for collaboration to share information, develop common solutions, and establish a community of practice among nuclear digital twin stakeholders. Of course, actual implementation of digital twin will require advancements in several key enabling technologies. Some of the most significant among these are data-informed and real-time multi-physics modeling and simulation.

And finally, while we are near the end of our future-focused research, the project's end is really only the beginning of the NRC's digital twin research. So what's next?

The NRC is busy working on three new reports and they'll be published soon, as well as the summary report for this project.

The new reports will explore regulatory considerations and opportunities, advanced sensors and instruments, and the safeguards and security issues associated with digital twins.

As this new research develops, the NRC

will continue to identify and address significant digital twin regulatory issues. Our aim is to engage early and to encourage discussion centered on common frameworks.

In closing, while much work remains, our future-focused research will successfully conclude this spring. We hope to have started a lasting regulatory conversation aimed at preparing both the NRC and nuclear stakeholders for the coming digital twin world.

Thank you for the opportunity share our work, and have a great rest of your day.

(Video playback stopped.)

MR. FURSTENAU: Well, thanks, everybody. That was really a great video and slide. It really gave an excellent overview. I think we are going to close the presentation here.

I think that Dr. Iyengar, he's the branch chief of our reactor and engineering branch, wanted to close the presentation with a few words. Raj? And then we'll get to questions.

DR. IYENGAR: Sure. Thank you so much, Ray. And so the story of growth in the future-focused research program.

So a couple of years ago when we initiated this, before ARPA-E funded the $14 million grants to several vendors in national labs, little did he think that the future is here and now. And here we are.

And we are so glad that we started preparing for it early enough to underscore the very theme of this RIC. Thank you so much to the Staff in putting in a lot of energy and enthusiasm, as you can see, into creating this successful program.

We owe this to the vision of Ray Furstenau because he was the one who conceived this. We did have a small seed, not a big one, not a acorn, a tiny one. And I think we have realized the vision of this program. And thank you so much.

MR. FURSTENAU: Well thanks, Raj. We do have time for one question here for you. How do we know that the digital twin is actually a twin?

That is, how do we know it's accurate enough to represent the plant and not missing any important subtleties?

MR. IYENGAR: That is an excellent question. And that's what we are hoping these enabling technical advances with enabling technologies will ensure.

It is not, although the digital twin might be similar to a simulator, you can have many simulators but it's not a twin. So in order to make that happen we need to be able to capture real-time data. And that's where the data analytics became important using advance sensors, and keep it concurrent with the real time.

That's what some of our folks call it state concurrence. It's a fancy term. But it just says that it's just real-time.

It's like looking at myself or looking at yourself in a mirror. And that's exactly what we will do. And the growth of the advancement of enabling technologies will help the industry achieve that.

MR. FURSTENAU: Okay. Well thanks, Raj. Unfortunately, that's all the time we have for questions right now on this and so, again, thanks to you and your team for this, the progress you made on this wonderful project so thank you.

MR. IYENGAR: Thank you.

MR. FURSTENAU: Next we're going to go to the next speaker to talk about the project he's leading on applying the Licensing Modernization Project Methodology on an Operating Reactor. Matt Humberstone is a Senior Reliability and Risk Analyst in the Division of Risk Analyst in our Office of Research here at the NRC.

The branch he's in uses risk assessments and insights to support a broad range of regulatory applications. And Dr. Humberstone's experience includes work on advance monitoring, diagnostics and prognostics for advance reactor and designs.

He also worked at Sandia National Labs in the advance reactor concepts division while he was in school. Dr. Humberstone received his bachelor's degree in engineering physics from New Mexico State University. A master's in statistics from the University of Tennessee. And a master's degree and PhD in nuclear engineering from the University of Tennessee.

And with that, Matt, I'll turn it over to you to talk to us about the application of the LMP methodology to an operating reactor. Matt.

DR. HUMBERSTONE: Great, thank you, Ray. I'm waiting for the slides to pop up. Okay, so, Spencer, are the slides going to pop up? Just a quick question. Oh, there it is. Okay, great. Just a little bit of a delay there.

Great. Well thanks, Ray. And hi everyone. My name is Matt Humberstone. I'm going to be talking about the licensing modernization project for operating reactors today.

This is one of several future-focused research projects that were approved back in 2020. And before I get started I definitely want to acknowledge Keith Compton and Trey Hathaway from the NRC and Kurt Vedros from Idaho National Labs for their support with this project.

Next slide please. So the overview of this project is really just looking for ways to leverage the NRC's Level 3 PRA results to look at how we can leverage the LMP methodology for light water reactors.

So the LMP methodology is a licensing approach for non-light water reactors. And it's described in NEI 18-04 and endorsed by the NRC's Regulatory Guide 1.233.

The LMP approach uses risk metrics that come out of a Level 3 PRA model. The NRC has been working on developing a Level 3 PRA model for several years now. So we thought this was a great opportunity to use those results to test the methodology and see what kind of insights we can get out of it.

When we first started this project the Level 3 model results we had available were for internal events only, so we started with that. And now that's expanded to include external events and other model enhancements. So now we're starting to dig into that. And throughout this presentation today I'll discuss some of the results of both of those.

Next slide please. So this project is a two phase project. Phase one we wrapped up about a year ago. And that used those initial results out of the Level 3 PRA model that was for internal events only.

Now we're moving on to Phase 2, which is these enhanced Level 3 PRA results. And we're also digging deeper into some of the other aspects, such as how uncertainty impacts the results and things like that.

Both Phase 1 and Phase 2 are part of the original scope of this project. But outside of that we're also looking at ways that some of these tools and insights can help support the NRC and their future needs. So we've been having those discussions.

Next slide please. So there are several objectives behind this project. First we wanted to look at the LMP's feasibility for light water reactors. We wanted to do that by piloting the LMP, using the NRC's Level 3 PRA model results.

While we are doing that we want to identify any challenges that come up and any issues that we have with it. We also want to see if there are any new risk insights on light water reactors that we can pull from this.

And then we're continuing to prioritize communication throughout this project by both sharing the progress and results with internal and external stakeholders.

Next slide please. So the technical approach is pretty simple. On the left here we have the NRC's Level 3 PRA model, which has decades of severe accident research behind it and light water reactor operating experience. That's all baked into that modeling.

And then on the right we have the LMP

methodology. Which is this risk-informed performance-based approach for non-light water reactor licensing. We're trying to merge those two together and see what kind of insights fall out of it.

Shown here in the middle is the LMP's frequency consequence curve. This is really the backbone of the LMP approach. And I have a slide later on where I'll go more into detail on how that's used.

Next slide please. So before we get into that though I want to discuss a little bit about NRC's Level 3 PRA modeling process. So this slide shows how the process starts with the Level 1 PRA and progresses out to the Level 3.

The Level 1 PRA is shown in red here. And that's the type of PRA model that both the NRC risk analysts and industry risk analysts are most familiar with. It starts with initiating events and progresses out to your different core damage states.

So that resulted in about 50,000 cut sets, which were then shed into the Level 2 portion of the model which narrowed those down by plant damage states. And then that ultimately was turned into 16 different release categories.

These 16 release categories were then feed into the Level 3 portion of the model where we modeled the release, dispersion and consequences associated with those different accidents.

Next slide please. So the NRC analysts used the MELCOR Accident Consequence Code System, the MACCS code, to do the assessment. The Level 3 portion. The MACCS code is a powerful tool with lots of different capabilities. We used it to generate the dose at the exclusion area boundary, which is the risk metric that's used in the LMP.

Next slide please. So we have those 16 release categories that come out of the Level 3 PRA model. And what we did is we broke them up into three separate groups.

We took the five release categories with the largest doses and we put them in the orange group. We took the five release categories with the lowest doses and put those in the blue group. And then we took the six release categories in the middle and we put them in the gray group.

There is a couple reasons why we did this. One was, it made the results easier to understand visually. And then the second reason is that the Level 3 PRA model results are non-public, so we can't share those directly with you, so this was the great way that we can still get some of the insights out there and get some of the discussion going.

Next slide please. So there are several different ways you can look at the release categories and compare them. Again, the largest dose group is in orange, the low dose group is in blue and the middle dose group is in gray.

If you look at the pie chart in the middle of this slide here, this is the contribution of these different groups to the overall consequence. So the largest dose group, as expected, dominates at about 90 percent.

If someone was to use this information to inform a review or determine where to focus their attention. It might ignore the low dose group because they say it doesn't even show up on this chart, it's less than .5 percent of the contribution to consequences.

But then if you move over and look at the frequency contributions you get a slightly different picture. You see that while the largest dose group, even though they dominate the consequences, you're much less likely to see one of those. Like one percent chance. So maybe we should focus our attention on the middle dose group and low dose groups since they dominate the frequency contribution.

So in reality we want to take both the consequence and the frequency into account. And that's taking a risk-informed approach to it.

A great took to do this is the LMP's frequency consequence curve. Next slide please. So here is the LMP's frequency consequence curve.

On the Y axis is the frequency, on the X axis is the dose at the exclusion area boundary. Going diagonally across this chart is the target line.

So for non-light water reactors and reactor designers their design objective is to push their accidents down towards the lower left. So they want to reduce the risk significance of their accident families.

As those accidents get closer to the target line and get above it, they become more risk significant. And for us that would be an indication that we should look at those a little closer.

So the LMP uses this curve to help the selection of licensing basis events, safety classification of SSCs and determination of defense-in-depth adequacy.

So we're not using, in this project, to do those three things since this is for an operating reactor technology. But we are using this powerful tool to see what kind of risk insights we can pull out and what we can learn on both the method and on light water reactors.

Next slide please. So when you take the Phase 1 results, remember, the Phase 1 results are the Level 3 PRA model results for internal events, and you put them onto the frequency consequence curve you get this slide.

Now these bubbles here represent where the different accident families within those groups falls on the chart. As you can see, all three of these groups have an edge of their bubble that gets close to the target line and an edge that's farther away.

So if an engineer or analyst at the NRC wanted to use this to help focus their attention on the more risk significant accidents, they would find accidents in all three of these groups to look at. And then there is also accidents in all three of these groups that are farther away from the target line and aren't as risk significant.

Next slide please. So that was the Phase 1 results. Then we got the Phase 2 results which were the Level 3 PRA model results that include external events, credit for FLEX, things like that.

And what we did is we took the results from Phase 2 results and we compared them to the Phase 1 results. And we saw that accidents, or different release categories, had different sensitivities to these changes.

So there are certain accidents that were sensitive to external events, so that would drive up some of their risks. There are other accidents that were more sensitive to credit for FLEX, and that would drive down the risks. And then there were several accidents that these changes did not have much of an impact on.

Some of examples of these type of accidents is if you look at a LOCA accident, that happens pretty quickly and so a credit for FLEX doesn't make very much of an impact because you don't have time to use your portable equipment. Where there are lots of offsite power accidents and station blackout accidents, they have a lot more impact, or FLEX has a lot more impact on those accidents and can drive down the risk there.

So out of the 16 different release categories, eight of them saw their risk increased a little bit, two of them saw their risk go down a little bit and six of them remained about the same.

Next slide please. So when you take those new results and you put them onto the frequency consequence curve you get this figure here. Again, we have the original solid bubbles from Phase 1.

Then when you put these new results in you see that it stretches the bubbles out a little bit. So some of the accidents get a little bit closer to the target line and some get a little farther away.

Again, these insights could be very useful for a reviewer or an engineer to be able to look at those accidents that are starting to get close to that target line so they can focus their attention there.

Next slide please. So, in conclusion,

the LMP methodology is a powerful tool and has the potential for use beyond the original intent and for light water reactors.

Some of the insights that we're gaining through this project, we have hope that they'll be able to support our future non-light water reactor licensing reviews. And then these results confirm that NRC's current expectations on operating reactor safety profile. So nothing popped out while we were doing this that really surprised us on that.

And then we're continuing to explore ways to leverage the LMP methods and tools to help support our future work.

Next slide please. So next steps. We're still working on Phase 2 now. So we got the preliminary results that you saw here on these slides, but we're still digging into those a little bit more to understand them.

We're wanting to look at some of the specific accident sequences to see how they change from Phase 1 to Phase 2. And then we're wanting to understand the uncertainty impacts a little bit more. So we're still working on that.

We're continuing to engage with NRC

technical staff on future needs. And looking for ways to leverage some of these tools and insights.

And then we'll continue to prioritize communication, both internally and externally, to share the results from this project. That concludes my presentation. Thanks, Ray.

MR. FURSTENAU: All right, thank you, Matt. I appreciate your presentation. Good stuff. We do have time for a couple of questions.

So the first one, Matt, how were the results from the Level 3 PRA project be expanded from one site to other sites in the fleet? It looks like a Level 3 PRA is redder and extensive composition.

DR. HUMBERSTONE: Yes, so that's a very important point, right?

Like, you can't just take your risk results from one plant and apply it to another. So that's one of the things we're looking at, at how we can leverage some of these insights to other plants even though you, like you just said, you can't just, one plant for one plant, there can be quite a bit of differences there.

So we have some proposals. Some ways we're thinking about doing that, but we're not really in a place yet to share that but we're still thinking through that. But great question.

MR. FURSTENAU: Okay, thanks, Matt. I think one more question. As you were doing the future-focused research project where there, when you were doing that I mean you probably went in with some thought of how the overlay on the LMP curve would appear. Were there any unexpected results or surprises to you?

DR. HUMBERSTONE: So actually, a lot of the results kind of aligned with what our expectation was. Which I guess is, it's a good thing.

But you're always kind of expecting to find something surprising you know. We didn't actually, a lot of the accidents that we expected to be more risk significant, like bypass events and things like that kind of did bubble up to the top.

So I think even though there were no surprises, I think that's definitely a great insight. You know, to kind of confirm some of our other projects we've done in the past that show the risk significance of the different accidents.

MR. FURSTENAU: Great. Well thanks, Matt. I think that's all the time we have for questions, but again, I really appreciate you presenting today on a really, to me, a very fascinating project. So thank you.

DR. HUMBERSTONE: Thank you.

MR. FURSTENAU: Next we'll go to the third future-focused research project. It's going to be presented by Stephanie Bush-Goddard and it's called, Drones and Virtual Reality Tools to Analyze Radiological Surveys in Decommissioning.

Dr. Bush-Goddard is a senior researcher in our Office of Nuclear Regulatory Research. And she provide leadership and project management to the radiation protection computer code analysis and maintenance program. And she leads advance reactor and future-focused research activities of course in this area.

Stephanie has held numerous leadership positions at the agency, including chief of the radiation protection branch, senior health physicist and nuclear engineer. She has more than 30 plus years of professional and educational experience in the material wastes, security and reactor arenas, including a wide range of radiation protection and regulatory framework issues. And performance assessment methods.

Dr. Bush-Goddard has a bachelor degree in mechanical engineering. And hold a master's in PhD and environmental health science with a minor in health physics.

She's an adjunct, was an adjunct assistant professor at the University of Maryland, and is a recipient of graduate level fellowships. With that, I'll turn it over to you, Stephanie. We look forward to your presentation.

MS. BUSH-GODDARD: Thank you, Ray. And thank you for the introduction. And hello to everyone. Can everyone hear me? Okay, great.

So as Ray said, my presentation is about unmanned vehicles for decommissioning, or UAV for short. But I'll probably hear me say drones in most of the presentation.

But before I begin, I want to give a special, special thanks to Pacific Northwest National Laboratory for providing the drone, the test bed facility, the radiation sources and numerous signatures for the drone flights to perform this work. And they did it in a very short period of time.

Next slide please. So, the agenda today is thoughts off of highlighting the major regulations and guidance documents developed to decommission and release NRC sites. Then I'll go into the driver, which really is a proof of concept to compare humans and drones.

The methodology is next. And then pictures of the actual drone flight, which was only four months ago. I'll present some preliminary data and results and potential next steps. This will be followed by final thoughts and recommendations for the future.

Next slide please. So, the major regulatory requirement is in 10 CFR Part 20, Subpart E. This regulation is called the radiological criteria for license termination.

And it states, among other things, that the total effect of dose equivalent to an average member of the critical group should not exceed 25 millirem a year and their radiation levels should be as low as reasonably achievable.

And so the key guidance for demonstrating a facility or site meets this regulation as provided here. One of the guidance documents is NUREG-1575.

It's the multi-agency radiation survey and site investigation manual, or MARSSIM.

Another one is 1507. It's the minimum detectable concentration with typical survey for instruments or various contaminates in field conditions. And trust me, that is a mouthful. And the third one, NUREG-1757, is consolidated decommissioning guidance.

Now, these are major guidance documents, among others, that form the basis of our proof of concept. Because a future goal is to incorporate drone usage into these guidance documents.

Next slide please. So before we start with the methodology, remember, the driver is to do a simple proof of concept.

So we're interested in three research questions. Did the observed UAV, or drone path, differ from the human path? And of course, if so, how much?

Did the survey path deviation impact the survey results?

And were radiological measurements from humans and drone significantly different?

Next slide please. So before we get into the results let's talk a little bit about the drone. The drone is an Aurelia X6 hexacopter, of course with six propellers, and can accommodate a payload of up to 25 pounds with, excuse me, 24 pounds, with a range of three miles and a flight time of approximately 30 minutes.

These specifications were very important when you are building the payload. Which includes what you see here. The data logger, a radiation detector. And in this course we had two different radiation detectors and a LiDAR system.

I also want to point out that the customized 3-D printing castings were used to attach the components securely in the figuration to maintain the center of gravity when flying.

Next slide please. So what you see here are the button or check sources used in the proof of concept were cobalt-60, cesium-137 and americium-241. The activities ranged from about 3.5 to 40 microcuries.

These radionuclides represent some, but not all, the major radionuclides found in decommissioning sites. Furthermore, these check sources are not representative of residual radioactivity in soil. Which would be ideal, but time and budget did not allow for that.

And if you remember, this is a proof of concept. So it's really to distinguish between humans and drones, not necessarily between a point source and an area source.

The flight path you see here has two parallel survey transits. One for radiological sources and of course one for background. They were marked with lumber to provide a visual pathway for the human surveyor, and for the drone pallet to follow. And the sources were randomly placed six meters apart on the lumber.

Next slide please. So a total of 12 radiological surveys were conducted. Four were conducted by humans and eight surveys were conducted by the UAV, or drone.

Of the eight drone surveys, the four surveys with the best performance were analyzed further. And the surveys were environmental conditions, for example, wind, all to the flight were excluded from further analysis. Also, each survey had a high or low altitude. The height, in centimeters, is shown here. And an average velocity of around .2 meters per second.

As a result, there were four major scenarios analyzed. For example, scenario three is just the first drone survey compared to the second human survey.

Next slide please. Now prior to conducting the radiological surveys the drones went through a series of pre-survey air worthiness, if you will. And this was required by PNNL and the Department of Energy.

These tests included a center of gravity determination, which include adjusting attachments to the drone such that the weight of the components were equal around the center of gravity. Other tests were for vibrational and structural integrity.

And the navigational systems evaluations were conducted with both a surrogate payload and without a payload. These were hot tests, a navigational test and a low altitude test.

Next slide please. So the radiation detector and data loggers were characterized in an indoor lab facility.

This included testing various functionalities, including of course, but not limited to just testing the cables to detectors to ensure that they were working, testing the loggers to laptops. And these steps were repeated several times to establish the robustness of the detector and data logger.

After the successful characterization and the indoor locations, field testing was conducted. Data was collected for a period of 15 minutes at each location using both detectors at 20 and 50 centimeters above the test, the test ground. And the average background count rate, in terms of TPMs, at least at each location, is presented in the table.

This picture on your left, I think if you are looking at your screen, represents a drone flying to collect activity over the background flight path.

Next slide please. So to keep all instruments consistent between the drone and human surveys, a field survey cart was used for the human survey. This allowed the UAV to be placed on the cart so that the same instrumentation for measuring location, the GPS, could be used for all surveys.

The cart was modified to include a adjustable cantilever to extent the detector and

LiDAR over the survey area to present distances above the ground. This picture shows a modified cart used to mimic radiological survey technics conducted by humans for comparison with the drone radiological surveys.

Next slide please. And this is just a close-up look configured with the radiation detected payload. You can see that at the time the photo was taken an activity of 1.6 counts per minute, which is the typical background activity around this facility was detected. And can be seen on the data logger.

Next slide please. So, let's move on, with caution, because the drone flights were in November, the analysis in December and the report, including data and results, are in draft form at this time.

Nevertheless, earlier in the presentation I discussed the three main questions. Question one was, did observe drone paths differ from human paths, and if so, how much?

This question addresses any deviation between the drone and human survey paths as a function of the GPS coordinates. This question was answered by estimating regression parameters. Specifically the slope of a regression line.

In this context, the slope of a path is defined as the upward movement along a path per one meter of lateral deviation. As a result, for all surveys for a lateral movement of one meter, the upward movement along a path was between 11 and 12 meters.

The path for scenario four had the greatest absolute slope difference of 1.9, while the path of scenario one had the smallest absolute difference of 0.36.

But the, what the takeaway is that the maximum difference between the human and the drone over a source location is around 4.5 centimeters.

Next slide please. Question two was, did survey path deviations impact survey results? This question addresses any deviation between the drone and human survey path as a function of altitude.

This question was quantitatively evaluated by examining the altitude data and to determine if the detector had the sensitivity to record reliable radiological measurements at the altitude for both the human and drone surveys.

Based on the minimum detectable

concentrations that were determined before the flight, the activity level measurements of some sources would be difficult to detect above 100 centimeters for this particular scenario. But overall, the survey path deviations were not statistically significant.

I want to point out a correction to this slide. The velocity variation said that it's between 1.8 and 2.6 meters per second. But that should be .18 and .26 meters per second. And the target, if you recall, is 0.2. I'm sorry, the velocity is 0.2 meters per second. Excuse me.

Question three was, were radiological measurements from human and drone surveys significantly different? This table here brings together questions one, two and three to summarize and answer this question.

If the drone and human surveys are the same, taken into account velocity, height, path difference, their statistical distributions will be statistically equivalent. For this project an alpha error of 0.05 was selected for statistical testing. At a peak value less than this number results in the findings that the two distributions are statistically different. Because all peak values were greater than the alpha error, there are no significant differences in the drone and human surveys using this radiological detector.

Next slide please. Now because we are still analyzing results, the conclusions are inconclusive at the moment. However, one limitation of the study was a limited scope and schedule that did not accommodate additional analysis for precision, accuracy and reproducible data.

Additionally, the background trends six, the background pathways were not analyzed in there. Also, what you considered on a check source by check source basis, some of the drawing in human path segments exceeded the GPS uncertainty. Nevertheless, the preliminary statistical analysis results indicates that drones and human paths are statistically equivalent.

From these direct exact results, the detectors worked well, the equipment is commercially available to perform surveys and the drone does have potential in decommissioning.

Next slide please. Okay, that was the overall conclusions. Can you go to the next slide?

So recommendations for next steps. Recommendations for next steps, the next phase of this project includes optimizing the radiation detection instruments, surveying complicated terrain, including subsurfaces, considering other radionuclides found at decommissioning sites, such as strontium-90 and tritium, and addressing the impact of environmental conditions like wind and moisture.

And we are looking at numerous ways to team up and collaborate with other future-focused research. For example, the first project mentioned in this RIC session, digital twins, shows promise in decommission space.

However, we are looking to expand the use of drones to reach beyond decommissioning as we see promise in reactor operation applications. Including unmanned radiation surveys, emergency response applications, security assessments and a number of applications where humans are normally used.

Next slide please. Finally, I want to recognize the project team listed here. I want to give an extra shout out, again, to PNNL for recruiting a number of subject matter experts, including statisticians, aviation personnel, including a pilot, payload and safety specialists, health physicists and a number of decommissioning specialists.

I also want to thank Cynthia Barr and Boby Eid in NMSS for being champions. Their report on this flight and data analysis should be out in a few months.

Thank you. And that is the end of my presentation.

MR. FURSTENAU: All right, thank you, Stephanie. I think we have time for one quick question, Stephanie.

In the work that you have doing so far, do you think drone usage offers, or can offer, some accuracy benefits compared to human surveys?

MS. BUSH-GODDARD: So, definitely some accuracy benefits. But you know, in our area, radiation protection, one of the big philosophies is ALARA. As low as reasonable achievable.

And that means it's reasonable to the human. So if we can substitute drones for humans, we are eliminating that radiation dose altogether.

And I just think that's a big health benefit. Although doses, we regulate the safe uses of all our facilities so doses are low. But it's just a benefit because if we can put in drones than we can eliminate the human.

And as far as the statistical accuracy, the equations have to put in some type of human sensitivity and some kind of human error. And I think if we can take that variable out then we can really move toward a more accurate calculation actually. A measurement, a radiological measurement. I hope that helps.

MR. FURSTENAU: Yes. Yes, thanks, Stephanie. That's a good answer. I think we're going to have to close your presentation now and go on to the next one, but thanks so much. This seems like a really exciting, and probably enjoyable project as well.

MS. BUSH-GODDARD: Yes. Yes, it was. Thank you.

MR. FURSTENAU: So, all right, thanks.

MS. BUSH-GODDARD: Thank you.

MR. FURSTENAU: So we'll go on to the fourth FFR project that we were going to highlight today. It's on the review of advance manufacturing technologies, or AMTs for fusion reactor material.

And our speaker today will be Amy Hull, but she'll also be accompanied by Shah Malik for the question, question and answer sessions.

Dr. Amy Hull is a senior materials engineer, again, in the Office of Nuclear Regulatory Research in the Division of Engineering. And most of her career was spent at Argonne National Laboratory with publications on nuclear power plant aging and licensing renewal, reactor materials, transportation on hazardous materials, advance engineered materials and metals, fusion reactor materials and high level radioactive waste disposal.

She holds a B.S. degree from Iowa State University and a PhD from Northwestern.

And for Shah Malik, Shah is also a senior materials engineer in the Office of Nuclear Regulatory Research. And most of his career was spent at the NRC working on nuclear power plant aging and reactor materials and reactor pressure vessel and primary system piping.

He previously worked at the General Electric on high temperature materials, structural integrity for jet engines. He holds a B.S. and M.S. from India and a PhD from the Ohio State University.

So, thanks for joining us, Amy and Shah. And, Amy, I'll turn the presentation over to you. Thanks.

DR. HULL: Thank you, Ray. Developing materials in advance manufacturing technologies is required to withstand the extreme environment of a fusion reactor. The U.S. National Academy has considered this as one of the three main needs to establish the scientific and technical basis for a fusion pilot plant by the 2040s.

The Nuclear Energy Innovation and Modernization Act, NEMA, requires NRC to develop a technology inclusive regulatory framework that encourages greater technological innovation for advance nuclear reactors, including fusion reactors. A just released NRC SECY-22-008 documents the work that NRC is undertaking to address regulatory pathways for potential fusion energy systems.

Slide 2 please. As shown in this slide, we focus on the nexus between advanced manufacturing technologies known at NRC by their acronym AMTs, and fusion reactor materials, particularly for plasma-facing components, usually referred to as PFCs.

The qualification of the materials and

components that surround the plasma and are exposed to fusion irradiation is still a significant challenge. The United States is heavily invested in these PFC materials, but to date, no credible PFC engineering solution has been demonstrated to meet requirements for fusion power plants.

The NRC staff have studied various AMTs for the present, and soon to be licensed, nuclear power plants and non-power facilities. The NRC needs to be prepared to review both the materials and AMTs proposed for license applications for fusion energy systems and ensure that the regulatory infrastructure would support, understanding the viability of a specific application.

Slide 3 please. There are three main fusion reactor approaches. Magnetic confinement, magneto-inertial and inertial fusion.

Our main focus is magnetic-confinement fusion energy systems, such as the international thermal nuclear experimental reactor commonly known as ITER, and more generally, tokamaks.

Different fusion energy systems will have different operating conditions and thus different values for heat load, radiation load and temperature gradients. Those given here were postulated for ITER.

As mentioned in the last bullet of this slide, PFC candidate materials included silicon carbide, molybdenum, niobium, tungsten, beryllium and graphite.

Slide 4 please. ITER is being built at Cadarache, France through an international partnerships. It is scheduled to start operations in 2025, although it won't be fueled with the power producing isotope tritium until 2035.

The ITER divertor region, mentioned in the last bullet on the previous slide, and shown in the right-hand schematic, and further on the next slide, also includes copper and stainless steel structure and heat removal tubing.

Novel materials and AMTs offer a potentially transformational advance for fusion energy by improving fusion plasma and engineering subsystems. Development of capabilities to better address risk and challenges associate with these advances would enable NRC to have a better awareness of fusion reactor materials when reviewing applications for fusion energy systems.

As an example of AMTs being considered for fusion reactor materials, Czech researchers are exploring cold spray deposition of thick tungsten and tungsten based armor that could be produced directly on the first wall surface.

Cold spray technology is an AMT option that NRC has been investigating recently with researchers from Pacific Northwest National Laboratory. The figure on the left is a schematic of a tokamak fusion reactor.

For dimensional perspective of this schematic, the red central solenoid will be about 13 meters high in ITER. The figure on the right shows a cross section of the ITER vacuum vessel and the plasma-facing components.

Tungsten was a candidate to be applied in the lower part of the reactor vessel, the divertor region, while tungsten alloys were considered for the remaining surfaces of the first wall. In ITER, due to its unique physical properties, such as low plasma contamination, low fuel retention, beryllium has been chosen as the element to cover the first wall, as shown here.

Slide 5 please. Every two years the

international conference on fusion reactor materials is held. This past October a researcher from Japan gave an invited paper examining radiation damage in tungsten alloys, which must withstand very harsh conditions. And thus high thermal conductivity and good mechanical properties are needed.

As mentioned earlier, tungsten is a plasma-facing material candidate because of its high melting temperature, high resistance to sputtering and high thermal conductivity.

This schematic illustrates the tokamak blanket cooling channels in the first wall and divertor block. The blanket, depicted in the lower left, is a modular component consisting of a shield block for neutron shielding and the plasma-facing detachable first wall panel that directly faces the plasma and removes the plasma heat load.

The purpose of the divertor is to extract the power coming from the conductive/convective heat flux and radiation while maintaining the plasma purity. It needs to tolerate high heat loads as the main interfacing components between the plasma and material surfaces, while at the same time providing neutron shielding for the vacuum vessel and super conducting magnets in the vicinity of the divertor.

Slide 6, please. Arranged in a circle at the bottom of the vacuum vessel, the divertor is made up of 54 cassette assemblies. Each one is formed from an actively cooled structural backbone cassette body each weighing about five metric tons in austenitic steel and copper alloy and plasma-facing elements covered in tungsten tiles. The full scale divertor dome assembly prototype, Number 5 in this schematic, was delivered to ITER in December 2021.

This followed a multi-year manufacturing, testing, and qualification program at the Efremov Institute. At the ITER site in France it will enter the first divertor integration trials where full-scale prototypes of all divertor components produced by Europe, Japan, and Russia will be assembled for the first time.

Simulations by researchers at Purdue University indicate that there may still be challenges, including unexpected and significant heat loads and damage to sensitive components other than the expected damage to the original divertor plate and first walls. Thus, we are looking at AMT alternatives and solutions for first wall or divertor regions.

Slide 7, please. During the past four months since the inception of this part-time project, we have been working to identify the technical challenges associated with novel advanced materials for fabricating components, especially plasma-facing components, in the context of NRC's current and evolving regulatory framework.

As an example of codes and standards activities we have participated in ASME Boiler and Pressure Vessel Section 3, Division 4, fusion related meetings. We are reviewing the draft standard, ASME FE.1-2022, Rules for Construction of Fusion Energy Devices.

We attended the 6-day long 20th international conference on fusion reactor materials and reported back information concerning the state of art for fusion reactor materials and specialized AMTs and challenges to their deployment.

Since the ITER first plasma is scheduled for 2025 followed by progressive ramp-up of the machine, it would be good to review the ITER databases while still available online from different contributing countries to anticipate any safety challenges before start-up.

This materials and manufacturing safety review must be done as soon as possible since some detailed information may no longer be available after 2025.

Slide 8, please. We are reviewing ongoing research, such as that sponsored by DOE, addressing silicon carbide, bimetallic joints, and refractory metals.

As an example of some of the exciting American work underway, the Advanced Research Projects Agency - Energy, ARPA-E, and DOE Fusion Energy Sciences have recently joined forces to prioritize R&D that helps establish viability of fusion-enabling technologies, including novel fusion materials and advanced manufacturing.

As alluded to in the third bullet, an Oak Ridge team will fabricate tungsten plasma-facing components using electron beam powder bed fusion and use in situ process monitoring for optimizing process structure relationships.

Silicon carbide materials are promising candidates for functional components in fusion reactors. Work is underway exploring different AM fabrication techniques, such as binder jetting.

The second bullet alludes to work done at Efremov Institute under ITER cooperation presented last October at ICFRM-20.

Slide 9, please. Our key planned activities will be integrated and refined to also be complementary and supportive of work done elsewhere at NRC.

As researchers our primary focus is to identify the technical challenges for fusion reactor materials and to investigate applications of AMTs to potential materials selected for different components.

We will prepare a technical assessment report which details the technical challenges and potential regulatory gaps related to using and licensing AMT fusion reactor materials for nuclear applications.

Also of interest are upcoming public meetings. NRC hopes to have a fusion public meeting on March 23, 2022. The White House will have a 2-hour summit on March 17 that will be live streamed at YouTube "developing a bold decadal vision for commercial fusion energy."

Slide 10, please. To summarize, Dr. Malik and I are investigating AMTs that may allow overcoming design restrictions and exploit new problem solving strategies.

Our focus is on exploring innovative materials and manufacturing technologies to meet material needs for fusion energy systems. A goal is to develop a gap analysis of material requirements versus risk in the context of NRC's current and evolving regulatory framework as needed.

We wanted to close with words of inspiration from the White House. "Fusion, the same reaction that powers the sun has the potential to be a game-changing technology to help us achieve net zero by 2050, protect national security, and enhance U.S. technology leadership.

Decades of public investment, rapid growth and private investment, and major recent scientific advances suggest that now is the time to quickly move towards demonstrating commercial fusion energy."

As part of the Office of Nuclear Regulatory Research we will help proactively provide the technical bases for technical flexibility for future Commission direction on the appropriate approach for licensing and regulating specific fusion energy systems. Thank you for your attention.

MR. FURSTENAU: Thank you, Amy. We have time for just one quick question because we are a little bit behind, but maybe you and Shah can take this question.

Technological advances like high temperature super conducting magnets suggest that commercial fusion power plants are likely to be far smaller than ITER.

Does NRC's analysis on this project apply equally to the smaller scale devices with smaller fuel requirements and smaller overall machines? Is different analysis required for these smaller devices?

DR. HULL: I will defer that to Shah. But since this project is just getting off the ground and since ITER is having problems, I really think it would behoove us to look more generally beyond ITER because that's very expensive.

It's not going to be really functional till 2035 and it has inherent difficulties from my perspective. Shah? Maybe I am speaking out of turn.

DR. MALIK: Yes, I think that is correct, and there are some smaller companies, both in U.S. and Canada and England, that are working on a smaller reactor both for the magnetic fusion and inertial fusion and magneto-inertial fusion.

And we will see who wins the race, so to speak, and comes up with a viable plant, a smaller plant, it will be smaller, within this intervening period between 2022 through 2035 or something like that.

There are companies in the U.S. as well as Canada that are really actively participating and they are also including EPRI, or Electric Power Research Institute, which is an international organization of affinities (phonetic) and they are cooperating along with that in that area.

MR. FURSTENAU: All right. Well, thank you, and thank you both for an exciting project as we look ahead to new technologies. Thank you both.

DR. MALIK: Thank you very much.

MR. FURSTENAU: Our final project is a Zero Trust Paradigm for Cyber Security in New Reactors and the presentation is going to be by Anya Kim and during the Q&A Kim Lawson-Jenkins will also join the stage.

Dr. Kim joined the NRC in 2020 as a computer scientist in the Instrumentation Controls and Electrical Engineering Branch.

Shortly after she joined us she kick-started several cybersecurity research projects such as characterizing the attack surface of digital assets in nuclear power plants, investigating novel technology implementations in nuclear power plants, and developing a zero trust paradigm for the nuclear industry, I think which you will hear about here in a few minutes.

Dr. Kim has 20 plus years of experience in cybersecurity. For 19 years she was at the U.S. Naval Research Lab working on research and development in cyber defenses technology for the Navy.

Some of her work includes developing a cyber event prioritization algorithm for Navy defense systems, creating the NRL security oncology and developing the security architecture for a vehicle tracking system. She holds a Ph.D. in cybersecurity from George Washington University.

Kim Lawson-Jenkins is an information

technology specialist in the Cyber Security Branch in the Office of Nuclear Security and Incidence Response.

So it's good to see we have two offices working together on this. That's a good thing to do with these future-focused research projects.

Kim has been with NRC since 2014 and prior to that she worked ten years as a computer scientist in the Center for High Assurance Computing Systems at the U.S. Naval Research Lab in Washington, D.C.

Kim received her bachelor's degree in computer science from the University of Illinois at Urbana-Champaign and a master's in applied sciences degree and computer science from George Washington University here in Washington, D.C.

So with that, Anya, I will turn it over to you for the presentation.

DR. KIM: Thank you, Ray. Good afternoon, everyone. Seeing as we are short on time I am going to jump right in. Can we go to the next slide, please? Thank you.

Okay, so first I would like to walk you through or give you a background on why we chose to look at zero trust for the nuclear industry. In fact, isn't this timing auspicious?

Because if you have been paying attention, just over a month ago the White House announced its zero trust architecture strategy for the federal government.

So zero trust has been getting a lot of publicity, but there is still a lot of misunderstanding about what it is, so it's worthwhile to spend some time going over the concept.

Then I will finally present how we plan to apply zero trust to the nuclear industry, in particular for new and advanced reactors.

Next slide, please. Okay, so let's look at the current state of the art first. Some of you may be familiar with this. For those of you who aren't this picture shows a logical model of a cybersecurity architecture in use at nuclear power plants today, or at least a version of it.

It's a flat network and it is segmented by security levels. This perimeter-based defensive architecture includes five concentric cybersecurity defense levels.

Each are separated by boundary devices,

such as firewalls or data diodes, and at each interface the digital communications are monitored and restricted.

Systems that require the greatest degree of security are located within the most secure levels, the innermost level of the defense architecture, which in this case would be Levels 3 and 4.

And this is typical of most networks today, whether it's IT or OT. So despite having access control applied or enforced at each security level, once you are inside that level you are considered trusted and so once you are in that level moving through a network or security level is very easy, especially if you are an attacker.

In cybersecurity that concept is called lateral movement. So lateral movement becomes easy. Could we move to the next slide, please?

So is that kind of perimeter-based defensive architecture we just saw, is that successful? For the most part yes, especially when you are talking about commercial U.S. nuclear power plants.

So if you are a U.S. nuclear plant

licensee go ahead and give yourself a pat on the back. However, whether you are a cybersecurity person or not you have probably heard of Stuxnet, you have heard of SolarWinds, and more recently log4j (phonetic), okay.

Despite perimeter-based defenses similar to the one we just saw, those attacks and the associated malware have found ways to successfully penetrate that perimeter and cause damage and sometimes remain undetected for a long period of time.

So, you know, supply chain vulnerabilities, insider threats aside, the one problem with this perimeter-based defense is that once the attacker gets inside that lateral movement I just talked about gets much more easy, okay.

And then on top of that we have new and advanced reactors that propose to use various technologies, such as artificial intelligence, AI, wireless and remote operations.

And, by the way, I would like to state for the record that the technologies depicted here are just examples of technologies that are proposed by the industry and the fact that they are here is in no way, shape, or form an endorsement of these technologies for use, okay.

So, okay, as I was saying, these new and advanced reactors are using these various technologies and they do not fit neatly within the traditional physical security and regulatory framework.

In fact, advanced reactors may rely less on physical security, putting more of an emphasis on the cybersecurity aspect.

So given all this, new threats, new technologies, new reactor designs, what was once working for us may not work in the future so we need a new way of thinking about security.

Next slide, please. So obviously that new way that we are thinking of is zero trust. So then what is zero trust?

Okay, first off let's start with what it isn't. It isn't a solution. It isn't a product that you can buy. It's not a specific technology. It's not even a set of technologies, okay.

There is no one size fits all magic bullet that you can pull off the shelf and neatly tie up with a ribbon and use.

Rather, zero trust is a strategy or a philosophy with a set of guiding principles, assertions, tenets, whatever you like to call them, and then actually, well, depending on what organization you ask, they state their principles or tenets a little bit differently.

But despite that, despite how they list these guidelines or tenets, the underlying concept is the same, okay. The zero trust model assumes that a breach is inevitable or it has already occurred, okay.

It's what I like to call a paranoia, despite the unsavory connotations with the word "paranoia." In fact, depending on which side of the fence you are on it's either paranoia or you are facing reality, okay.

Regardless, when you assume your network is hostile, okay, and a breach is going to occur, you realize there is no safe area inside the perimeter. So what do you do? You won't implicitly trust anymore, okay.

Trust is always going to be explicitly granted and dynamic, okay, requiring every user, every device, ever request, every data flow to be authenticated, authorized, and inspected for abnormal behavior or activity.

Okay, let me stress that again. Trust, explicit, and dynamic, okay. So even after that, even after access has been granted, that you only grant the least amount of privilege needed for them to get their job done, and that would limit the lateral movement that we talked about of the attacker.

At that same time you continuously monitor the network to detect threats, look for anomalous or malicious, okay. This paranoid mindset, that's what zero trust is about, okay.

And how you plan to address this paranoia is a separate issue, okay. So while buzz words that you may have heard like multi-factor authentication, virtual LAN, micro-segmentation, they are technologies that support your zero trust strategy, but they themselves are not zero trust, and just because you use them it doesn't mean you implement zero trust.

So my bank requires me to present my PIN and my ATM card when I attempt to check my balance, okay. That's an example of multi-factor authentication because I am using two things, a PIN and an ATM card, and sometimes even a driver's license, but does that mean the bank can claim they employ a zero trust security model? I don't think so.

Next slide, please. Okay, so this is a diagram or a figure from the NIST zero trust architecture document and it depicts the logical components of a zero trust architecture.

Obviously not all architectures have to look like this, and I definitely won't go over all of these components here, but I just wanted to relate these components to what we just talked about, okay.

So in the very bottom you see there is a user who is using their device to access a resource and in the middle there is a policy enforcement point.

So in zero trust you cannot, the policy enforcement point will enforce any access decisions that are made and you can't even know the existence of a resource without going through this policy enforcement point.

And remember we said that with the paranoid mindset, again, we require every device, every user, every transaction, every flow has to be authenticated and authorized and inspected.

So that's where the policy decision point above that comes in. The policy decision point is responsible for the ultimate decision to grant access or not to a resource given a subject against a dynamic policy.

It makes that decision based on a lot of information, okay. And so the two, the bunches, the eight squares on the side are potential sources of information that would be used to make this decision, for example, the CDM on the top left on your side would be the continuous diagnostics and mitigation system.

That gathers and provides information about the enterprise asset's current state, like what kind of patches have been applied to an operating system, what kind of vulnerabilities exist on the system, essentially what we would call the cybersecurity posture of that system would be maintained there and then the information would be sent.

The policy decision point would also obtain information about newly-discovered attacks for vulnerabilities from threat intelligence feeds that could be external to the network, okay.

And then the SIEM, on the low bottom right side, Security Information and Event Management System, that collects and provides security-censored information in a uniform manner that the policy decision point can use.

Also, data access policies can be used to make least privileged assessments or determinations. So, in other words, to make proper authorization decisions you really need to know what you have, like the network visibility asset identification, what their current state is, and then continuously monitor and log, support least privileged strategy, have threat detection data analytics, et cetera.

All of that is needed to determine whether in this current state with this particular request you can grant the access or not.

Next slide, please. So now we have a good idea of what zero trust is I hope and we know that the federal government is adopting it for its IT networks. But can we use this in the nuclear industry?

So while we believe, Kim and I believe, that a zero trust paradigm is a potential strategy or alternative that can be applied to protect reactors, particularly new and advanced reactors, there are issues that need to be researched.

The characteristics of a nuclear industry and, in fact, the other industrial control systems as well, need to be considered and reflected in any zero trust model that we adopt.

In particular, the intersection of safety and security and how to maintain that needs to be considered. The zero trust assertions and concepts in an IT environment may not be applicable as is in the nuclear industry.

So those are the things that we have to consider if we decide to apply this. And, furthermore, if this paradigm does work for the nuclear industry and nuclear power plants, then to adopt this new paradigm what do we need, we need applicable guidance, right.

So this research will also provide the basis for any future guidance in this area as well as performance criteria against which to measure from.

Next slide, please. So like Ray said, this is a newly-started future-focused research. We don't actually have any results. We are currently surveying the zero trust landscape.

And, obviously, most of the current body of work knowledge out there focuses on IT networks. So we are keeping that in mind as we do our investigations and identifying any challenges and gaps that may exist for our use.

And based on that we plan to develop a zero trust framework that is suitable for the nuclear industry where we look at concepts, principles, implementation strategies, and so on.

We are running out of time. So I also want to mention that of the principles I mentioned earlier there is one I stressed particularly, and that was the concept of explicit and dynamic trust.

That one, unlike other ones, such as least privileged or knowing what you have, continuous monitoring, is a little bit different. So Kim and I really would like to look at that.

That one requires using multiple sources to compute a trust score based on your acceptable level of risk.

What we want to do is develop or look at what kind of data analytics are needed to compute that trust score, what kind of data, how to collect it, how to measure it and combine, and so on and so forth.

Next slide, please. So this picture, I just wanted to show you where that trust algorithm, the component is usually called the trust engine, belongs.

So while I implied that all this information goes into the policy decision point, it actually goes into that trust algorithm that takes all this information to compute the trust score and then that, along with the request, is what the policy decision point would use to make decisions.

Next slide, please. The expected results and benefits are that we hope that this would form the basis for any future regulatory guidance documents and provide an alternative way to have as a defensive architecture when applied to new reactors.

Also, what we learned here would be useful for any industry or control system building or a safety critical industry in general as well as providing like a common ground for a discussion among staff, licensees, applicants, and vendors and inspectors regarding the zero trust paradigm.

Next slide, please. So thank you very much. If you have any questions, please let me know.

MR. FURSTENAU: Yes, thank you, Anya. I do have some good questions coming in. I know we are past time, but I think if everybody can stick around for just a few more minutes I would like to ask you guys some questions.

The first question is what are your thoughts on distributive ledger technology and quantum communications? Can these be applied?

MS. LAWSON-JENKINS: This is Kim Lawson-Jenkins. It's probably going to be one of the technologies that we will possibly be looking at in the future as a way of implementing these things.

It is not something that is currently covered in -- unlike a lot of the controls that are normally used, those specific things about blockchain, for instance, those kind of things aren't mentioned in cybersecurity plans, okay.

So that will probably require some new guidance that is outside of the current set of security controls we are using, okay, but that will probably be one of the things we'll be looking at for the future, okay.

There is new Part 53 work that is going on right now and that no doubt will probably be one of the things that, one of the controls or a set of controls that they will be looking at in the future guidance.

MR. FURSTENAU: Okay. Kim, one more question before we close the session, because it's a really interesting topic.

The zero trust concept is interesting, but does this introduce a single point of vulnerability on the architecture if the policy enforcement point is compromised? How would the NRC regulate such an approach? So who would like to take that question?

MS. LAWSON-JENKINS: I don't think it's going to be an issue if the policy decision point is compromised because there are, for better or worse there are going to be multiple ways of getting into a system, just as you have today.

So it's really in the end the decision point that's going to be making the decision on how you are going to enforce a policy. So I honestly can't see a single point of failure because this is going to be more of a philosophy and a strategy rather than a device.

As Anya said, this is not a device that is going to be installed that is going to do everything. So I don't think a single point of failure will be applicable for this type of architecture.

MR. FURSTENAU: Okay.

DR. KIM: Could I add to that just a little bit, do we have time?

MR. FURSTENAU: Sure.

DR. KIM: Okay. So, yes, the picture was just a logical picture when we were looking at the NIST picture, but it is very simple.

I didn't go into detail, but there is a data plan and a control plan and it's very similar to software-defined networks.

So I think the person who asked the question is probably familiar with that because that's how it is divided and in software-defined networks there is that single point of failure that goes where -- And there are many attacks where they try to take out I guess the control plane on top, in our case it would be the policy decision point.

But there are ways architecturally to get around it and you have multiple policy enforcement points and you have redundant and resilient policy decision points, and that's just one way.

So what Kim said, totally, because, you know, it's not an actual device. We don't even know what it's called this way. This is just conceptual. But even if it were to look this way there are ways to get around those kinds of vulnerabilities, but it's great that you pointed that out.

MR. FURSTENAU: Okay. Well, thank you. Thank you, Kim and Anya. I would like to thank all of the presenters. I wish we would have had more time for Q&A.

We could go on for at least another hour yet, but I think we really do need to close this session since I'll get in a bit of trouble if we don't.

So, again, thanks everybody and thanks for all of you who signed up and participated and asked questions today. It was really a fascinating, exciting, and enjoyable session I know for me.

So thanks again, everybody. With that, I will close the session. Thank you.

(Whereupon, the above-entitled matter

went off the record at 2:36 p.m.)